Utility-based Sensor Selection*

Fang Bian David Kempe Ramesh Govindan

Department of Computer Science
University of Southern California
Los Angeles, CA 90089
{bian, dkempe, ramesh}@usc.edu

ABSTRACT

Sensor networks consist of many small sensing devices that monitor an environment and communicate using wireless links. The lifetime of these networks is severely curtailed by the limited battery power of the sensors. One line of research in sensor network lifetime management has examined *sensor selection* techniques, in which applications judiciously choose which sensors' data should be retrieved and are worth the expended energy. In the past, many ad-hoc approaches for sensor selection have been proposed. In this paper, we argue that sensor selection should be based upon a trade-off between application-perceived benefit and energy consumption of the selected sensor set.

We propose a framework wherein the application can specify the utility of measuring data (nearly) concurrently at each set of sensors. The goal is then to select a sequence of sets to measure whose total utility is maximized, while not exceeding the available energy. Alternatively, we may look for the most cost-effective sensor set, maximizing the product of utility and system lifetime.

This approach is very generic, and permits us to model many applications of sensor networks. We proceed to study two important classes of utility functions: submodular and supermodular functions. We show that the optimum solution for submodular functions can be found in polynomial time, while optimizing the cost-effectiveness of supermodular functions is NP-hard. For a practically important subclass of supermodular functions, we present an LP-based solution if nodes can send for different amounts of time, and show that we can achieve an $O(\log n)$ approximation ratio if each node has to send for the same amount of time.

Finally, we study scenarios in which the quality of measurements is naturally expressed in terms of distances from targets. We show that the utility-based approach is analogous to a penalty-based approach in those scenarios, and present preliminary results on some practically important special cases.

*This material is based upon work supported by the National Science Foundation under Grants No. 0520235 and 0121778. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

IPSN'06, April 19–21, 2006, Nashville, Tennessee, USA.

Categories and Subject Descriptors: C.3 [Special-purpose and Application-based Systems]: Real-time and Embedded systems, F.2.1 [Analysis of Algorithms and Problem Complexity]: Optimization.

General Terms: Algorithms, Theory.

Keywords: Sensor networks, Utility, Sensor Selection, Complexity, Sub-modular, Super-modular.

1. INTRODUCTION

Sensor networks consist of many small sensing devices that monitor an environment and communicate using wireless links. The limited energy of sensor nodes curtails the lifetime of sensor network deployments. Therefore, a large body of literature has examined methods for extending sensor network lifetime by carefully managing communication and computation resources.

One thread has focused on techniques to maximize sensor network lifetime by using routing or topology control [1, 2, 3, 4, 5, 6, 7, 8]. Somewhat orthogonal to this thread is the approach of asking which and how much data can be collected over the lifetime of a network. Work on this question is usually centered around one of two paradigms: (1) Maximize the total amount of data collected, or (2) Collect data from all sensors for as long as possible. The first of these assumes that all data is equally interesting or important to the application. In particular, optimizing for the total amount of data will give undue preference to data measured close to a base station, as that data can be extracted at relatively low cost. On the other hand, the second objective is based on the assumption that data collection is only worthwhile if data from *all* sensors is being collected, and as soon as one node's energy is depleted, the network may as well not collect any data.

Clearly, both approaches are oversimplifications of reality. Which data is useful depends heavily upon the specific application and its needs. In many cases, individual sensor readings will be of little use; the appeal of sensor networking lies in the ability to aggregate and correlate sensor readings from different locations. In other scenarios, readings from a group of proximate sensors may be considered redundant, and it would suffice to obtain a reading from one of them.

In this paper, we argue that the algorithms for deciding which data to retrieve should be sufficiently generic to let the *application* specify how useful particular measurements are. We call the usefulness of a concurrently measured sensor set S its *utility* $u(S)$. Given a network topology and application-specified utility function $u(\cdot)$, it is then the algorithm's decision how to trade off the utility and energy consumption of sensor sets in an optimal way, and maximize the total utility extracted from the network until the network ceases to function. Such a *utility-based sensor selection* approach has also been proposed recently by Byers and Nasser [9].

Ideally, one would like to be able to find an (approximately) optimal sequence of sets and associated communication scheme to measure for arbitrary monotone utility functions. This goal seems very ambitious: as we show, the problem of selecting an optimal sequence of sets is NP-hard in many settings. In this paper, we explore three natural and practically important classes of utility functions in more detail. Specifically, we focus on submodular functions (with returns for additional sensors diminishing for larger sets), supermodular functions (with returns increasing for larger sets), and a general framework of geometric covering objectives. In the latter case specifically, we show that the utility-based approach is analogous to a *penalty-based* approach that characterizes the penalty of a sensor set as its collective distance from the targets to be measured.

We show that the optimum sequence of sets for submodular functions can be found in polynomial time, while optimizing the cost-effectiveness of supermodular functions is NP-hard. For a practically important subclass of supermodular functions, we present an LP-based solution if nodes can send for different amounts of time, and show that we can achieve an $O(\log n)$ approximation ratio if each node has to send for the same amount of time. Finally, for geometric covering objectives, we show that finding the best sensor set is NP-hard, and unless P=NP, the optimum solution cannot be approximated to within better than a multiplicative factor of 1.822.

The paper is structured as follows. Section 2 formally describes the utility-based sensor selection problem. Sections 3 and 4 then examine the feasibility of solving several variants of the utility-based sensor selection problem for submodular and supermodular functions. Section 5 discusses geometric coverage objectives. Finally, Section 6 discusses related work and Section 7 summarizes our contributions.

2. NETWORK AND PROBLEM SETUP

2.1 Network Model

Formally, our network model can be described as follows. The network is considered as a directed graph $G=(V,E)$ with $n=|V|$ nodes including a special root node $r \in V$. Each non-root node v has a finite *initial energy* E_v. If node v is to sense data, it incurs a *sensing cost* σ_v. To transmit data along the link $e=(v,w) \in E$, node v incurs a *transmission cost* of τ_e per unit of data. The total of sensing cost and transmission cost may not exceed a node's energy. When the two are equal, we say that the node's energy is *depleted*.

This network model is consistent with current practice in sensor network deployment. Most existing deployments, including the James Reserve habitat monitoring network [10], the Great Duck Island network [11, 12], and the Extreme Scaling network [13] are *tiered*: they consist of a large number of small battery-powered motes, sending data (perhaps after some local processing) to a smaller number of well endowed upper-tier 32-bit embedded nodes (*e.g.,* Stargates). The small-form-factor motes are most often battery equipped, and thus constrain the system lifetime, while the upper-tier nodes are much more powerful, and rarely impose additional constraints. For our purposes, this means that once data has reached the upper-tier, it can be extracted from the network. This allows us to consider all upper-tier nodes as one *base station* or *root*, for the purposes of theoretical treatment.

This network model is also consistent with at least one proposed sensor network architecture [14]. That architecture proposes to place application-specific functionality in the upper-tier nodes. Applications can generically *task* collections of motes. Each mote senses the environment, processes the sensed data as specified in a task, and returns the data to an upper-tier node. Indeed, our problem formulation is motivated by this architecture: applications may use our utility-based sensor selection framework to determine which set(s) of sensors to task.

2.2 Utility-Based Sensor Selection

We are now ready to characterize the optimization problem of utility-based sensor selection. In addition to the network with costs and initial energies for nodes, we are given a monotone *utility function* $u: 2^V \to \mathbb{R}^+$. The meaning of this function is that if all nodes of S are measured (nearly) concurrently, then the application derives a benefit of $u(S)$ from this measurement. The goal is to calculate a sequence of measurements to give the maximum total benefit until the sensor nodes' energy is depleted.

Utilities naturally capture application requirements in many practical sensor network scenarios. For example, in a beam-forming application, a non-collinear set of sensors has a higher utility than one that is collinear. In a surveillance application, a set of (approximately) evenly spaced sensors around the periphery of the sensed area has a higher utility than a clustered set of sensors in one corner of the sensed area. In a structural monitoring application, a set of sensors in which none are on the *nodes* (nulls in the structural mode shapes) has a higher utility than a set where some are.

More formally, a *feasible solution* consists of a collection $\mathscr{S} = \{S_1,\ldots,S_k\}$ of subsets of V to be measured, with a non-negative amount of data t_i to be transmitted from each set S_i. In addition, the solution must specify a flow f of value $\sum_{i:v\in S_i} t_i$ from each node v to the root node r, such that the total cost of flows and sensing does not exceed any node's energy. (This flow captures the delivery of all relevant data from node v to the root.) The goal is then to find a feasible solution of maximum total utility.

We can state this as a linear program (with exponentially many variables) as follows. For each set $S \subseteq V$, we have a variable t_S specifying the amount of data sent from S to the root. (Here, we denote by $\delta^+(v)$ and $\delta^-(v)$ the set of outgoing and incoming edges of v, respectively.)

$$
\begin{array}{llll}
\text{Maximize} & \sum_{S\subseteq V} u(S)\cdot t_S & & \\
\text{subject to} & \sum_{e\in\delta^+(v)} f_e - \sum_{e\in\delta^-(v)} f_e & = \sum_{S:v\in S} t_S & \forall v\neq r \\
 & \sum_{e\in\delta^+(v)} \tau_e f_e + \sum_{S:v\in S} \sigma_v t_S & \leq E_v & \forall v \neq r.
\end{array}
$$

The first constraint characterizes flow conservation: at every node, the total amount of flow leaving exceeds the amount entering exactly by the amount of data that the node itself is sending. The second constraint simply states that the total of transmission and sensing cost for no node exceeds the available energy.

Notice that the LP formulation does not contain a "time" component. Indeed, the order in which the sets transmit their data does not affect the total utility derived from a sequence. It is also not difficult to see that the flow f can be decomposed into different flows f^S for each sending set S, while keeping the total amount $\sum_{S\subseteq V} f^S$ transmitted along any link e the same as f_e. Hence, we are justified in considering the simplified formulation above.

We term the problem as formalized above the Utility-based Sensor Selection Problem (USSP). Notice that the formulation is very similar to that proposed by Byers and Nasser [9], but the focus of our work is significantly different from theirs (Section 6). Several natural variants of the problem may be of practical importance:

- In the Integral Utility-based Sensor Selection Problem (IUSSP), all data transmission amounts t_S are required to be integral. This variant is important in applications which require data collection in discrete data units (for example, [15] describes a scenario in which a structure's response needs to be sampled for no less than 1 second at high enough frequency in order to have enough samples to compute a FFT).

- In the Utility-based One Set Sensor Selection Problem (UOSSP), the amount of data is allowed to be positive for at most one set, *i.e.*, we are interested in collecting data only from one set. This case is of importance in determining the set with the best utility-to-energy tradeoff.
- Naturally, the previous two variants can be combined, by requiring that only a single set transmit data, and do so for an integral amount of time. We then obtain the Integral Utility-based One Set Sensor Selection Problem (IUOSSP).

Notice that even the LP for the basic USSP problem is unlikely to be solvable explicitly for arbitrary utility functions, due to its exponential size. Indeed, even specifying a utility function completely would require giving a value for each set S, and thus space exponential in the number n of sensors. These two observations motivate studying restricted classes of utility functions of practical importance. In the following sections, we investigate more closely submodular and supermodular functions, as well as utility functions applicable to geometric coverage settings.

3. SUBMODULAR UTILITY FUNCTIONS

In game theory, submodular functions are frequently studied as a natural restriction of utility functions [16]. Recall that a function $u: 2^V \to \mathbb{R}$ is *submodular* if $u(S_1 \cup S_2) + u(S_1 \cap S_2) \leq u(S_1) + u(S_2)$ for all sets S_1, S_2. The practical importance of submodular functions is more evident from an equivalent characterization: u is submodular if and only if $u(S_1+v) - u(S_1) \leq u(S_2+v) - u(S_2)$ for all sets $S_2 \subseteq S_1$ and elements $v \notin S_2$. This inequality captures the idea of "diminishing returns": the benefit of adding the element v to a set is non-increasing as the set that v is added to increases. In sensor networks, utility functions will be submodular if the data measured by different sensors is (partially) redundant; it will cease to be so if multiple measurements yield superlinear benefit.

For example, when sensors are deployed for measuring the average temperature, one natural notion of utility is the expected reduction in variance of the estimate of average temperature. Under some natural assumptions, this function is submodular. In order to maximize the system lifetime utility, one will then be trading off accuracy in the estimate against taking measurements for a longer time.

In this section, we show that USSP can be solved in polynomial time for arbitrary submodular functions. The key observation is that, without loss of generality, an optimal solution never measures two sensors at the same time. This observation then allows us to reduce the size of the linear program to polynomial, and solve it explicitly.

LEMMA 1. *Without loss of generality, the optimum solution to USSP only retrieves data from singleton sets,* i.e., $t_S = 0$ *for all non-singleton sets* S.

Proof. Let $((t_S),(f_e))$ be an optimal solution for USSP, and assume that the set $S = \{v_1, v_2, \ldots, v_k\}$ (with $k \geq 2$) transmits $t_S > 0$ amount of data. Consider a solution which sets $t'_S := 0$, and instead increases $t'_{\{v_i\}} := t_{\{v_i\}} + t_S$ for $i = 1, \ldots, k$. Each node v in the new solution still needs to transmit the same total amount of data as before, so f is still a valid flow, and the new $((t'_S),(f_e))$ is a feasible solution.

From the diminishing returns property, we can inductively derive that $u(S) \leq \sum_{i=1}^{k} u(\{v_i\})$, and as the $t_{S'}$ values for all other sets S' remain unchanged, the total utility of the solution changes by exactly $t_S \cdot (\sum_{i=1}^{k} u(\{v_i\}) - u(S)) \geq 0$. In particular, the new solution is at least as good as the original one. By continuing in this way, we eventually arrive at a feasible solution in which only singleton sets have non-zero amounts of data transmitted, and which has at least the same total utility. ■

Based on the observation in the lemma, we can remove all variables for non-singleton sets from the LP for USSP, arriving at the following modified LP:

$$\begin{array}{llcll} \text{Maximize} & \sum_{v \in V} u(\{v\}) \cdot t_v & & & \\ \text{subject to} & \sum_{e \in \delta^+(v)} f_e - \sum_{e \in \delta^-(v)} f_e & = & t_v & \forall v \neq r \\ & \sum_{e \in \delta^+(v)} \tau_e f_e + \sigma_v t_v & \leq & E_v & \forall v \neq r. \end{array}$$

This new LP has polynomial size, and can thus be solved explicitly in polynomial time. Notice also that it is enough for the application to specify the utilities of singleton sets. Thus, our observation allows us to circumvent the problem of needing exponential space to describe the utility function completely — a complete description is not necessary to find the optimum solution.

4. SUPERMODULAR UTILITY FUNCTIONS

Supermodularity in a sense captures the opposite of submodularity: the benefit of individual nodes is *larger* as they are added to an already larger set; more generally, the benefit from combining two (disjoint) sets is at least as large as the sum of the individual benefits. Formally, we say that a function $u: 2^V \to \mathbb{R}$ is *supermodular* if $u(S_1 \cup S_2) + u(S_1 \cap S_2) \geq u(S_1) + u(S_2)$ for all sets $S_1, S_2 \subseteq V$. In sensor network applications, supermodular functions will occur when the application requires the combination of diverse data measured (nearly) simultaneously.

For example, in an aquatic monitoring application, sensors of different types (temperature, light, chlorophyll, ...) may be deployed in an expanse of water. An application may be interested in a correlated measurement of all light sensors, a correlated measurement of all sensors at depth 10ft, and a measurement of all sensors that had measured a certain phenomenon the previous day. Naturally, there may be overlap between these sets: for instance, a light sensor may be at a depth of 10ft, and have measured the phenomenon the previous day. If the application is such that missing even one of the relevant sensors makes the measurement (nearly) useless, for instance because correlations are to be analyzed in detail, then the utility derived from a particular set S is simply the number of categories (light, 10ft, phenomenon) that are completely contained in S. This function is supermodular - in fact, it is a special case of the category of *set-weighted functions* studied in more detail below.

Notice that a function can be submodular and supermodular at the same time; this happens exactly when the function is of the form $u(S) = \sum_{v \in S} u(\{v\})$, *i.e.*, the utility of any set is simply the sum of utilities of the individual elements in the set. Though one can solve the same LP used for the sub-modular case here, notice that the resulting optimum singletons can be combined freely to get another optimum solution.

The property of supermodularity allows us to infer an interesting property of the optimal solution: without loss of generality, the sending sets are nested. In other words, if S_1, S_2 are sets sending data, then either $S_1 \subseteq S_2$, or $S_2 \subseteq S_1$.

LEMMA 2. *Without loss of generality, the optimum solution for USSP with a supermodular utility function consists of nested sets.*

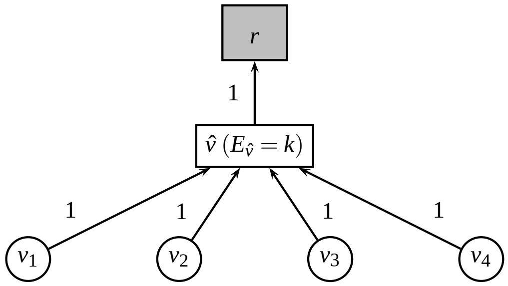

Figure 1: SETCOVER **graph with** $n = 4$

Proof. Let $((t_S),(f_e))$ be an optimal solution for USSP, and assume that there are sets S_1, S_2 such that neither $S_1 \subseteq S_2$ nor $S_2 \subseteq S_1$, yet $t_{S_1} > 0$ and $t_{S_2} > 0$. Without loss of generality, assume that $t_{S_1} > t_{S_2}$. Consider a solution which instead sets $t'_{S_2} = 0$, $t'_{S_1} = t_{S_1} - t_{S_2}$, and $t'_{S_1 \cup S_2} = t_{S_1 \cup S_2} + t_{S_2}$, as well as $t'_{S_1 \cap S_2} = t_{S_1 \cap S_2} + t_{S_2}$. In the new solution, each node still needs to transmit the same amount of data, so $((t'_S),(f_e))$ is still feasible.

Since each set other than $S_1, S_2, S_1 \cup S_2$, and $S_1 \cap S_2$ still transmits the same amount of data, the change in the total system utility is

$$\begin{aligned} & u(S_1 \cup S_2) \cdot t_{S_2} + u(S_1 \cap S_2) \cdot t_{S_2} - u(S_1) \cdot t_{S_2} - u(S_2) \cdot t_{S_2} \\ = \quad & t_{S_2} \cdot (u(S_1 \cup S_2) + u(S_1 \cap S_2) - u(S_1) - u(S_2)) \\ \geq \quad & 0, \end{aligned}$$

by the supermodularity of u. This reduces the number of pairs (S_1, S_2) of non-nested sets by one, as it can be easily verified that whenever $S_1 \cup S_2$ is not nested with some set S', then at least one of S_1, S_2 is not nested with S', either. Thus, we can continue this process inductively, and arrive at a nested solution of at least the same total utility. ■

While this lemma shows that the optimum solution gathers data from at most n sets, it unluckily does not restrict the class of sets that these could be drawn from, and thus, does not restrict the size of the LP. Hence, it does not suggest a polynomial-time algorithm. Indeed, we conjecture that USSP is NP-complete for supermodular functions. While we currently have no proof of this conjecture, we can show that the version requiring integral data amounts to be sent is not only NP-hard, but cannot be approximated unless P=NP.

THEOREM 3. *IUSSP with supermodular utility functions is NP-hard. Unless P=NP, it cannot be approximated to within any multiplicative factor.*

Proof. We will prove that it is NP-hard to decide if the system admits a solution with positive total utility. That proves both the NP-hardness and approximation hardness result. We reduce from the decision version of SETCOVER: Given a universe $U = \{1,\ldots,n\}$, and a collection of subsets $S_i \subseteq U$ for $i = 1,\ldots,m$, is there a collection of k sets S_{i_j} covering the entire universe, *i.e.*, such that $\bigcup_{j=1}^{k} S_{i_j} = U$.

Given an instance of SETCOVER, we construct an instance of IUSSP with a supermodular utility function. The graph has one node v_i for each set S_i, as well as a *bottleneck node* $\hat{v}$ and the root r. There is a directed link from each v_i to $\hat{v}$, and from $\hat{v}$ to r. All links e have transmission cost $\tau_e = 1$, and all sensing costs are 0. The initial energy of each v_i is 1, while the initial energy of $\hat{v}$ is k. The resulting graph for four sets is shown in Figure 1.

To complete the reduction, we need to specify the utility function. For any node set $T \subseteq \{v_1,\ldots,v_m\}$, we define the utility $u(T) := 2^{|T|}$ if $\bigcup_{i:v_i \in T} S_i = U$, and $u(T) := 0$ otherwise. This reduction can be computed in polynomial time. The utility function, although defining the value for 2^m subsets, can be specified in polynomial space (and time) by giving the collection of sets S_i. It is also not difficult to verify that u is in fact supermodular, by considering the three cases that neither T_1 nor T_2 defines a set cover, that exactly one of them does, and that both of them do.

We now claim that a positive total utility can be achieved in this setting if and only if the SETCOVER instance had a solution of size at most k. If $\{S_{i_j}\}$ is a set cover of size at most k, then it is easy to see that the node set $\{v_{i_j}\}$ leads to a feasible solution, of utility $2^k > 0$. Conversely, if $\{v_{i_j}\}$ is a node set contributing total utility to an IUSSP solution, we first observe that by feasibility, it has size at most k. Furthermore, since the utility is positive, $\{S_{i_j}\}$ must be a SETCOVER by the definition of u. Thus, we have established the claim, proving both hardness results. ■

Notice that the same hardness results carry over to the case when we are seeking a single set to send an integral amount of data, i.e., the integral version of UOSSP.

4.1 Set-weighted functions

While Theorem 3 shows a very strong hardness result for IUSSP with supermodular functions, the functions it uses in the reduction are arguably not very natural in the context of sensor networks. In particular, one would assume that most utility functions specified by applications have a more explicit compact representation. Here, we focus on a further restriction on supermodular functions.

We term the class of functions *set-weighted* utility functions. A set-weighted utility function is characterized by a collection $\mathscr{P}$ of sets, and for each set $P \in \mathscr{P}$ a non-negative weight w_P. The utility of a set S is then defined as $u(S) = \sum_{P \subseteq S, P \in \mathscr{P}} w_P$. If the collection $\mathscr{P}$ is small enough (for instance, polynomial in n), then it provides a natural and compact representation of u. Also notice that any function u defined in this way is always supermodular.

For set-weighted utility functions, we can in fact find the optimum solution to USSP by solving a linear program. The insight for this linear program is based on Lemma 2: the optimum solution for any supermodular utility function without loss of generality consists of a nested collection of sets. In particular, we can think of this collection as being ordered by decreasing size. Then, any solution can be fully characterized by specifying how much data each node sends, or, equivalently, how long each node sends. Suppose that up to some time t, all nodes in S still send data (or, the nodes in S each send t or more units of data). Then, these nodes up to time t contribute utility $t \cdot u(S) = t \cdot \sum_{P \subseteq S, P \in \mathscr{P}} w_P$. Based on this observation, and changing the order of summation in the objective function, we can rewrite the linear program for USSP as follows:

$$\begin{array}{lllll} \text{Maximize} & \sum_{P \in \mathscr{P}} w_P \cdot t_P & & & \\ \text{subject to} & t_P & \leq & s_v & \forall P \in \mathscr{P}, v \in P \\ & \sum_{e \in \delta^+(v)} f_e - \sum_{e \in \delta^-(v)} f_e & = & s_v & \forall v \neq r \\ & \sum_{e \in \delta^+(v)} \tau_e f_e + \sigma_v s_v & \leq & E_v & \forall v \neq r. \end{array}$$

Notice that we are introducing variables s_v for the amount of data sent by each node v. By expressing the utility function in terms of contributions by subsets, we managed to eliminate the terms for sets other than the $P \in \mathscr{P}$. By solving this polynomial-sized LP, we can find an optimal solution to USSP in polynomial time.

4.2 Set-weighted single set selection

In the context of understanding the solutions to the USSP problem, an important question is which single set is most cost-effective, *i.e.,* gives the largest utility for its energy consumption. In other words, which set S maximizes the product of the set's utility $u(S)$ and the lifetime of the system if only S sends data? This is the problem termed UOSSP above. Here, we show that even for the special case of set-weighted supermodular functions, the UOSSP problem is NP-hard.

THEOREM 4. *UOSSP with set-weighted supermodular functions is NP-hard.*

Proof. We will give a reduction from the densest k-subgraph problem [17, 18]: Given a graph $G=(V,E)$ and a number k, as well as a *density requirement* ρ, is there a subset $S \subseteq V$ of size at most k containing at least $\rho \cdot k$ edges. (Here, we say that S contains an edge e if it contains both of its endpoints.)

Given an instance of the densest k-subgraph problem, we construct an instance of UOSSP with a set-weighted utility function as follows: The graph $G'=(V',E')$ consists of a sink node r, two bottleneck nodes $\hat{v}_1, \hat{v}_2$, and a node v for each v of the given graph G. Each such node v is connected to the bottleneck node $\hat{v}_2$, and both bottleneck nodes are connected to the sink r. All sensing costs are 0, and all transmission costs are 1. All nodes have initial energy of 1, with the exception of node $\hat{v}_2$, which has initial energy of k.

The collection $\mathscr{P}$ contains each pair (u,v) of G' such that $(u,v) \in E$ is an edge of the given graph, and assigns them a weight of 1. In addition, the node $\hat{v}_1$ is in $\mathscr{P}$ as a singleton set, having very large weight $1+|E| \cdot k$. We claim that the given graph contains a k-subgraph of density at least ρ if and only if there is a node set to sense that gives total utility at least $1+k\cdot(|E|+\rho)$.

For the first direction, observe that if the set S contains at most k nodes and has density ρ, then sending one unit of data each from nodes $\hat{v}_1$ and from all nodes $v \in S$ gives the desired total utility. Conversely, we first observe that the maximum utility that can be obtained without including $\hat{v}_1$ is at most $k|E|$, as the nodes can send at most k units of data, and the utility of any such set of nodes is at most $|E|$. Thus, the optimum solution will always send one unit of data from $\hat{v}_1$, and is thus also forced to send at most one unit of data from each other node v. But then, it will always send data from exactly k nodes other than $\hat{v}_1$ (else, some of the energy of $\hat{v}_2$ would go unused). Since the total utility was assumed to be at least $1+k\cdot(|E|+\rho)$, the utility of the set S of nodes other than $\hat{v}_1$ must be at least $\rho \cdot k$. This in turn means that the set S contains at least $\rho \cdot k$ edges in G. This proves the correctness of the reduction, and thus the NP-hardness of the problem. ■

While the UOSSP problem for set-weighted supermodular functions is NP-hard to solve optimally, it can be approximated to within a factor of $O(\log n)$ within polynomial time, using an LP-rounding based approach. The algorithm first solves the corresponding linear program for USSP (which is also the natural linear program for UOSSP) in polynomial time, and then considers all candidate node sets C_x of the form $C_x = \{v \in V \mid s_v \geq x\}$. Among all such sets C_x, it chooses the one maximizing the product of $u(C_x)$ and the system lifetime if all of C_x sends. Notice that the algorithm only needs to look at such sets C_x for values $x = s_v$ for some v. Thus, at most n different sets need to be considered. For each set, calculating the system lifetime is done in turn by solving the corresponding flow LP.

The proof of the $O(\log n)$ approximation guarantee is based on the following useful lemma.

LEMMA 5. *If $x_1 \geq x_2 \geq x_3 \geq \ldots \geq x_k \geq 0$ are k real numbers with associated non-negative rational weights $w_1, w_2, \ldots, w_k$, then $\max_{i=1}^{n}(x_i \cdot \sum_{j=1}^{i} w_j) \geq \frac{\sum_{i=1}^{k} w_i x_i}{H_k}$, where $H_k = \sum_{i=1}^{k} 1/i$ denotes the k^{th} Harmonic number.*

Proof. To prove the claim, we first consider the case when all $w_i = 1$ (thus, $\sum_{j=1}^{i} w_j = i$). Let i^* be the index maximizing $x_i \cdot i$. Thus, we have that $x_i \leq (i^* \cdot x_{i^*})/i$, and

$$\frac{\sum_{i=1}^{k} x_i}{H_k} \quad \leq \quad i^* \cdot x_{i^*} \cdot \frac{\sum_{i=1}^{k} 1/i}{H_k} \quad = \quad i^* \cdot x_{i^*} \quad = \quad \max_i (i \cdot x_i).$$

If the w_i are not all equal to 1, then we first multiply all of them by the common denominator to make them integers — notice that multiplying all weights by the same constant does not affect the validity of the inequality to be proved. Then, if the number x_i has weight $w_i \in \mathbb{N}$, we replace it by w_i copies of the number x_i, each with weight 1. Notice that this leaves both the left-hand side and right-hand side in the inequality $\max_{i=1}^{k}(x_i \cdot \sum_{j=1}^{i} w_j) \geq \frac{\sum_{i=1}^{k} w_i x_i}{H_k}$ unchanged; for the right-hand side, this follows because $w_i x_i$ is replaced by $\sum_{j=1}^{w_i} x_i$, and for the left-hand side, it can be seen that the maximum after the replacement is attained at an index i' such that $x_{i'}$ is the w_i^{th} copy of the number x_i. ■

Using this lemma, we can state and prove the approximation guarantee for the LP-rounding algorithm.

THEOREM 6. *The above approximation algorithm achieves an $O(\log n)$ approximation ratio.*

Proof. First, the utility of a fractional solution that can use at most one set is certainly upper-bounded by the solution that is allowed to use multiple sets, *i.e.,* the value of the LP-solution of USSP is an upper bound on the optimum UOSSP value.

If the algorithm selects a set S_x, then the lifetime of S_x is at least x (since the fractional solution was able to send at least x units of data for each node in S_x). Thus, the algorithm obtains total utility at least

$$x \cdot u(S_x) \quad = \quad x \cdot \sum_{P \subseteq S_x, P \in \mathscr{P}} w_P.$$

If we sort the sets $P \in \mathscr{P}$ by non-increasing t_P values, numbering them $P_1, \ldots, P_k$, then we can rewrite and lower-bound the algorithm's total utility as $\max_{i=1}^{k} t_{P_i} \cdot \sum_{j=1}^{i} w_{P_j}$.

On the other hand, the optimum fractional solution obtains total utility $\sum_{i=1}^{k} t_{P_i} w_{P_i}$. By the above lemma, the ratio between these two quantities is bounded by $H_k = O(\log k)$, which in turn is $O(\log n)$ whenever the size of the collection $\mathscr{P}$ is polynomial in n. ■

If we want to improve the approximation guarantee beyond the factor $O(\log n)$ proved above, the approach will have to be based on a different upper bound from the one given by the USSP LP. (Notice that the USSP LP is also the natural LP relaxation for UOSSP.) For the family of examples depicted in Figure 2 exhibits an $O(\log n)$ integrality gap in the LP. The graph has n nodes in addition to the root. Each node is connected to the root through an edge of transmission cost 1. All the other n sensor nodes have sensing cost 0, unit initial energy 1, and weights $1, 1/2, 1/3, \ldots, 1/n$. The fractional solution sends $1/i$ units of data from each node i, for a total utility of H_n. On the other hand, the UOSSP solution can only chose one set, all of whose nodes have to send the same amount of data. Since any set of i nodes can send at most $1/i$ units (one of them must be a node with energy at most $1/i$), the maximum utility for UOSSP is 1, proving an integrality gap of $H_n = \Omega(\log n)$.

4.3 Greedy Algorithms

The $O(\log n)$ LP-rounding algorithm given above implicitly shows that selecting the best single set gives utility within $O(\log n)$ of selecting the best sequence. For many types of problems (such as maximizing the value of a submodular monotone function $f(S)$ on sets S), simple greedy algorithms based on adding or removing one

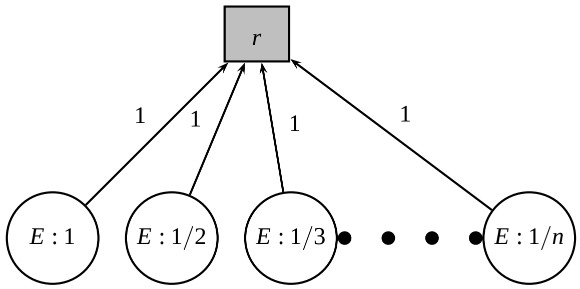

Figure 2: $O(\log n)$ Integrality Gap Example

element at a time lead to optimal solutions or good approximations. Here, we show that neither the greedy addition nor removal algorithm will give a good approximation for supermodular functions.

Specifically, we consider the following two algorithms. In the description, we let t_S be the maximum amount of data that the set S can send if no other nodes are sending data.

Addition Algorithm: Start with the empty set $S_0 = \emptyset$. In each iteration i (until all the nodes are selected), always add the node v_i such that the total utility $t_{S_{i-1}\cup\{v_i\}} \cdot u(S_{i-1}\cup\{v_i\})$ is maximized. (That is, set $S_i := S_{i-1}\cup\{v_i\}$.) Output the best of all the S_i.

Removal Algorithm: Start with all nodes $S'_0 = V$. In each iteration i (until no nodes are left), always remove the node v'_i such that the total utility $t_{S'_{i-1}\setminus\{v'_i\}} \cdot u(S'_{i-1}\setminus\{v'_i\})$ is maximized. (That is, set $S'_i := S'_{i-1}\setminus\{v'_i\}$.) Output the best of all the S'_i.

Even if the two algorithms are combined (*i.e.,* the better of their solutions is output), the approximation guarantee can be as bad as $\Omega(n)$. The example has a network consisting of the sink r, a bottleneck node $\hat{v}$, and n other nodes v_i. The bottleneck $\hat{v}$ is connected to all other nodes with edges of transmission cost 1, and it has an energy of 1. No other edges exist, and the energy of each v_i is also 1. Then, it is easy to see that the amount of data that can be sent by a set S is $1/|S|$, and hence, the optimum solution will be the one maximizing $\frac{u(S)}{|S|}$.

The supermodular utility function u is defined as follows: We partition the nodes $v_1,\ldots,v_n$ into $S_1 := \{v_1,v_2\}$ and $S_2 := \{v_3,\ldots,v_n\}$. Their utilities are $u(S_1) = \frac{1}{2} - \varepsilon n$ for some very small constant ε, and $u(S_2) = \frac{1}{2} + \varepsilon n$. The utility of the entire set is $u(\{v_1,\ldots,v_n\}) = 1$.

Each node in S_1 has individual utility $\varepsilon/2$, and each node in S_2 has utility ε. Any set $S \supseteq S_1$ has utility $u(S) = u(S_1) + \varepsilon \cdot |S\cap S_2|$, and any set $S \supseteq S_2$ has $u(S) = u(S_2) + \varepsilon/2 \cdot |S\cap S_1|$. Finally, all other sets S have utility $u(S) = \varepsilon \cdot |S\cap S_2| + \varepsilon/2 \cdot |S\cap S_1|$. It is not difficult to verify that this function u is indeed supermodular.

The greedy addition algorithm will add all nodes from S_2 first, never encountering the set S_1 as a possible solution. Similarly, the removal algorithm will remove the nodes from S_1 first. Hence, the output of any solution based on the greedy algorithm will be the set of all nodes, giving a total utility of $1/n$. On the other hand, the set S_1 can send half a unit of data, and thus achieves utility $\frac{1}{2}(1/2 - n\varepsilon) \approx \frac{1}{4}$, which is better than the greedy solution by a factor of $\Omega(n)$.

5. GEOMETRIC PENALTY FUNCTIONS

In the previous two sections, we investigated the special cases of submodular and supermodular utility functions. One of our main goals was to find a compact representation for utility, as arbitrary utility functions require specifying up to 2^n different function values.

Another approach to obtain compactly represented utility functions is based on the observation that the real-world application of sensor networks frequently involves monitoring geometric entities such as structures or habitats. In those cases, the quality of a solution can often be characterized in terms of the average or maximum distance of "interesting" points on the geometric entity from the measuring sensors.

Specifically, let $A \subseteq \mathbb{R}^2$ denote an *area* to be monitored, and $w : A \to \mathbb{R}^+$ a *weight function* measuring how "important" it is to measure any given point $x \in A$. Notice that A may be discrete or continuous. We let $d_{x,y}$ be a distance metric measuring how effectively a node at point x can measure the point y; most frequently, we will use the Euclidean distance $d_{x,y} = \sqrt{(x_1-y_1)^2+(x_2-y_2)^2}$, but our framework also permits using different metrics, or distance measures that are not truly metrics (such as the Euclidean distance with the additional provision that the $d_{x,y} = \infty$ if x and y are too far apart). Given a set S of sensors, we use $d_{S,x}$ to denote the distance of the point x from the sensor set, *i.e.,* $d_{S,x} = \min_{y\in S} d_{y,x}$. Based on this notion of distance, we can define the weighted maximum and average distance *penalty* as $\hat{p}(S) = \max_{x\in A} d_{S,x}$, and $\overline{p}(S) = \int_{x\in A} w_x \cdot d_{S,x}dx$. (If A is a discrete set, the integral should be replaced by a sum.)

Based on this penalty, we can now define the utility of a set to be inversely proportional to its penalty, *i.e.,* $u(S) := 1/\hat{p}(S)$ or $u(S) := 1/\overline{p}(S)$. Alternatively, we can phrase the problem as a penalty minimization instead of a utility maximization. In particular, in many monitoring applications, the goal will be to chose one best set of sensors that will monitor the environment. Thus, the goal would be to choose a set S such that $\hat{p}(S)/t_S$ or $\overline{p}(S)/t_S$ is minimized subject to the constraint that the set S can indeed send t_S units of data.

Unfortunately, the problem of choosing the best single set is NP-hard for both of these objectives, at least if the amount of data sent by all nodes must be integral.

THEOREM 7. *If t_S is required to be integral, then finding the set S minimizing $\hat{p}(S)/t_S$ or $\overline{p}(S)/t_S$ is NP-hard. In fact, unless P=NP, $\hat{p}(S)/t_S$ cannot be approximated to within better than 1.822.*

Proof. We give reductions from the Euclidean k-median and Euclidean k-center problems, respectively. The former was shown to be NP-hard by Papadimitriou [19], and the latter was shown to be hard to approximate to within 1.822 by Feder and Greene [20]. In both cases, we are given a finite set of points $A \subseteq \mathbb{R}^2$, and are asked to find a subset $S \subseteq A$ of size at most k. For the k-median problem, the objective is to minimize $\sum_{x\in A} d_{S,x}$, and for the k-center problem, to minimize $\max_{x\in A} d_{S,x}$.

For the reduction, in both cases, we let A be the area to be covered by sensors. There is a sensor node v_x with energy 1 located at each point $x \in A$, and all the v_x are connected to a bottleneck node $\hat{v}$ with edges of sending cost 1. The bottleneck in turn is connected to the root r via an edge of sending cost 1, and has energy k.

To prove that this is a correct (and approximation-preserving, in the case of k-center) reduction, first notice that if S is a solution for k-center (or k-median), then having each node v_x with $x \in S$ send one unit is feasible, and gives exactly the same objective value. Conversely, if S is the set of nodes sending data in a solution to the sensor selection problem, then by integrality and the energy constraints of 1 on nodes v_x, each node in S sends exactly one unit of data. But then, the objective function value of the set S as a solution to k-center or k-median is exactly its penalty in the corresponding sensor selection problem. This completes the proof. ∎

This result does not rule out the possibility that sending fractional amounts of data may make the problem tractable, or the existence of good approximation algorithms for the objective $\overline{p}(S)/t_S$. Indeed, as shown by Arora *et al.* [21], there is a Polynomial-Time Approximation Scheme (PTAS) for the Euclidean k-median problem.

6. RELATED WORK

Energy-conservation in sensor networks has received a tremendous amount of attention in the literature. For example, there is extensive work on overhearing avoidance and node duty-cycling at the MAC layer [22, 23, 24, 25]. As well, several pieces of work have focused on topology control [26, 27, 28, 29] to conserve energy in dense deployments. Our utility-based sensor selection is largely complementary, in that applications, rather than the system, indicate which set of sensors should be active. Also complementary, for a similar reason, is work on increasing network lifetime using energy-aware routing [1, 2, 3, 4, 5, 6, 7, 8].

Utility functions have significant prior history in networking, having been used to model architectural questions relating to network QoS [30] and network pricing [31]. In the sensor networks literature, however, we have found relatively few applications thereof. Perhaps closest to our work is that of Byers and Nasser [9], who first proposed utility-based sensor selection. However, they have analytically examined only the special class of utility functions that depend exclusively on the size of the set. More recently, Isler and Bajcsy [32] used utility functions to model the sensor selection problem for target localization. By contrast to both these pieces of work, we examine a broader class of utility functions. Finally, the work of Mainland et al. [33] is tangentially relevant to ours. They proposed a price-based resource management scheme for sensor networks. Their focus is more on the mechanistic aspects of resource management, whereas ours is algorithmic.

7. CONCLUSIONS

In this paper, we have examined a utility-based sensor selection approach that enables sensor network applications to express utilities associated with retrieving data from sets of sensors. We study the feasibility of determining, subject to network lifetime constraints, the sequence of sets whose data has the maximum total utility.

We explore fractional and integral variants of the utility-based sensor selection problem, as well as single-set variants thereof, on three important classes of utility functions. We find that many variants are NP-hard, and some hard to approximate. On the other hand, submodular functions can be optimized efficiently, and an important subclass of supermodular functions admits a fractional solution via solving an LP, and an $O(\log n)$-approximation when nodes are constrained to send the same amount of data. Finally, we show that geometric utilities can be cast into a penalty framework for which we are able to prove preliminary hardness results.

This paper should be considered a first step towards a much more thorough exploration of utility-based sensor set selection. The utility-based framework is a very natural way of expressing the true goal of a senor network application, and being able to select (approximately) optimal schedules for practical classes of utility functions is crucial in making the best use of a deployment. The geometric penalty framework appears to lead to more tractable classes than, *e.g.*, supermodular functions, and we would not be surprised if efficient approximations are possible. Similarly, it would be desirable to derive matching upper and lower bounds for supermodular functions, and to characterize more generally the classes of utility functions that would be useful for sensor network applications.

8. REFERENCES

[1] Q. Dong, "Maximizing system lifetime in wireless sensor networks," in *IPSN*, 2005.

[2] C. Schurgers and M. B. Srivastava, "Energy efficient routing in wireless sensor networks," in *MILCOM*, Vienna, VA, 2001, pp. 357–361.

[3] Q. Li, J. Aslam, and D. Rus, "Online power-aware routing in wireless ad-hoc networks," in *MobiCom '01: Proceedings of the 7th annual international conference on Mobile computing and networking*. ACM Press, 2001, pp. 97–107.

[4] J.-H. Chang and L. Tassiulas, "Fast approximate algorithms for maximum lifetime routing in wireless ad-hoc networks," in *NETWORKING '00: Proceedings of the IFIP-TC6 / European Commission International Conference on Broadband Communications, High Performance Networking, and Performance of Communication Networks*. Springer-Verlag, 2000, pp. 702–713.

[5] G. Xing, X. Wang, Y. Zhang, C. Lu, R. Pless, and C. Gill, "Integrated coverage and connectivity configuration for energy conservation in sensor networks," 2005.

[6] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. B. Srivastava, "Topology management for sensor networks: Exploiting latency and density," in *Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Lausanne, CH, 2002, pp. 135–145.

[7] I. Kang and R. Poovendran, "Maximizing static network lifetime of wireless broadcast adhoc networks," in *IEEE ICC*, Anchorage, Alaska, 2003.

[8] N. Ehsan and M. Liu, "Minimizing power consumption in sensor networks with quality of service requirement," in *to appear in Annual Allerton Confercence on Communications, Control and Computing (Allerton 2005)*, Allerton, IL, 2005.

[9] J. Byers and G. Nasser, "Utility-based decision-making in wireless sensor networks, Tech. Rep. 2000-014, 1 2000. [Online]. Available: citeseer.ist.psu.edu/byers00utilitybased.html

[10] "The Extensible Sensing System." [Online]. Available: http://www.cens.ucla.edu/ eoster/ess/

[11] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, and D. Culler, "An analysis of a large scale habitat monitoring application," in *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*. ACM Press, 2004, pp. 214–226.

[12] R. Szewczyk, E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin, "Habitat monitoring with sensor networks," *Commun. ACM*, vol. 47, no. 6, pp. 34–40, 2004.

[13] A. Arora, R. Ramnath, E. Ertin, and P. S. et. al., "Exscal: Elements of an extreme scale wireless sensor network," in *11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, 2005.

[14] R. Govindan, E. K. D. Estrin, F. Bian, K. Chintalapudi, O. Gnawali, S. Rangwala, R. Gummadi, and T. Stathopoulos, "Tenet: An Architecture for Tiered Embedded Networks," Tech. Rep., November 10 2005.

[15] J. Paek, K. Chintalapudi, J. Cafferey, R. Govindan, and S. Masri, "A wireless sensor network for structural health monitoring: Performance and experience," in *Proceedings of the Second IEEE Workshop on Embedded Networked Sensors (EmNetS-II),*, Syndney, Australia, May 2005.

[16] A. Vetta, "Nash equilibria in competitive societies with

applications to facility location, traffic routing, and auctions," in *FOCS*, 2002.

[17] U. Feige, G. Kortsarz, and D. Peleg, "The dense *k*-subgraph problem," in *Proc. 25th ACM Symp. on Theory of Computing*, 1993.

[18] U. Feige and M. Seltser, "On the densest *k*-subgraph problem," The Weizmann Institute, Rehovot, Tech. Rep., 1997.

[19] C. Papadimitriou, "Worst-case and probabilistic analysis of a geometric location problem," *SIAM Journal on Computing*, vol. 10, pp. 542–557, 1981.

[20] T. Feder and D. Greene, "Optimal algorithms for approximate clustering," in *Proc. 20th ACM Symp. on Theory of Computing*, 1988.

[21] S. Arora, P. Raghavan, and S. Rao, "Approximation schemes for euclidean *k*-medians and related problems," in *Proc. 30th ACM Symp. on Theory of Computing*, 1998.

[22] J. Polastre, J. Hill, and D. Culler, "Versatile Low Power Media Access for Wireless Sensor Networks," in *Proceedings of the ACM Conference on Embedded Networked Sensor Systems*, Baltimore, MD, November 2004.

[23] T. Dam and K. Langendoen, "An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks," in *Proceedings of the ACM Conference on Embedded Networked Sensor Systems*, Los Angeles, CA, November 2003.

[24] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient mac protocol for wireless sensor networks," in *Proceedings of the IEEE Infocom*, June 2002.

[25] M. Liu and C. fan Hsin, "Network coverage using low duty-cycled sensors: Random and coordinated sleep algorithms," in *IPSN*, 2004.

[26] S. Bhattacharya, G. Xing, C. Lu, G.-C. Roman, B. Harris, and O. Chipara, "Dynamic wake-up and topology maintenance protocols with spatiotemporal guarantees," in *International Conference on Information Processing in Sensor Networks (IPSN)*, Los Angeles, CA, 2005.

[27] A. Cerpa and D. Estrin, "ASCENT: Adaptive self-configuring sensor networks topologies," in *Proceedings of the IEEE Infocom*. New York, USA: IEEE, June 2002.

[28] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, "Span: An Energy-efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks," in *Proceedings of the IEEE Infocom*. IEEE Computer Society Press, 2001, pp. 85–96.

[29] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*. Rome, Italy: ACM, July 2001, pp. 70–84.

[30] S. Shenker, "Fundamental design issues for the future internet," September 1995.

[31] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," in *Journal of the Operational Research Society*, vol. 49, 1998. [Online]. Available: citeseer.csail.mit.edu/kelly98rate.html

[32] V. Isler and R. Bajcsy, "The sensor selection problem for bounded uncertainty sensing models." in *IPSN*, 2005, pp. 151–158.

[33] G. Mainland, D. C. Parkes, and M. Welsh, "Decentralized, adaptive resource allocation for sensor networks," in *In Proceedings of the 2nd USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI)*, 2005.