An Integrated Approach to Sensor Role Selection

Mark Perillo and Wendi Heinzelman, *Member, IEEE*

Abstract—Many sensor network applications require consistent coverage of the region in which they are deployed over the course of the network lifetime. However, because sensor networks may be deployed randomly, node distribution and data redundancy in some regions of the network may be lower than in others. The sensors in the sparsest regions should be considered more critical to the sensor network application since their removal would likely result in unmonitored regions in the environment. For this reason, sensors in the more densely deployed regions should be considered more favorable as candidates to route the traffic of other nodes in the network. In this work, we propose several coverage-aware routing costs that allow traffic to be routed around the sparsely deployed regions so that the coverage of the environment can remain high for a long lifetime. We also propose an integrated route discovery and sensor selection protocol called DAPR that further lengthens network lifetime by jointly selecting routers and active sensors, again with the goal of minimizing the use of sensors in sparsely covered areas. Simulation results show the effectiveness of our approach in extending network lifetime nearly to the extent that can be reached using a centralized approach based on global network knowledge.

Index Terms—Wireless Sensor Networks, Protocol Architectures, Routing Protocols.

1 Introduction

As wireless sensor networks continue to attract attention for use in numerous commercial and military applications, there have been many efforts to improve their energy efficiency so that they can operate for very long periods with no manual maintenance. Because of the limited energy supplies of typical microsensors, however, achieving long network lifetimes has been a very challenging task. A great deal of research has focused on power reduction in several areas, such as hardware, operating system, and low-level protocol design, in order to increase network lifetime. However, further steps must be taken in order to *balance*, as well as reduce, energy consumption so that sensor networks will be able to realize their maximum potential lifetime.

As the cost of manufacturing sensor nodes continues to decrease and large-scale networks consisting of thousands of sensors become realizable, the redundancy that exists among the data generated by the sensors can be exploited. Recent work in this area has focused on techniques such as dynamic sensor selection, in-network aggregation, and distributed source coding that reduce the amount of data generated by the network but ensure that the cumulative data from the sensor network at any given time meets the sensor network's application quality of service (QoS) requirements. In this work, we focus on networks in which data flow is reduced by dynamically selecting only a subset of the sensors in the network to generate data at a given time. The generated data are routed back to a single base station within the sensor network, where it may be processed locally or sent to an end user via a dedicated communication channel.

Depending on the nature of deployment, it may be the case that certain sensors are more important than others in a sensing role due to non-uniformities in sensor deployment, sensing capabilities, and initial energy resources. The loss of these critical sensors could lead to unattended regions of the network during early stages of the network lifetime. To avoid this situation, these sensors' use in network roles such as data routing should be avoided whenever possible in order to maximize their lifetime and in turn, the useful lifetime of the sensor network. In this work, we propose DAPR, an integrated routing and sensor selection protocol for wireless sensor networks that attempts to avoid these critical sensors by assigning novel routing costs that incorporate coverage overlap and choosing sensors to actively sense and generate data with the knowledge of the effects that this has on potential routers. Our proposed routing costs are the first that attempt to avoid routing through sensors that are critical in the sense of meeting application QoS requirements.

This paper expands on previous work, proposing variations on the routing costs originally presented in [1] and providing additional simulations and analysis to demonstrate the efficiency of our approach. Furthermore, we compare our distributed solutions with a centralized approach that uses global knowledge about the network topology, sensing capabilities, initial energy of each node, and base station locations over the course of the network's lifetime. Compared to this centralized approach, DAPR performs very well, not quite reach-

- *The authors are with the Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY, 14627.*
E-mail: {perillo,wheinzel}@ece.rochester.edu

This work is supported in part by the NSF under grants CNS-0448046 and ECS-0428157.
Manuscript received April 17, 2007; revised October 28, 2008.

ing the lifetime afforded by the centralized approach, but coming much closer than existing non-integrated approaches and using distributed methods.

The rest of this paper is organized as follows. Section 2 addresses related work. Section 3 formally presents the problem that we are addressing. Section 4 presents the proposed coverage-aware routing costs, and Section 5 presents the DAPR integrated sensor selection and routing protocol. Section 6 provides simulation results and analysis. Section 7 compares DAPR with a centralized approach to sensor scheduling and routing using global information. Section 8 concludes the paper and suggests future work in this area.

2 Related Work

A great deal of research has been dedicated to different areas of role section in wireless sensor networks. In this section, we give an overview of some relevant work in the development of distributed protocols for two areas of role selection - sensor selection and router selection.

2.1 Sensor Selection

Several energy-efficient coverage preservation protocols have been developed to provide consistent environmental coverage and robustness to unexpected sensor failures. In PEAS [2], sleeping sensors periodically enter a probing state, querying all sensors within a probing range (based on communication and/or connectivity requirements), and become active if no active sensors exist within the desired probing range. In [3], the problem of sensor selection was modeled as a Gur game, where sensors operate as finite state machines and change states (sending traffic only in selected ones) based on feedback from the base station, which is based on the state of the network's current data resolution. The authors in [4] propose a round scheduling scheme in which sensors exchange reference times and schedule themselves around their own reference time, guaranteeing that the environment is completely covered at all times. In [5], the authors propose a distributed selection algorithm for coverage preservation in sensor networks, in which a sensor measures its neighborhood redundancy as the union of the sectors/central angles covered by neighboring sensors within the sensor's sensing range. In CCP (Coverage Configuration Protocol) [6], sensors consult an eligibility rule, in which each sensor finds all intersection points between the borders of its neighbors' sensing radii and considers itself eligible for deactivation if each of these intersection points is covered with the desired sensing degree.

The aforementioned protocols generally aim to provide consistent coverage while ignoring the impact that active sensors will have on other sensors in the network, specifically the additional sensors that are required to route the data generated by the active sensors. The algorithms presented in [7], [8] consider routing implications when activating sensors. The goal of these algorithms is to find a minimum set of sensors and additional routers necessary in order to cover a given geographical region. Each iteration of the algorithm finds the sensor with the best combination of a short path to the active subset and a large number of additional unique sections covered. That sensor and those along its path are selected for inclusion in the sensor set. While this solution provides some integration of sensor and router selection, it uses a different model for achieving energy efficiency than the one used in this paper. Specifically, the protocol in [7] assumes that the activation of a sensor as a data generator or routers incurs a constant cost on that sensor, while we assume that power consumption is affected by the amount of traffic transmitted and received. Also, while the algorithm presented in [7] considers the potential routers when selecting which sensors to activate, it does not consider a node's role in the sensing application during route selection, and hence it does not avoid selecting critical sensors as routers, as is achieved in our approach.

2.2 Routing Protocols

The field of ad hoc routing has been explored extensively. Initially, protocol design focused on efficiently finding shortest path routes in the presence of node mobility [9]. Later research addressed the need for energy-based metrics to be used in energy-efficient ad hoc routing protocols. Singh et al. proposed several routing costs based on the residual energy of individual nodes [10]. Chang et al. proposed a routing cost that was a combination of residual energy, normalized residual energy, and required transmission energy and found an optimal combination of these parameters [11]. We build on this work by developing a routing cost for use specifically in wireless sensor networks, where the property of node redundancy is important. Our proposed coverage-aware routing cost is based not only on a sensor node's residual energy, but also the residual energy of redundant neighboring sensors, in order to ensure that the most critical sensors are avoided and live long enough to maintain high fidelity over long periods of time.

3 Problem Formalization

In this work, we assume the use of an application where the entirety or a portion of an area $\underline{\mathbf{A}}$ needs to be monitored by any one or multiple sensors that are within their sensing range of that location. We refer to the complete set of sensors as $S = \{s_1, \ldots, s_{N_s}\}$, where N_s represents the number of sensors deployed in the network. If we require the network to perform at some predetermined level of QoS, or fidelity, we can assign a nominal sensing range to the sensors so that sensors can adequately monitor activity within this sensing range (e.g., the signal-to-noise ratio exceeds a given threshold at this range). While the region that a sensor is able to cover does not necessarily need to be restricted to a disc-shaped region, we consider this model throughout this paper. In general, let $A(s_i)$ represent the sensor s_i's coverage area.

In order to provide coverage of the entire region that is being monitored, it is possible to activate the sensors in many combinations, each combination constituting a cover set. A cover set c_i is defined as any set in which the constraint given in Equation 1 is satisfied.

$$\bigcup_{j:s_j \in c_i} A(s_i) \supseteq \underline{\mathbf{A}} \tag{1}$$

We refer to the set of N_c cover sets as $C = \{c_1, \ldots, c_{N_c}\}$, where for each cover set, $c_k \subseteq S$. It should be noted that our problem formalization does not preclude the use of non-disjoint cover sets. The sensor network is periodically queried by a single data sink, and data from the sensors in the current cover set are routed back to this data sink. In this work, we consider the general scenario where the data sink's location does not remain constant over the course of the entire network lifetime. In general, we assume that there are N_{sink} data sink locations, represented as $S^* = \{s_{N_s+1}, \ldots, s_{N_s+N_{sink}}\}$. The total number of nodes in the network (sensors and data sinks) over the course of the network lifetime is $N_t = N_s + N_{sink}$.

We use the variable F to represent the traffic flow. F consists of components f_{ijm}, which represent the total traffic flow that sensor s_i forwards toward data sink s_{N_s+m} using s_j as its next hop. The variable T represents the scheduling of cover sets and component t_{km} represents the scheduled time of cover set c_k during the operation of data sink s_{N_s+m}. Furthermore, R, P^{sense}, represent the bit rate and sensing power consumption of an active sensor, respectively. E^{tx}, E^{rx}, and E^{init}. Elements e_{ij}^{tx}, e_{ij}^{rx}, and e_i^{init}, are vectors representing the energy consumption to transmit a bit from sensor s_i to sensor s_j, the energy consumption to receive a bit sent from sensor s_i to sensor s_j, and the initial energy of sensor s_i, respectively. Typically, the elements of E^{tx} depend on the distance of the link over which data is being transmitted, while the elements of E^{rx} do not. The variables used in this paper are summarized in Table 1.

In our problem, we are limited by several constraints, including constraints ensuring the conservation of data flow (i.e., that the sum of a node's incoming data and its generated data must equal its outgoing data), expressed in Equation 3.

$$\sum_{j=1}^{N_t} f_{ijm} = \sum_{j=1}^{N_s} f_{jim} + \sum_{k:s_i \in c_k} R \cdot t_{km} \tag{2}$$

$$\forall i \in \{1, \ldots, N_s\}, \quad m \in \{1 \ldots N_{sink}\} \tag{3}$$

Furthermore, we are limited by energy consumption constraints that limit the total amount of time any sensor can route and actively sense data by that node's initial energy.

$$\sum_{m=1}^{N_{sink}} \sum_{k:s_i \in c_k} P^{sense} \cdot t_{km} + \sum_{m=1}^{N_{sink}} \sum_{j=1}^{N_t} f_{ijm} e_{ij}^{tx} + \sum_{m=1}^{N_{sink}} \sum_{j=1}^{N_s} f_{jim} e_{ji}^{rx} \le e_i^{init} \qquad \forall i \in \{1, \ldots, N_s\} \tag{4}$$

Over the course of the network lifetime, queries may arrive from a number of locations within the network, either from multiple data sinks within the network or from a single mobile sink roaming throughout the network. The fraction of queries Q that will propagate from each data sink location impose additional constraints given in Equation 5.

$$\sum_{k=1}^{N_c} t_{km} = q_m \sum_{k=1}^{N_c} \sum_{m'=1}^{N_{sink}} t_{km'} \qquad \forall m \in \{1, \ldots, N_{sink}\} \tag{5}$$

The network lifetime L is the combined operating time of all individual cover sets.

$$L = \sum_{k=1}^{N_c} \sum_{m=1}^{N_{sink}} t_{km} \tag{6}$$

In this work, we propose distributed methods for determining the cover sets from C to use as the network operates over time. Furthermore, for the sensors included in any currently active cover set during the data sink's query, we choose routes from those sensors to that data sink, essentially building up the values in F over time. We choose the cover sets and routes in such a way to maximize the network lifetime L. Formally, our works attempts to solve the problem given in Figure 1. Section 4 and Section 5 describe a distributed solution to this problem by presenting routing costs that can be used so that the selected routes will build up the elements of F, and presenting a protocol for discovering these routes and choosing the cover sets over time. Section 7 shows how this problem can be solved to maximize lifetime in a centralized fashion using linear programming and compares the lifetime of our distributed protocol to this linear programming solution.

4 Coverage-Aware Routing Costs

Our work is motivated by the intuition that for collaborative sensor networks, application goals should play a role in many of the network decisions, such as which sensors to activate and how to route the data. Specifically, sensors that are more important to the sensing application as data generators (e.g., those that are located in sparsely deployed areas) should not be chosen as routers over those that are less important to the application (e.g., those with more redundant neighbors). Because these nodes are expected to be used most often to sense and generate data, they should be considered critical sensors and avoided as routers for other sensors that are generating data. Guided by this intuition, Given:

$S = \{s_1, \ldots, s_{N_s}\}$, the set of sensors.
$S^* = \{s_{N_s+1}, \ldots, s_{N_s+N_{sink}}\}$, the set of data sinks.
$A(s_i)$, each sensor's coverage area.
e_{ij}^{tx}, the energy consumption to transmit a bit from sensor s_i to sensor s_j.
e_{ij}^{rx}, the energy consumption to receive a bit sent from sensor s_i to sensor s_j.
e_i^{init}, the initial energy of sensor s_i.
q_m, the fraction of queries originating from each data sink.
R, the average traffic rate of a sensor included in an active cover set.
P^{sense}, the average sensing power consumption of a sensor included in an active cover set.

Compute:

$C = \{c_1, \ldots, c_{N_c}\}$, a set of cover sets.
t_{km}, the amount of time for which each cover set c_k is scheduled during the operation of data sink s_{N_s+m}.
f_{ijm}, the total traffic flow that sensor s_i forwards toward data sink s_{N_s+m} using s_j as its next hop.

Goal:

Maximize L, the network lifetime, as given in Equation 6, under the constraints given in Equations 1, 3, 4, and 5.

Fig. 1. Problem Formalization.

Description	Matrix	Elements	Size
Set of Sensors	S	s_i, $1 \leq i \leq N_s$	N_s
Set of Data Sinks	S^*	s_i, $N_s < i \leq N_t$	N_{sink}
Set of Cover Sets	C	c_k	N_c
Traffic Flow	F	f_{ijm}	$N_t \times N_t \times N_{sink}$
Cover Set Scheduling	T	t_{km}	$N_c \times N_{sink}$
Initial Energy	E^{init}	e_i^{init}	N_s
Transmit Energy (bit)	E^{tx}	e_{ij}^{tx}	$N_s \times N_t$
Receive Energy (bit)	E^{rx}	e_{ij}^{rx}	$N_s \times N_s$
Sink Time Fraction	Q	q_m	N_{sink}
Bit Rate	R		1
Sensing Power	P^{sense}		1
Lifetime	L		1

TABLE 1
Summary of variables in the optimization program.

we propose the use of "coverage-aware" routing costs that consider the importance of the node to the sensor network's application QoS requirements. Since certain nodes are more critical than others as data generators, using a coverage-aware routing cost allows the network to identify and avoid these sensors as routers.

A common "energy-aware" routing cost used in wireless ad hoc networks is the inverse of a node's residual energy e_i^{res}.

$$C_{ea}(s_i) = 1/e_i^{res} \tag{7}$$

With the use of this routing cost, nodes with little energy remaining are unlikely to be used to route the traffic of other nodes and, consequently, this increases the time before the first nodes die. In the application model that we are considering, certain nodes are expected to be used more often than others as data generators, meaning that on average, their energy consumption will exceed that of the other nodes in the network. As the network progresses into its final stages, these nodes will have the lowest remaining energy and will be avoided as routers. However, this happens too late and these nodes may die prematurely, as they are required to generate traffic very often. In order to improve network lifetime, these nodes should be avoided as routers even in the initial stages of the network.

Because redundancy exists between the coverage of the sensors, each location x is characterized by a sensor set $S(x) \subseteq S$ that is capable of monitoring it. We will denote the total energy of all sensors that have location x within their coverage area as $E(x)$.

$$E(x) = \sum_{s_i \in S(x)} e_i^{res} \tag{8}$$

We can define several cost functions based on $E(x)$ that allow nodes to indicate their unwillingness to route traffic even before their residual energy drops significantly below other nodes in the network. While these coverage-aware costs can be used for the sensor network QoS model considered in this work (i.e., coverage), other methods for determining application QoS-aware costs may be used for sensor network applications that do not conform to this model. In developing an application QoS-aware routing cost, the general goal is to provide information about the importance of the individual sensors to the sensing application.

4.1 Worst Coverage-Based Cost

In some applications, it may be critical that the entirety of the region being monitored is covered as long as possible. In other words, the utility of the application drops significantly as the coverage falls from 100% to just below 100%. For such situations, we define a worst coverage-based cost $C_{wc}(s_i)$

$$C_{wc}(s_i) = \frac{1}{\min_{x \in A(s_i)} E(x)} = \max_{x \in A(s_i)} \frac{1}{\sum_{s_j \in S(x)} e_i^{res}} \tag{9}$$

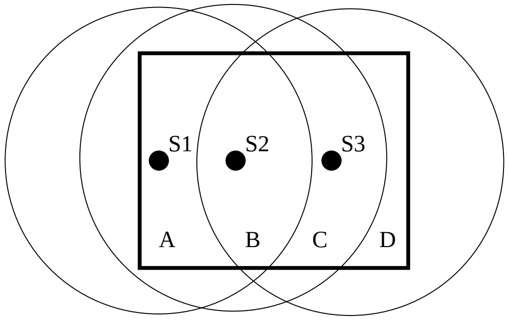

Fig. 2. Example sensor network. Since s_3 is the only sensor that can cover region D, its worst coverage-based cost $C_{wc}(s_i)$ is the highest in the network. The comprehensive coverage-based cost $C_{wc}(s_i)$ gives a more complete encapsulation of a sensor's value to the sensing task and considers the area and redundant energy of each subregion.

This cost assignment method finds the least-covered subregion (in terms of energy) of each node's coverage area and sets the node's cost equal to the inverse of the sum of the energy of the individual sensors capable of monitoring that critical subregion.

Consider the scenario illustrated in Figure 2, where the rectangular area is the region to be monitored and sensors s_1, s_2, and s_3 are capable of monitoring the regions within the circles representing their respective sensing ranges. For simplicity, we assume that all sensors have a single unit of energy. Any point in region A, which we will refer to as x_A, can be covered by 2 sensors – s_1 and s_2. Thus, $E(x_A) = 2$ and similarly, $E(x_B) = 3$, $E(x_C) = 2$, and $E(x_D) = 1$. Sensor s_1 can monitor regions A and B and since the coverage in region A is the poorest in terms of total energy, s_1's cost is set to $C_{wc}(s_1) = \frac{1}{E(x_A)} = \frac{1}{2}$. Similarly, $C_{wc}(s_2) = \frac{1}{2}$ and $C_{wc}(s_3) = 1$.

Note that several sensors, whose least redundantly covered portions of the monitored region consist of overlapping portions, will have identical application costs, regardless of their individual residual energy. This follows the intuition of our design, since these sensors are equally effective at monitoring this critical region of the environment.

4.2 Comprehensive Coverage-Based Cost

In some scenarios, the utility of a sensor network application may degrade gracefully with the amount of area that is covered. To account for this, we propose another routing cost $C_{cc}(s_i)$ that considers the comprehensive coverage in the regions that a sensor can monitor instead of the single least-covered region. This comprehensive coverage-based cost is set as a weighted sum of $1/E(x)$, weighted by the area of each subregion. In other words, to obtain $C_{cc}(s_i)$, we integrate the inverse of $E(x)$ over s_i's coverage region.

$$C_{cc}(s_i) = \int_{A(s_i)} \frac{dx}{E(x)} = \int_{A(s_i)} \frac{dx}{\sum_{s_j \in S(x)} e_j^{res}} \qquad (10)$$

Again, consider the scenario illustrated in Figure 2. Sensor s_1 will set its cost as $C_{cc}(s_1) = \int_{A(s_i)} \frac{dx}{E(x)} = \int_A \frac{dx}{2} + \int_B \frac{dx}{3} = \frac{area(A)}{2} + \frac{area(B)}{2}$. Similarly, $C_{cc}(s_2) = \frac{area(A)}{2} + \frac{area(B)}{3} + \frac{area(C)}{2}$ and $C_{cc}(s_3) = \frac{area(B)}{3} + \frac{area(C)}{2} + \frac{area(D)}{1}$. This comprehensive coverage-based routing cost provides a more balanced view of a node's importance to the sensing task.

4.3 Combining Several Cost Functions

So far, we have proposed two coverage-aware cost functions that capture the importance of individual nodes to the sensing of the environment. However, the usefulness of sensors is not limited to their ability to sense the environment and generate data; they are useful for routing the data of other sensors as well. While the objective of our proposed costs is to use the sensors that are not important as data generators more liberally as routers, some combination of the coverage-aware costs proposed in this work and a connectivity cost could ensure that these sensors are not used too liberally. Consider a network in which a number of nodes that can serve as routers but do not have any sensing capabilities are deployed in addition to the microsensors that we have considered thus far. Using $C_{wc}(s_i)$ and $C_{cc}(s_i)$ as routing costs, these nodes will be assigned a cost of 0. Large amounts of traffic will be routed through these nodes and sent as far toward the data sink as transmission ranges permit, even when large distances between these nodes and the data sink cause energy inefficient transmissions. Clearly, this routing strategy is not optimal. In an energy efficient solution, these router-only nodes would be used more conservatively as routers and a greater portion of their energy would be saved for use in the later stages of the network.

Thus, we propose the use of a routing cost that considers a node's importance as a router as well as a data source. Here, we simply use the energy-aware routing cost $C_{ea}(s_i)$ to help balance the importance of a node. It should be noted that this is a very coarse approximation of the importance of a node to maintaining connectivity, but it is typically a closer approximation than the coverage-aware routing costs. In future work, we plan to develop a connectivity cost that measures the importance of individual sensors in routing data and maintaining good network connectivity.

We considered several methods for combining the energy-aware routing cost and the comprehensive coverage-aware routing cost, including the weighted arithmetic mean, the weighted geometric mean, and the weighted maximum. Simulation results have shown that using the maximum value of the worst coverage-based cost and a weighted value of the energy-aware cost is most effective in extending network lifetime with 100% coverage.

$$C(s_i) = \max(C_{wc}(s_i), \beta C_{ea}(s_i)) \tag{11}$$

Similarly, the use of the maximum value of the comprehensive coverage-based cost and a weighted value of the energy-aware cost is effective in providing long network lifetimes with graceful degradation.

$$C(s_i) = \max(C_{cc}(s_i), \beta C_{ea}(s_i)) \tag{12}$$

In each case, the parameter β can be optimally tuned to maximize network lifetime, as we will show in Section 6.3.

5 DAPR - Distributed Activation with Predetermined Routes

We have designed a simple distributed protocol called DAPR (Distributed Activation with Predetermined Routes) that integrates the services of sensor selection and route discovery. Most architectures proposed for use in coverage-preserving wireless sensor network applications use a modular approach where sensor selection and routing are performed independently. Even in those that use an integrated approach (e.g., [7]), the integration is rather loose, as the sensor selection algorithm considers the effect of the potential routers, but the routers are not chosen with any consideration of the sensor selection algorithm. In the proposed DAPR protocol, route discovery and sensor selection are performed separately, but decisions made in each process are influenced by the other. The premises for the design of DAPR are twofold — that sensors critical to the sensing applications as data generators should be avoided as routers and that the selection of a sensor for active sensing affects its potential routers as well as the sensor itself.

In DAPR, finite-length queries, which are triggered by the sending of Query packets, are processed for a predetermined query length by a subset of the sensors available in the network. Before the query is processed, the network undergoes a Route Discovery Phase, which is followed by a Sensor Selection Phase. Upon completion of the Sensor Selection Phase, sensors process the query and provide data to the querying node for the duration of the query. In previous work [1], we considered a round-based approach where a data sink collects data for long periods of time, and sends Round Start messages periodically so that roles are updated regularly and energy is balanced throughout the lifetime of the network. The single-query approach proposed here is simply a more generic version of this protocol and can be made equivalent by requiring queries to be sent at the correct interval.

During the Route Discovery Phase, the Query packets are broadcast throughout the network, with one copy of the packet broadcast by each node, so that a spanning tree, rooted at the data sink, is formed. As the packets are flooded throughout the network, each node updates a cost field within the packet, adding the cost of the link to its parent node. Routing costs such as those proposed in the previous section are assigned to individual sensors, and the cost of a link is a weighted sum of the effort that each sensor must put forth to transfer the data. Specifically, the cost of a link is calculated as

$$C_{link}(s_i, s_j) = C(s_i)e_{ij}^{tx} + C(s_j)e_{ij}^{rx} \tag{13}$$

where e_{ij}^{tx} represents the energy that is required by s_i to transmit a bit to s_j and e_{ij}^{rx} represents the energy that is required by s_j to receive a bit from s_i. The cumulative cost of a sensor's route is

$$C_{route}(s_i) = \sum_{(s_j,s_k)\in p(s_i)} C_{link}(s_j, s_k) \tag{14}$$

where $p(s_i)$ represents the set of links along the chosen optimal path from s_i to the sink that minimizes $C_{route}(s_i)$.

After the Route Discovery Phase, each sensor must decide in the Sensor Selection Phase whether or not it is necessary to actively sense and generate data. After initially assuming that it will actively sense and generate data to process the query, each sensor will attempt to deactivate itself if possible by sending a deactivation beacon. To ensure that sensors with the highest route costs are given the highest priority to deactivate, each node backs off before broadcasting its deactivation beacon, with backoff delays set according to a decreasing function of the route costs. The intuition behind prioritizing sensors based on route costs is based on the fact that a sensor's activation affects its potential routers as well as itself. If, after its backoff delay, a sensor infers that its coverage region is entirely covered by its neighbors that have not yet sent a deactivation beacon, the sensor sends its deactivation beacon, informing its neighbors that it has decided not to actively sense or generate data. It should be noted that this deactivation is for sensing and data generation purposes only. A node that sends a deactivation beacon must remain available for routing purposes since routes have already been determined by this time.

Implementation Issues

The calculation of our proposed coverage-aware routing costs assumes that nodes have location information of neighboring nodes with redundant coverage regions. This information can be exchanged between neighbors after being obtained through GPS or any number of proposed location estimation algorithms in the current literature [12], [13], [14]. Since DAPR was designed for networks of static sensor nodes, location updates must be performed only a single time at the beginning of network operation, or very infrequently in the worst case. Note that the need for location information is not a drawback of the DAPR protocol or the proposed coverage-aware routing costs specifically – this information is necessary in any coverage-preserving protocol. Also, very loose time synchronization is required so that nodes can identify the beginning and end of query periods.

The coverage-aware routing costs also depend on information about the residual energy of neighboring nodes. This information can be conveyed within the Query messages that are forwarded. Before forwarding these messages, which each node should do once per query, a node simply fills in a field in the packet header that is reserved for residual energy information. Since a node must know its own routing cost before forwarding a Query message, it must calculate this value from information obtained during the previous query. As long as the query length is not so long that nodes may use a significant portion of their initial energy during a single query, the residual energy information should not be too stale to calculate near-optimal routes. Alternatively, two packets could be sent during the Route Discovery Phase – the first containing only residual energy information and the second containing route cost information.

We have assumed that nodes are able to begin the dissemination of a query immediately after the data sink broadcasts the initial Query packet. In practice, this means that sensor nodes must listen to the channel in an idle listening mode until receiving these packets. Since it has been shown that power consumption in the idle listening mode is typically comparable to that in the receive mode, this can severely impact network lifetime. If moderate delays are acceptable, then a low power wakeup system may be used to inform nodes about a predetermined time at which the Route Discovery Phase will start [15]. However, idle listening during the Route Discovery and Sensor Selection Phases is unavoidable, as sensors do not know when their neighbors will send the Query messages and deactivation beacons. Since the Route Discovery and Sensor Selection Phases are expected to be very short compared to the query length, this will not greatly impact the energy efficiency of DAPR. Also, this is not a requirement of the coverage-aware routing costs specifically, and would be required under a similar protocol using other routing costs. For short-lived queries where the Route Discovery and Sensor Selection Phases contribute significant overhead in terms of energy consumption, DAPR should not be used.

For normal network operation during the processing of the query, we assume that a schedule-based MAC protocol is used so that idle listening does not contribute significantly to overall energy consumption. The development of such a MAC protocol is beyond the scope of this work, but the reader is referred to [16], [17] for some examples.

The determination of the existence and size of overlapping coverage regions during the calculation of the proposed routing costs and decisions concerning deactivation can potentially be very computationally intensive. Our implementation uses an approximation in which sensors create a grid of locations within their sensing ranges and, point-by-point, observe the redundancy of their neighbors. However, for the deactivation decision, any of the coverage preserving rules described in the current literature could be used in place of this method [5], [6].

The deactivation beacons may be sent over a single hop if it is assumed that the transmission range is at least twice as great as the sensing range. If this assumption is not valid, the beacons must be forwarded through controlled flooding until they reach all sensors that redundantly cover at least some portion of the sending sensor's coverage region (i.e., those within twice the sensing range).

6 Simulations and Analysis

In this section, we present simulation results measuring the performance of the proposed coverage-aware routing costs and the DAPR protocol. The simulations in this section were performed using Matlab, and they focus on the routing and application layer while simplifying MAC and physical layer implications.

In these simulations, a data sink sent periodic queries, which were processed by sensors that sent constant bit rate traffic to the sink. We assumed that queries were generated from different locations in the network throughout the network lifetime. This helped to avoid rapid energy drain in the nodes surrounding the data sink.

The energy model that was used in our simulations was similar to that used in [18], in which the energy required by s_i to send a bit to s_j separated by a distance of d_{ij} was

$$e_{ij}^{tx} = E_{elec} + \varepsilon \ d_{ij}^{\alpha} \tag{15}$$

where E_{elec} represents the energy associated with the radio electronics, ε characterizes the power amplification component, and α represents the path loss exponent. The energy required by s_j to receive a bit from s_i was

$$e_{ij}^{rx} = E_{elec} \tag{16}$$

Under ideal conditions (e.g., very high density), power consumption is minimized by sending packets over distances of d^* [19], [20], where

$$d^* = \sqrt[\alpha]{\frac{2E_{elec}}{(\alpha-1)\varepsilon}} \tag{17}$$

Using our power consumption numbers, given in Table 3, the optimal transmission distance d^* was approximately 32 m. While ideal conditions were not seen in our simulations and sensors may send traffic along circuitous routes in order to avoid routing through critical sensors, this distance, along with the geographic size of the networks that were simulated, provides some indication of the amount of routing that must be performed on the data generated within the network.

Scenario	Uniform	Clustered	Video
Mean coverage overlap (Number of sensors)	9.0	8.9	6.4
Standard deviation of coverage overlap (Number of sensors)	3.0	6.6	2.2

TABLE 2
Coverage overlap statistics for the three simulated deployment scenarios.

In these simulations, we considered three deployment scenarios. The first was a uniform deployment scenario, in which sensor locations were selected uniformly from a circular deployment region. In this scenario, coverage nonuniformities were generally not very severe. While the coverage-aware costs were not designed for such networks, we include analysis of their performance in these types of deployments for thoroughness. The second scenario that we considered was a clustered deplyment scenario, in which small groups of sensors were deployed in a normal distribution around a number of locations chosen randomly from within the network. In this scenario, more coverage nonuniformities existed as a result of deployment nonuniformity. The third scenario was a video network, in which cameras were mounted in a grid deployment on four walls, each of which was required to be monitored at all times. Each camera was randomly tilted horizontally and vertically between -45 and 45 degrees. The simulations of this scenario helped to measure the performance of the coverage-aware routing costs when the sensors' physical proximity to each other did not necessarily determine their coverage overlap. Examples of deployment patterns for each scenario are given in Figure 3. The coverage nouniformities are summarized in Table 2, which shows the mean and standard deviation of the coverage overlap throughout the region to be monitored. While both the uniform and clustered scenarios have an average overlap of about 9 sensors, the standard deviation is more than twice as high in the clustered scenario.

The rest of the parameters used in our simulations are summarized in Table 3. All simulation results were averaged over 25 trials. It should be noted that we did not compare our results against other coverage-preserving rules that exist in the literature. The reason for this is that the major contributions of our work are the incorporation of coverage information into the routing protocol and the priority (as opposed to the rules) for sensor selection. In fact, the rules by which nodes determine whether or not they need to actively sense and generate data to preserve coverage in the network is not very important. Any of the coverage-preserving decision rules in the literature could be used with our protocol.

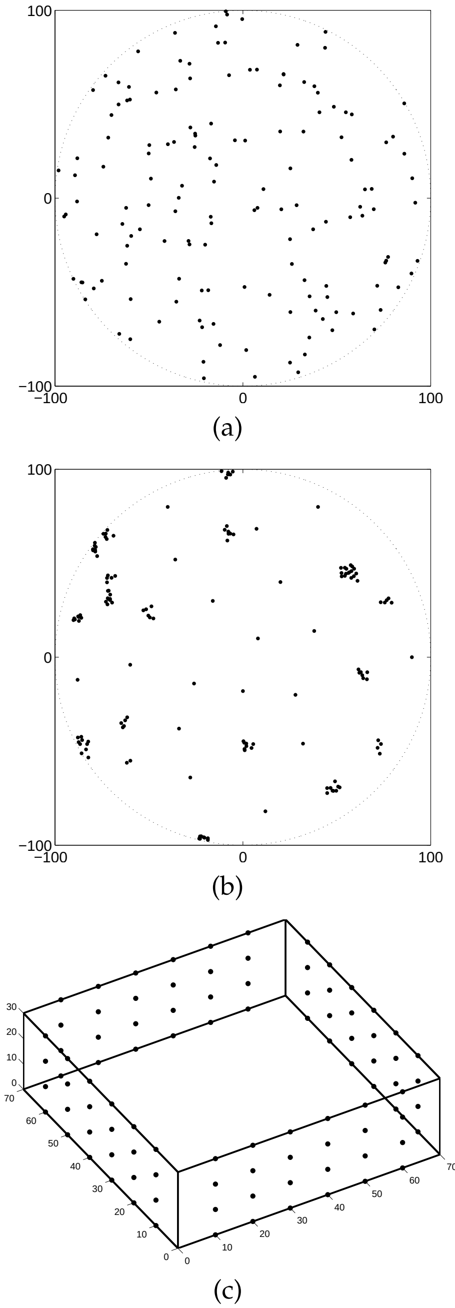

Fig. 3. Example sensor deployment patterns for the uniform deployment scenario (a), clustered deployment scenario (b), and video scenario (c).

6.1 Performance of Coverage-Aware Routing Costs

In this section, we analyze the performance of our proposed coverage-aware routing costs as alternatives to traditional energy-aware routing, where $C(s_i) = C_{ea}(s_i)$, and minimum power routing, where $C(s_i) = 1$, using the DAPR protocol for sensor and router selection. All networks in this section consisted of 150 sensor nodes.

Figure 4 shows the coverage degradation over time for the uniform deployment scenario. Although the coverage-aware routing costs were not designed for such

Parameter	Value
Packet Size	20 $bytes$
Packet Rate	1 $packet/sec$
α	2
E_{elec}	50 nJ/bit
ε	100 $pJ/bit/m^2$
Query Length	24 hr
Initial Node Energy	1000 J
Sensing Range (Uniform, Clustered)	25 m
Deployment Radius (Uniform, Clustered)	100 m
Surveillance Radius (Uniform, Clustered)	90 m
Room Width (Video)	70 m
Room Height (Video)	30 m
Sensor Spacing (Video)	10 m
Sensor Field of View (Video)	30 degrees

TABLE 3
Default simulation parameters.

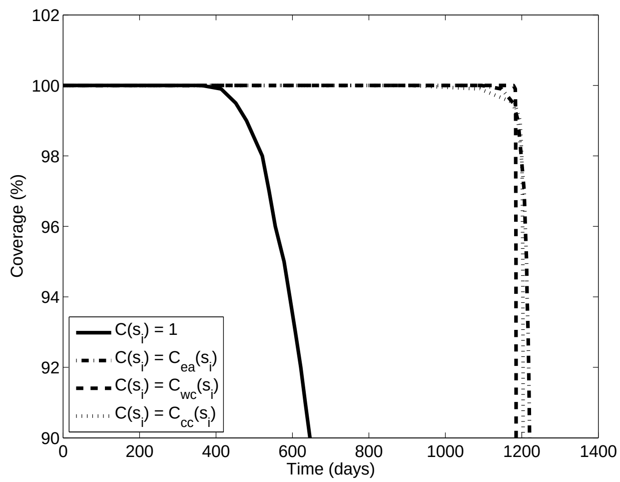

Fig. 4. Coverage degradation over time for different routing costs in the uniform deployment scenario.

networks in which node redundancy is approximately equivalent throughout the network, it can be seen that the coverage-aware routing costs perform very similar to the energy-aware cost, and even slightly better. From this plot and the results summarized in Table 4, we can see that the lifetime before the first break in coverage is highest for the worst-coverage-based routing cost, giving an improvement of 7% over the energy-aware routing cost. Networks using the comprehensive coverage-based cost, which was designed so that coverage degrades more gracefully, were the last to drop below 98%, although the gain over the energy-aware routing cost was minimal in this scenario.

Figure 5 and Table 5 present the results for the clustered deployment scenario. Because coverage is less uniform throughout the network, the gains that can be obtained from the use of the coverage-aware routing costs are higher than in the case of the uniform deployment scenario. The worst coverage-based routing cost gives an improvement of 48% over the energy-aware routing cost in terms of lifetime before the first break in coverage. The comprehensive coverage-based cost gives an improvement of 49% in lifetime before coverage drops below 98% over the energy-aware routing cost.

$C(s_i)$	1	$C_{ea}(s_i)$	$C_{wc}(s_i)$	$C_{cc}(s_i)$
100% coverage lifetime (days)	362	1094	1178	904
98% coverage lifetime (days)	521	1198	1184	1200

TABLE 4
Simulation results for different routing costs in the uniform deployment scenario.

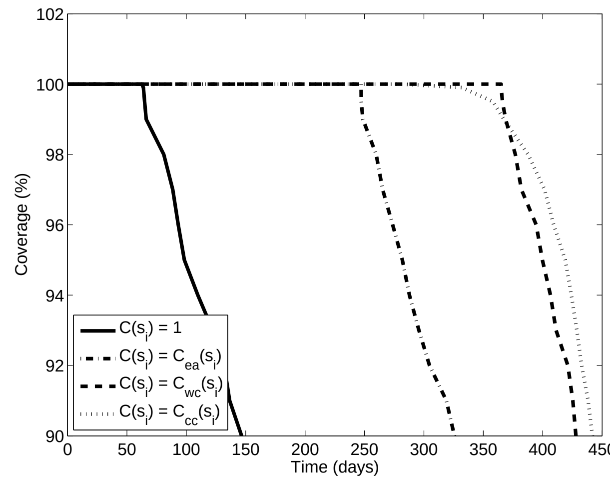

Fig. 5. Coverage degradation over time for different routing costs in the clustered deployment scenario.

Figure 6 and Table 6 present the results for the video scenario. The worst coverage-based cost gives a significant gain in network lifetime before the first break in coverage, increasing lifetime by 24%. However, the comprehensive coverage-based cost performs very poorly in this scenario. We suspect that the reason for this is that this cost does not consider the utility of a node as a router, but rather as a sensor only. Nodes that should be kept alive for routing purposes may be used too liberally, causing them to die and forcing other sensors in the region to use suboptimal routes for the remainder of the network lifetime. This is not a problem in the uniform and clustered deployment scenarios since a node's importance as a sensor and as a router are both

$C(s_i)$	1	$C_{ea}(s_i)$	$C_{wc}(s_i)$	$C_{cc}(s_i)$
100% coverage lifetime (days)	62	247	365	376
98% coverage lifetime (days)	81	260	377	388

TABLE 5
Simulation results for different routing costs in the clustered deployment scenario.

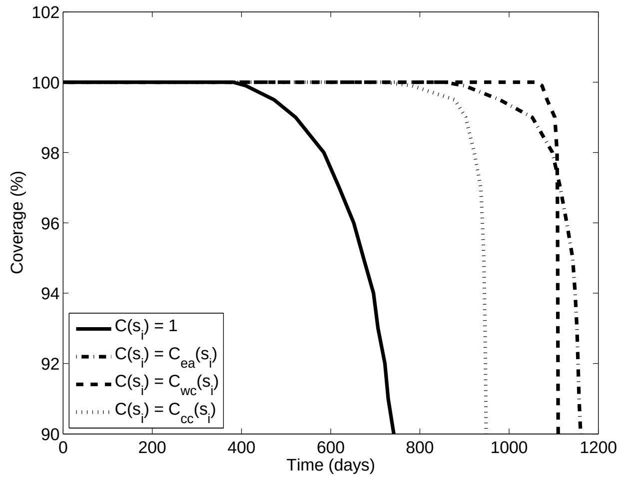

Fig. 6. Coverage degradation over time for different routing costs in the video scenario.

$C(s_i)$	1	$C_{ea}(s_i)$	$C_{wc}(s_i)$	$C_{cc}(s_i)$
100% coverage lifetime (days)	381	855	1063	717
98% coverage lifetime (days)	585	1097	1108	921

TABLE 6
Simulation results for different routing costs in the video scenario.

tied to its location. One way to avoid this problem for video networks is to use a combined routing cost, as discussed in Section 6.3.

6.2 Effect of Sensor Selection Criteria

In this section, we explore the effect of the sensor selection criteria when using the worst-coverage based routing cost. Recall that sensors are deactivated by sending a deactivation beacon to neighboring sensors after a backoff timer expires. In these simulations, we compare network lifetime when setting the backoff timer according to three different criteria – randomly, based on the sensor's individual cost, and based on the sensor's cumulative route cost, given in Equation 14. As the activation or deactivation of a sensor affects its routers as well as itself, we expect lifetime to be highest when setting the backoff timer according to the cumulative route cost. As shown in Table 7, choosing sensors based on their cumulative cost improves network lifetime over using sensors' individual costs by a modest 8% in the uniform deployment scenario, almost nothing in the clustered deployment scenario, and a more significant 33% in the video scenario. These values can be explained by the fact that nearby sensors typically have very similar routes to the data sink. For this reason, when selecting which sensors to deactivate among multiple nearby sensors that cover the same region, the choice will probably affect only the sensors being deactivated,

Selection Criteria	Random	Individual Cost	Cumulative Routing Cost
Uniform	1036	1088	1178
Clustered	364	365	365
Video	818	800	1063

TABLE 7
Network lifetime (days) before first coverage break when using the worst coverage-based routing cost with different selection criteria.

but few or none of the routers, since they are probably the same for all sensors under consideration. This is the case in the clustered deployment scenario and to less of an extent, in the uniform deployment scenario. However, in the video scenario, two sensors that cover the same region may have very dissimilar routes to the base station. In this situation, the choice of which sensor to deactivate will affect different groups of sensors. Thus, the gain in network lifetime is highest in this scenario.

6.3 Combining Routing Costs

In this section, we explore the effectiveness of combining the coverage-aware routing costs with the energy-aware routing cost in order to account for both coverage and connectivity requirements. We simulated similar networks as in the previous sections as well as heterogeneous networks, in which 150 sensors capable of sensing the environment and generating data for the data sink were deployed along with additional nodes that could only be used to route data (50 for the uniform and clustered scenarios and 32 for the video scenario). We ran simulations in which we set the nodes' routing costs to $C(s_i) = \max(C_{wc}(s_i), \beta C_{ea})$ and others in which we set the nodes' routing costs to $C(s_i) = \max(C_{cc}(s_i), \beta C_{ea})$, as described in Section 4.3, and we tuned β to maximize network lifetime.

Network lifetime (for 100% coverage) when setting the routing cost as the combined cost $C(s_i) = \max(C_{wc}(s_i), \beta C_{ea})$ are summarized in Table 8. The results show that network lifetime before coverage degrades below 100% is typically maximized or very nearly maximized when β is set around 0.25 in these scenarios. The improvement is most significant in the heterogeneous networks since the router-only nodes' value is most misrepresented by the coverage-aware cost in this case. In the heterogeneous networks, the use of the combined routing cost with this value of β improves network lifetime by 20% over the use of the worst coverage-based routing cost and by 17% over the use of the energy-aware routing cost for the uniform deployment scenario. For the clustered deployment scenario, these numbers grow to 21% and 43%, respectively. In the video network, the combined cost improves lifetime by 8% over the use of the worst coverage-based routing cost and by 20% over the use of the energy-aware routing cost.

β	0 C_{wc}	0.05	0.25	0.5	1 C_{ea}
Uniform	1178	1192	1224	1093	1094
Clustered	365	365	360	245	247
Video	1062	1082	1106	853	855
Uniform (add. routers)	1368	1572	1635	1549	1402
Clustered (add. routers)	476	567	576	525	403
Video (add. routers)	1214	1299	1306	1083	1083

TABLE 8
Network lifetime (days) before first coverage break for heterogeneous networks when using a combination of worst coverage-based energy-aware routing costs.

β	0 C_{cc}	$20m^2$	$100m^2$	$400m^2$	$\beta \to \infty$ C_{ea}
Uniform	1203	1210	1240	1212	1212
Clustered	419	419	418	346	281
Video	943	1064	1163	1142	1143
Uniform (add. routers)	1392	1554	1670	1646	1645
Clustered (add. routers)	487	569	629	561	493
Video (add. routers)	1232	1595	1747	1629	1629

TABLE 9
Network lifetime (days) before coverage drops below 98% for heterogeneous networks when using a combination of comprehensive coverage-based energy-aware routing costs.

The effects of the combined cost are less dramatic in the scenarios that contain only sensing-capable nodes, as in the simulations of the previous sections. The network lifetime improvement is about 4% for the uniform and video scenarios when using a β value of 0.25. Meanwhile, the clustered scenario does not benefit at all from the use of the combined cost.

Network lifetime (for 98% coverage) for each scenario when setting the routing cost as $C(s_i) = \max(C_{wc}(s_i), \beta C_{ea})$ are summarized in Table 9. The results show that network lifetime before coverage degrades below 98% is typically maximized or nearly maximized when β is set at $100m^2$ for each scenario. Again, the impact is greatest in the heterogeneous networks. The use of the combined routing cost with this value of β improves network lifetime by 20% over the use of the comprehensive coverage-based routing cost and by 2% over the use of the energy-aware routing cost for the uniform deployment scenario. In other words, we gain very little in using the combined cost over using the energy-aware cost alone. For the clustered deployment scenario, these numbers grow to 29% and 28%, respectively. For the video scenario the combined cost improves network lifetime by 42% over the use of the comprehensive coverage-based routing cost and by 7% over the use of the energy-aware routing cost.

As in the case of the worst coverage-based cost, the improvements are not as great in the networks containing only sensing-capable nodes.

7 Centralized Role Selection

If global information about the network topology, sensing capabilities, initial energy of each node, and base station locations throughout the lifetime of the network is available, it is possible to optimize the sensor scheduling and data routing so that network lifetime is maximized [21], [22], [23], [24]. In this section, we present a centralized approach incorporating an optimization program that is used as a baseline to measure the performance of DAPR.

In Section 3, we described several constraints that we are subjected to in our sensor network model. If all cover sets could be enumerated, a linear program with constraints given by Equations 3, 4, and 5 and a goal of maximizing L (given in Equation 6) could be run to find the maximum achievable network lifetime for a given network scenario. However, the problem of finding all cover sets may be computationally infeasible for networks consisting of a large number of sensors and a large amount of sensing redundancy. Rather than enumerating all of these cover sets, we would like to find a subset of C whose optimal scheduling would yield a similar lifetime as the optimal scheduling of C. Berman et al. showed how the calculation of the cover sets can be accomplished simultaneously with the scheduling of the sets for a single-hop sensor network [25], using an algorithm proposed by Garg and Könemann [26]. The Garg-Könemann algorithm yields a scheduling solution whose lifetime is arbitrarily close (within some factor ϵ) to the optimal value while considering only a subset of the cover sets. Once the subset of cover sets have been chosen and scheduled, the schedule may be modified through a linear program optimization, and the rescheduled lifetime may match or nearly match the true upper bound when a small enough value of ϵ is used.

While Berman et al.'s work looked at the scheduling of cover sets in single hop sensor networks, we are interested in the maximum obtainable lifetime of multihop wireless sensor networks. We wish to develop a schedule that determines the length of time that each cover set should be used and the fraction of traffic that each sensor should route toward each of its neighbors. In our centralized approach, we initially generate a diverse group of cover sets using the single hop scheduling algorithm of Berman et al. Once we have determined a large enough group of cover sets, we run the optimization program with Equations 3, 4, and 5 as constraints and the maximization of L in Equation 6 as a goal. Note that the lifetime resulting from this optimization program is not necessarily the true upper bound on network lifetime since all feasible cover sets have not been enumerated. However, the lifetime does give us a good baseline against which to test our distributed protocol.

In some situations (e.g., if this program is being used to plan the actual operation of a wireless sensor network), it may not be reasonable to expect that the values of q_m in Equation 5 (the fraction of queries originating from each data sink location) are known a priori. However, we are performing these optimizations for the purposes of providing a baseline for the performance of our proposed algorithm. Thus, we use the fraction of time that the data sink was located in each of its locations in the simulations and proceed with these numbers to find the maximum lifetime that could have been obtained with ideal sensor and route selection.

7.1 Comparison of DAPR and Centralized Approach

Our simulation results of DAPR with the worst coverage-based cost are compared with the lifetime obtained via the centralized approach in Figures 7(a), 7(b), and 7(c) for the uniform, clustered, and video scenarios, respectively. We compare several approaches here:

1) a typical non-integrated approach - setting node costs as $C_{ea}(s_i)$ (energy-aware routing) with sensor selection based on the individual sensors' costs,
2) the non-integrated approach, but setting node costs as $C_{wc}(s_i)$,
3) DAPR, using energy-aware routing cost $C_{ea}(s_i)$,
4) DAPR, using the worst-coverage routing cost $C_{wc}(s_i)$, and
5) DAPR, using a combination of the energy-aware routing cost and the worst-coverage routing cost, as described in Section 4.3, with β set to 0.25.

In the uniform scenario, the use of the combined worst-coverage and energy-aware cost with DAPR gives a total network lifetime gain of 14% over the non-integrated approach, closing the gap with the centralized solution by 56%. The results for the clustered deployment scenario show that DAPR with the worst-coverage routing cost performs especially well in this scenario, improving lifetime by 56% and closing the gap between the non-integrated approach and the centralized solution by 77%. Most of this improvement in this scenario is due to the use of the coverage-aware routing cost rather than the selection of sensors based on the cumulative route cost. Finally, the use of DAPR with the combined worst-coverage and energy-aware cost in the video scenario gives a total network lifetime gain of 50% and closes the gap between the non-integrated approach and the centralized solution by 76%. Most of the improvement in this case comes from the selection of sensors based on the cumulative route cost.

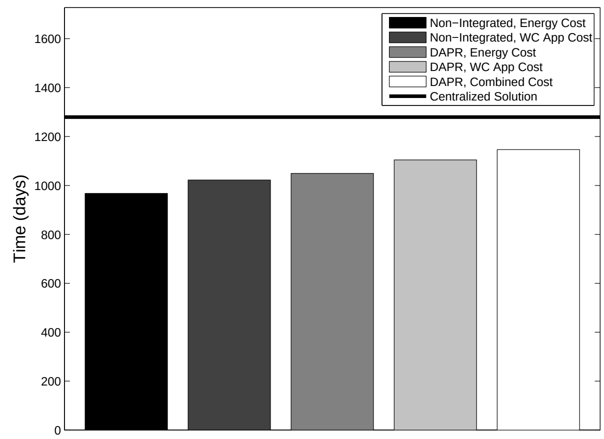

(a) Uniform Scenario

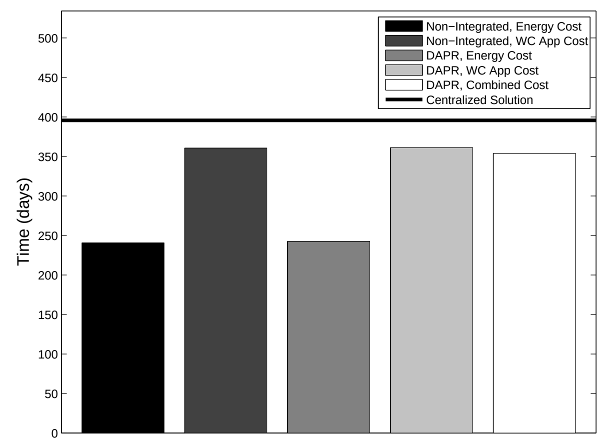

(b) Clustered Scenario

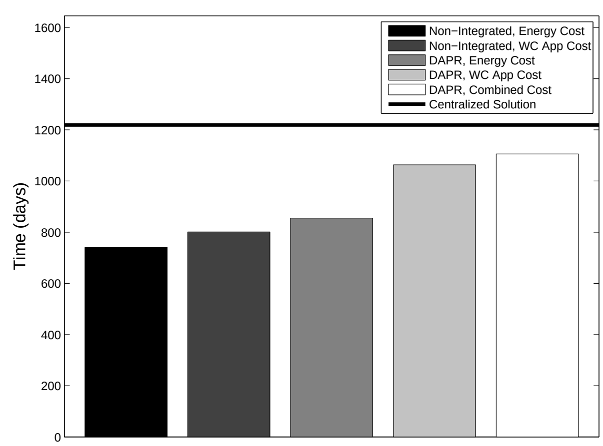

(c) Video Scenario

Fig. 7. Comparison of DAPR with the centralized approach.

8 Conclusions and Future Work

In this work, we have proposed the use of coverage-aware routing costs and a distributed, integrated protocol for sensor selection and routing for use in coverage-preserving wireless sensor network applications. Our approach integrates application layer and routing layer functionality in two ways - by assigning routing costs with awareness of each node's importance to the application goals, and by choosing sensors that participate in the application with the knowledge of the effects that this has on potential routers. The proposed worst coverage-based cost aims to maintain 100% coverage for the maximum lifetime by finding each node's worst-covered subregion and assigning costs inversely proportional to the energy of nodes covering that region. The comprehensive coverage-based cost gives a more balanced interpretation of a node's value to the sensing task by considering all – not just the worst-covered – subregions. Because these routing costs avoid the nodes that will be used most often as data generators even in the early stages of the network, their use can significantly improve sensor network lifetime. Our simulation results have shown that the gains in network lifetime from using these coverage-aware costs become highest when many nonuniformities exist in the sensing redundancy (e.g., the clustered deployment scenario).

In addition to the advantages of using a coverage-aware routing cost, DAPR considers the effect that sensor selection has on potential routers by selecting sensors to actively sense and generate data based on their cumulative route costs. Our simulation results have shown that the gains in network lifetime from using this approach become highest when sensing overlap between sensors is not as directly tied to physical proximity (e.g., the video scenario).

We have compared the lifetime achieved by the DAPR protocol with that achieved with a centralized approach using a large-scale optimization program. Results show that DAPR can significantly close the gap between existing non-integrated and non-coverage-aware approaches and the lifetime of the centralized approach.

In this work, we considered a sensing model in which sensors make a simple decision of whether to turn on or off depending on the current quality of coverage in their neighborhood. One aspect of our future work is to develop similar coverage-aware routing costs for other sensing models, including

- the CEO problem [27], in which a group of sensors obtain noisy measurements of the same process and send measurements independently to the data sink, where the process must be reconstructed subject to some distortion criteria,
- the reconstruction of a data image, where application quality is measured by the signal-to-noise ratio of the reconstructed signal,
- edge detection, where application quality is measured by the uncertainty in the edge approximation, and
- target tracking, where multiple sensors must be able to detect and estimate the range of a target within their coverage region.

The design of application QoS-aware routing costs for these applications, as well as for networks consisting of multi-mode sensors, is a challenging task.

References

[1] M. Perillo and W. Heinzelman, "DAPR: A protocol for wireless sensor networks utilizing an application-based routing cost," in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2004.

[2] F. Ye, G. Zhong, J. Cheng, S. Lu, and L. Zhang, "PEAS: A robust energy conserving protocol for long-lived sensor networks," in *Proceedings of the Twenty-Third International Conference on Distributed Computing Systems (ICDCS)*, 2003.

[3] R. Iyer and L. Kleinrock, "QoS control for sensor networks," in *Proceedings of the IEEE International Conference on Communications (ICC)*, 2003.

[4] T. Yan, T. He, and J. Stankovic, "Differentiated surveillance for sensor networks," in *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.

[5] D. Tian and N. Georganas, "A node scheduling scheme for energy conservation in large wireless sensor networks," *Wireless Communications and Mobile Computing Journal*, vol. 3, no. 2, pp. 271–290, Mar. 2003.

[6] X. Wang, G. Xing, Y. Zhang, C. Lu, R. Pless, and C. Gill, "Integrated coverage and connectivity configuration in wireless sensor networks," in *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.

[7] H. Gupta, S. Das, and Q. Gu, "Connected sensor cover: Self-organization of sensor networks for efficient query execution," in *Proceedings of the Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2003.

[8] Z. Zhou, S. Das, and H. Gupta, "Connected k-coverage problem in sensor networks," in *Proceedings of the Thirteenth International Conference on Computer Communications and Networks (ICCCN)*, 2004.

[9] E. Royer and C. Toh, "A review of current routing protocols for ad-hoc mobile wireless networks," *IEEE Personal Communications*, vol. 6, no. 2, pp. 46–55, Apr. 1999.

[10] S. Singh, M. Woo, and C. Raghavendra, "Power-aware routing in mobile ad hoc networks," in *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 1998.

[11] J. Chang and L. Tassiulas, "Energy conserving routing in wireless ad hoc networks," in *Proceedings of the Nineteenth International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2000.

[12] N. Bulusu, J. Heidemannm, and D. Estrin, "GPS-less low-cost outdoor localization for very small devices," *IEEE Personal Communications*, vol. 7, no. 5, pp. 28–34, Oct. 2000.

[13] A. Savvides, C. Han, and M. Srivastava, "Dynamic fine-grained localization in ad-hoc sensor networks," in *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2001.

[14] D. Niculescu and B. Nath, "Ad hoc positioning system (APS)," in *Proceedings of the Global Telecommunications Conference (GLOBECOM)*, 2001.

[15] C. Guo, L. Zhong, and J. Rabaey, "Low power distributed mac for ad hoc sensor radio networks," in *Proceedings of the Global Telecommunications Conference (GLOBECOM)*, 2005.

[16] K. Arisha, M. Youssef, and M. Younis, "Energy-aware TDMA-based MAC for sensor networks," in *Proceedings of the IEEE Integrated Management of Power Aware Communications, Computing and Networking (IMPACCT)*, 2002.

[17] M. Sichitiu, "Cross-layer scheduling for power efficiency in wireless sensor networks," in *Proceedings of the Twenty-Third International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2004.

[18] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan, “An application-specific protocol architecture for wireless microsensor networks,” *IEEE Transactions on Wireless Communications*, vol. 1, no. 4, pp. 660–670, Oct. 2002.

[19] Y. Chen, E. Sirer, and S. Wicker, “On selection of optimal transmission power for ad hoc networks,” in *Proceedings of the Thirty-Sixth Hawaii International Conference on System Sciences (HICSS-36)*, 2003.

[20] I. Stojmenovic and X. Lin, “Power aware localized routing in wireless networks,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 12, no. 11, pp. 1122–1133, Nov. 2001.

[21] M. Bhardwaj and A. Chandrakasan, “Bounding the lifetime of sensor networks via optimal role assignments,” in *Proceedings of the Twenty-First International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2002.

[22] M. Perillo and W. Heinzelman, “Simple approaches for providing application qos through intelligent sensor management,” *Elsevier Ad Hoc Networks Journal*, vol. 1, no. 2–3, pp. 235–246, 2003.

[23] K. Kalpakis, K. Dasgupta, and P. Namjoshi, “Maximum lifetime data gathering and aggregation in wireless sensor networks,” in *Proceedings of the 2002 IEEE International Conference on Networking (ICN)*, 2002.

[24] F. Ordonez and B. Krishnamachari, “Optimal information extraction in energy-limited wireless sensor networks,” *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 6, pp. 1121–1129, 2004.

[25] P. Berman, G. Calinescu, C. Shah, and A. Zelikovsky, “Power efficient monitoring management in sensor networks,” in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2004.

[26] N. Garg and J. Könemann, “Faster and simpler algorithms for multicommodity flow and other fractional packing problems,” in *Proceedings of the IEEE Symposium on Foundations of Computer Science*, 1997.

[27] T. Berger, Z. Zhang, and H. Viswanathan, “The CEO problem,” *IEEE Transactions on Information Theory*, vol. 42, no. 3, pp. 887–902, May 1996.

PLACE PHOTO HERE

Mark Perillo is a research engineer at Syracuse Research Corporation in Syracuse, NY. He received B.S., M.S., and Ph.D. degrees in Electrical and Computer Engineering from the University of Rochester in 2000, 2002, and 2007, respectively. His research interests lie in the areas of wireless communications, ad hoc and sensor networks, and synthetic aperture radar.

PLACE PHOTO HERE

Wendi Heinzelman is an associate professor in the Department of Electrical and Computer Engineering at the University of Rochester, and she holds a secondary appointment as an associate professor in the Department of Computer Science. Dr. Heinzelman also serves as Dean of Graduate Studies for Arts, Sciences and Engineering at the University of Rochester. Dr. Heinzelman received a B.S. degree in Electrical Engineering from Cornell University in 1995 and M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from MIT in 1997 and 2000, respectively. Her current research interests lie in the areas of wireless communications and networking, mobile computing, and multimedia communication. Dr. Heinzelman received the NSF CAREER award in 2005 for her research on cross-layer architectures for wireless sensor networks, and she received the ONR Young Investigator Award in 2005 for her work on balancing resource utilization in wireless sensor networks. She is an Associate Editor for the IEEE Transactions on Mobile Computing, and she is an Area Editor for ACM Mobile Computing and Communications Review (MC2R). Dr. Heinzelman is currently the Vice Chair for the Systems and Applications track for DCOSS 2009. She is a member of Sigma Xi and the ACM, she is a senior member of the IEEE, and she is co-founder and current co-leader of the N^{2}Women (Networking Networking Women) group.