Moving Multiple Sinks Through Wireless Sensor Networks for Lifetime Maximization*

S. Basagni
ECE Dept.
Northeastern University
basagni@ece.neu.edu

A. Carosi, C. Petrioli
Dip. di Informatica
Univ. di Roma "La Sapienza"
{carosi,petrioli}@di.uniroma1.it

C. A. Phillips
Sandia National Labs
Albuquerque, NM
caphill@sandia.gov

Abstract

We propose scalable models and centralized heuristics for the concurrent and coordinated movement of multiple sinks in a wireless sensor network (WSN). The proposed centralized heuristic runs in polynomial time given the solution to the linear program and achieves results that are within 2% of the LP-relaxation-based upper bound. It provides a useful benchmark for evaluating centralized and distributed schemes for controlled sink mobility.

1 Introduction

Recent research on wireless sensor networks (WSNs) has clearly shown that moving some of the network components can reduce energy consumption and prolong the time the network is able to perform its functions (*network lifetime*). In particular, researchers have shown through analysis and experiments that moving the network data collection points, called the *sinks*, leads to the highest values for network lifetime. Recent papers [2,3,8,10] have presented analytical models and distributed heuristics for controlling the mobility of a single sink, achieving up to fivefold improvement over the case of one optimally-placed static sink.

In this paper we investigate the problem of controlling the concurrent and coordinated mobility of *multiple* sinks for prolonging the lifetime of a WSN, a case that presents significant challenges compared to the single mobile sink case. We make two contributions: 1) We define a scalable analytical model that provides a provable upper bound on the maximum lifetime of a WSN. The model allows varying parameters, data routing schemes and sensor deployments. Our model is scalable to realistic network sizes (i.e., WSNs hundreds of nodes), so it can serve as a benchmark for current and future distributed heuristics. 2) We present a centralized heuristic that runs in polynomial time given the solution to the linear program. This heuristic scales to realistic network sizes (400 nodes, 64 different positions where the sink can sojourn, 5 sinks concurrently roaming throughout the network) and gives sink schedules that are within 2% of the LP-relaxation-based upper bound. See [4] for a more complete discussion of these results.

2 Problem Formulation

We envision a two-tier scenario consisting of a small number s of mobile sinks deployed to collect data from a large number of resource-constrained sensor nodes (or sensors for short) that are monitoring the deployment area. The sinks, e.g., mobile robots or unmanned flying drones, roam through the network moving among a finite number of designated sink sites. Mobile sinks are resource-rich nodes, i.e., they are not particularly constrained in terms of energy, mobility, computational resources, and storage. We assume that they have multiple radio interfaces. They collect data from the sensors by using wireless sensor technologies and they coordinate among themselves over longer distances by means of higher data rate technologies. Each sensor p in a set N of sensors has an initial energy e_p. It generates data at a rate of r_p packets per second. These packets are routed either directly or via a multi-hop path to the sink sojourning at the closest site (in terms of number of hops) according to a given routing protocol. Sensor p requires α_p Joules per packet for sensing, creating, and transmitting packets it generates *locally*. It requires β_p Joules per packet to *receive* and *relay* a packet for another sensor. Since the sinks are mobile, the sensors must learn the closest sink that is avail-

* This work was partially supported by NSF grant #0738720, by the FP7 EU project "SENSEI, Integrating the Physical with the Digital World of the Network of the Future," Grant Agreement Number 215923, www.ict-sensei.org, and by the International FIRB contract RBIN047MH9. Sandia is a multipurpose laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. The authors are grateful to Bob Carr of Sandia National Laboratories for useful comments.

1-4244-2575-4/08/$20.00

able to receive their packets. More specifically, at any time, a sink is either at a sink site or it is moving. If it is at a sink site, it can be either *active*, meaning it is available to receive sensor data, or *inactive*. When a sink at a sink site becomes active, it must broadcast its availability to the sensors. Similarly, when it becomes inactive, it must broadcast its unavailability. These broadcasts require energy from the sensors and are sink-site dependent. An active sink must deactivate before moving. We require that at least one sink is active at all times. This design choice allows the sensors to always have a destination for their packets. This avoids high packet end-to-end latencies since nodes do not need to buffer packets while the sinks are moving. To ensure that routes can stabilize, we require that a *configuration* (i.e., a set of active sinks associated with sink sites) holds steady at least t_{min} seconds.

Our objective is to find a schedule of sink movements that maximizes the network lifetime, which is reached when the first sensor runs out of battery power.

3 Mathematical Model

In this section we describe a linear program (LP) whose solution provides an upper bound on the maximum lifetime of a sensor network with s mobile sinks. The LP is a relaxation of the mobile sinks problem. It selects a set of configurations and the time for each configuration. It ignores the ordering of configurations and the transitions between them, minimum times for configurations to hold, and energy costs for broadcasting sink (de)activations.

The LP has only one type of variable. Let $t_c \geq 0$ be the time that configuration $c \in C$ holds. There are $\sum_{k=2}^{s} \binom{|V|}{k}$ ways to place at most s sinks on $|V|$ sites, which is exponential in s. However, most of these variables will be 0 and not explicitly represented. The LP model is to maximize $\sum_{c \in C} t_c$ subject to:

$$\sum_{c \in C} \left(\alpha_p r_p t_c + \sum_{q \in N, w \in V} r_q \beta_p \rho_{qpw} y_{qwc} t_c \right) \leq e_p,$$

$\forall p \in N$. The objective function maximizes network lifetime. The constraints require all sensors to be alive throughout the network lifetime. Given a sensor $p \in N$ and a configuration $c \in C$, the expression in the parentheses on the left side of the constraint for sensor p gives the energy cost for sensor p while the system is in configuration c. The second term computes the energy sensor p expends routing packets for other sensors during configuration c. Each term in this sum depends on the data rate at sensor q, which might route through sensor p depending on sensor q's choice(s) of sink and choice(s) of route.

Solving the LP. One can optimize an LP (the primal) by instead optimizing its dual [9]. This is another LP, derived mechanically from the primal, which has the same optimal objective value. The dual of our LP has an exponential number of constraints, one for each possible configuration.

We can satisfy an exponential-size set of constraints by only explicitly enforcing a polynomial number of them provided we have a *separation algorithm*. Suppose we wish to enforce a set of constraints $\mathcal{X}$. A separation algorithm accepts a solution vector for the LP. It then either says that all constraints in $\mathcal{X}$ are satisfied by the solution, or it gives a constraint $x \in \mathcal{X}$ that is violated. This separation algorithm must exploit the structure of the constraint family; it cannot enumerate the constraints.

To use a separation algorithm to solve a large LP, we begin by explicitly listing none of the constraints in $\mathcal{C}$. We then solve the LP and pass the solution to the separation algorithm. If all the constraints are satisfied, we are done. Otherwise, we add the violated constraint it returns, solve the new LP, and so on. This is guaranteed to finish in a polynomial number of iterations [11].

By considering the structure of a violated constraint of our LP, it is possible to see that we can find a violated constraint by solving a p-median problem [12]. The p-median problem is NP-complete. However, there are exact algorithms based on math-programming formulations that can effectively solve problems with ten thousand or more nodes and hundreds of thousands of customers [5]. This is well within or beyond the range of current sensor network sizes. Thus, although we cannot solve the LP in guaranteed polynomial time, we can solve it in practice for the sizes of LP we require. This determines the upper bound on the wireless sensor network lifetime.

A centralized heuristic. The high-level operations for the centralized heuristic are as follows:

1. Solve the LP to obtain a solution t_c^*. Let $B = t_{min}$.
2. Let $\mathcal{C} \subseteq C = \{c \in C | t_c^* \geq B\}$. That is, we select the set of configurations for which the LP assigns a time of at least B (initially t_{min}).
3. Order the configurations in $\mathcal{C}$, preferably with more closely-related configurations near each other.
4. Compute transient configurations between each pair of adjacent configurations.
5. Solve another final LP to adjust the times for each configuration, enforce minimum times on configurations, and account for sink-movement broadcast costs.
6. If the final LP is infeasible, increase B (e.g., to remove the next-shortest configuration). Return to step 2.

We now consider each step of the centralized algorithm, starting with step 3. We use a simple traveling salesperson (TSP) model. Two (ordered) configurations $c_i, c_j \in C$ have distance equal to $1 + \gamma$, where γ is the minimum number of intermediate configurations required to transition between c_i and c_j. We wish to find a traveling salesman path (not a closed tour) among the chosen configurations. The optimal

TSP minimizes the number of configurations we must add, which heuristically minimizes the amount of time spent in these added configurations. These are tiny and easy problems for the free TSP code Concorde [7]. However, one could also use a polynomial-time approximation algorithm for TSP such as Christofides' heuristic [6].

We now consider how to compute the transient configurations in step 4. Because sinks require time to move between sink sites and because we require at least one active sink at all times, we forbid certain types of transitions that appear to "teleport" sinks. We can transition between any pair of configurations using at most two intermediate configurations. Suppose we wish to compute the transition states between two states c_i and c_{i+1} that are adjacent in the ordering computed in step 3. Here we assume this transition requires 2 intermediate steps, since the other cases are almost completely determined. The initial LP chooses a set of configurations for which the vectors of energy costs for each sensor (weighted by t_i^*) pack well into the vector of initial sensor energies. We try to keep the transient configurations between a pair of configurations as close to these selected configurations as possible. We estimate the routing distance between each pair of sink sites in c_i. We then find a minimum-weight maximum-cardinality matching in the complete graph with a node for each site in c_i, and edges weighted by this pairwise distance. We then pick an element from each matched pair arbitrarily and move this set of sinks M. Our hope is that if sites v_i and v_j are matched, the nodes sending to v_i will be instead redirected to v_j (and vice-versa), inducing energy consumption patterns that approximate those in configuration c_i. We then compute a similar matching in c_{i+1} and use that to pick the new locations L for the sinks from sites M to move to. So the transient configurations are $c_i - M$ and L.

We now consider the final LP. Because we have selected the precise set of configurations (steps 2 to 4), we now no longer have to allow for zero values of t_c. So we can enforce minimum times for configurations. Because we know the order of the configurations, we can account for broadcast costs. Because we consider only a polynomial number of configurations, we can solve this LP in polynomial time.

4 Performance Evaluation

In this section we evaluate the performance of the centralized heuristic (CEN) we have proposed by comparing it to the upper bound on an optimal solution (OPT), to random sink mobility (RND), and to the case where sinks are static and optimally placed (STATIC). Our simulator considers realistic parameters of wireless sensor networking.

We have considered the following scenarios. 400 wireless sensor nodes are deployed on a 20×20 grid over a square deployment area of side $L = 475$m. The node transmission radius is 25m. This allows the nodes to have at most 4 neighbors. Each node has an initial energy of 50J. Each node generates packets of size 512B at a rate of $r = 0.5$bps. They are sent to the closest sink according to a (hop-based) shortest path routing. The channel data rate is 250Kbps (consistent with IEEE 802.15.4). The transmission power and the receiving power are 0.0144W and 0.0125W, respectively, according to the specifications of the TR 1000 radio transceiver from RF Monolithics [1]. Sinks can sojourn at any of the sites of a 8×8 grid. We vary the number of sinks in the range $[2, 5]$. Protocol related parameters are configured as follows. The time sinks are forced to stay at a site, t_{min}, ranges in the set $\{50, 100, 250\}$Ks.We considered the energy spent by the nodes during the route maintenance due to sink movements. We ran 100 experiments for each choice of parameters, which achieves a 95% confidence level within a 5% precision.

Table 1 shows the network lifetime (in millions of seconds, rounded) induced by the various protocols when varying t_{min} and the number s of sinks. Each table entry shows the absolute lifetime (in millions of seconds) and the percentage decrease with respect to OPT (in parenthesis).

Table 1. Lifetime (and % gap from OPT)

s	t_{min}	OPT	CEN	RND	STATIC
2	50K	79.51	78.8 (0.9)	39.2 (50.7)	14.2 (82.1)
	100K	79.51	78.8 (0.9)	39.2 (50.7)	14.2 (82.1)
	250K	79.51	78.4 (1.4)	38 (52.2)	14.2 (82.1)
3	50K	105.9	105.2 (0.6)	50.7 (52.1)	20.8 (80.3)
	100K	105.9	105.3 (0.5)	50.7 (52.1)	20.8 (80.3)
	250K	105.9	105.2 (0.6)	49.3 (53.4)	20.8 (80.3)
4	50K	131.4	130.6 (0.6)	63.4 (51.7)	27.7 (78.9)
	100K	131.4	130.6 (0.6)	63.6 (51.6)	27.7 (78.9)
	250K	131.4	130.5 (0.6)	61.5 (53.2)	27.7 (78.9)
5	50K	150.1	149.2 (0.6)	75.6 (49.6)	34.3 (77.1)
	100K	150.1	149.2 (0.6)	75.9 (49.4)	34.3 (77.1)
	250K	150.1	149.2 (0.6)	73.9 (50.7)	34.3 (77.1)

The centralized heuristic CEN achieves network lifetimes that are remarkably close to the optimum. The gap from OPT is always below 2%. OPT selects a set of configurations that are particularly effective at balancing energy use among the network nodes. CEN spends most of the time in a subset of OPT's configurations. CEN spends minimum time in intermediate configurations designed to mimic these selected configurations. The increase of CEN lifetime over STATIC can be as high as 335%.

Sink mobility is advantageous even when unaware of any nodal status. The random movements of the sinks in RND double the network lifetime with respect to STATIC. However, controlled mobility is much more effective, inducing network lifetimes that are over twice the RND-induced lifetime. OPT and the other heuristics induce longer lifetimes when the number of sinks increases. Because the network traffic is partitioned among a larger set of sinks, the sink's neighbors receive fewer packets so their energy consumption is reduced. However, the lifetime improvements are not linear. Suppose we charge the cost of sending or relaying a packet to the packet's destination sink. Achieving a double (threefold, etc.) improvement in lifetime when doubling (tripling, etc.) the number of sinks would require a set of configurations that consistently reduces the traffic to each single sink and perfectly balance the total energy consumption among the nodes close to the sinks. The energy costs of any given configuration depend on the routing strategy. Routes may not pack well in this way. Furthermore, the fixed set of sink sites somewhat limits the set of choices. OPT only achieves a 33.2% lifetime increase when we deploy 3 sinks instead of 2, and a 65.26% improvement when increasing the number of sinks from 2 to 4.

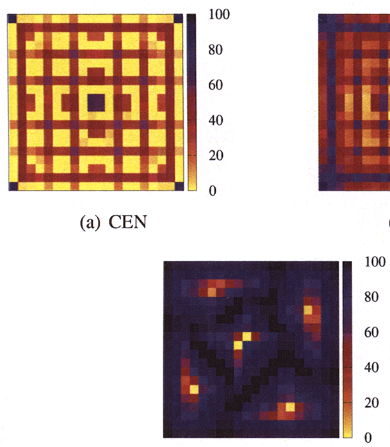

(a) CEN (b) RND

(c) STATIC

Figure 1. Residual energy at lifetime

Beyond achieving improved network lifetime, CEN also results in a more even distribution of nodal residual energies compared to RND and STATIC. Figure 1 displays the residual energy of the nodes at network lifetime for a given run in scenarios when 5 sinks can select among 64 sites (performance does not significantly change when considering different runs). A lighter color means a lower percentage of residual energy in that area of the network. CEN results in better load balancing than RND, which in turn improves over STATIC. At lifetime CEN shows an impressive percentage of nodes with very little energy left, a witness of its good energy drainage balancing property. The fraction of nodes with less than 20% (40%) residual energy is 52.5% (80%) in CEN. These figures reduce to 5.41% and 23.7% for RND and to 1.75% and 3.5% for STATIC.

References

[1] ASH transceiver designer's guide. www.rfm.com, 2004.

[2] A. P. Azad and A. Chockalingam. Mobile base station placement and energy aware routing in wireless sensor networks. In *Proceeding of the IEEE Wireless Communications and Networking Conference, WCNC 2006*, volume 1, pages 264–269, Las Vegas, NV, April 3–6 2006.

[3] S. Basagni, A. Carosi, E. Melachrinoudis, C. Petrioli, and Z. M. Wang. Controlled sink mobility for prolonging wireless sensor networks lifetime. *ACM/Springer Wireless Networks*, 2008. To appear.

[4] S. Basagni, A. Carosi, C. Petrioli, and C. A. Phillips. Coordinated and controlled mobility of multiple sinks for maximizing the lifetime of wireless sensor networks. Technical Report SAND2008-4670J, Sandia National Laboratories, 2008.

[5] J. Berry, W. Hart, C. Phillips, J. Uber, and J.-P. Watson. Sensor placement in municipal water networks with temporal integer programming models. *Journal of Water Resources Planning and Management*, 132:218–224, 2006.

[6] N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Carnegie-Mellon University, Graduate School of Industrial Administration, 1976.

[7] Concorde, 2005. www.tsp.gatech.edu/concorde.html.

[8] S. R. Gandham, M. Dawande, R. Prakash, and S. Venkatesan. Energy efficient schemes for wireless sensor networks with multiple mobile base stations. In *Proceedings of IEEE Globecom 2003*, volume 1, pages 377–381, San Francisco, CA, December 1–5 2003.

[9] F. S. Hillier and G. J. Lieberman. *Introduction to Operations Research*. McGraw Hill, Boston, MA, eighth edition, 2005.

[10] J. Luo and J.-P. Hubaux. Joint mobility and routing for lifetime elongation in wireless sensor networks. In *Proceedings of IEEE Infocom 2005*, volume 3, pages 1735–1746, Miami, FL, March 13–17 2005.

[11] L. L. M. Grötschel and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.

[12] C. Revelle and R. Swain. Central facilities location. *Geographical Analysis*, 2:30–42, 1970.