

Dynamic Node Activation in Networks of Rechargeable Sensors

Koushik Kar, Ananth Krishnamurthy and Neeraj Jaggi

Abstract—We consider a network of rechargeable sensors, deployed redundantly in a random sensing environment, and address the problem of how sensor nodes should be activated dynamically so as to maximize a generalized system performance objective. The optimal sensor activation problem is a very difficult decision question, and under Markovian assumptions on the sensor discharge/recharge periods, it represents a complex semi-Markov decision problem. With the goal of developing a practical, distributed but efficient solution to this complex, global optimization problem, we first consider the activation question for a set of sensor nodes whose coverage areas overlap completely. For this scenario, we show analytically that there exists a simple threshold activation policy that achieves a performance within a factor of $\frac{3}{4}$ of the optimum over all possible policies. We extend this threshold policy to a general network setting where the coverage areas of different sensors could have partial or no overlap with each other, and show by simulations that the performance of our policy is very close to that of the globally optimal policy. Our policy is fully distributed, and requires the sensor nodes to only keep track of the node activation states in its immediate neighborhood. We also consider the effects of spatial correlation on the performance of the threshold activation policy, and the choice of the optimal threshold.

I. INTRODUCTION

A. Background and Motivation

Due to major technological innovations in recent years, development of tiny, low-cost sensor devices has become possible. Such sensor devices can be deployed in large numbers in different environments for monitoring and data gathering purposes [1]. These sensor devices, although cheap, are typically unreliable. Moreover, sensor devices are limited by battery energy. Therefore, a sensing device can remain powered on (and be sensing) only for a limited amount of time, until it runs out of battery energy [12]. In many scenarios, sensors can be recharged, but recharging is often a very slow process and the rate of recharging could be significantly less than the rate of energy depletion during the sensing period. As a result, a sensor could need to spend most of its lifetime in the “off” state, when it is not sensing, but recharging. These factors motivate redundant deployment of sensors to cover the area of interest. Each sensor being unreliable, sensing reliability increases if more number of sensors are sensing the same area at the same time. If larger number of sensors are deployed, it is likely that more number of these sensors would remain charged (and hence can be used for sensing) at any given time. Thus the

K. Kar and N. Jaggi are with the Department of Electrical, Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180, USA (email: koushik@ecse.rpi.edu, jaggin@rpi.edu). A. Krishnamurthy is with the Department of Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, NY 12180, USA (email: krisha@rpi.edu).

This work was supported by the NSF grant ECS-0330203.

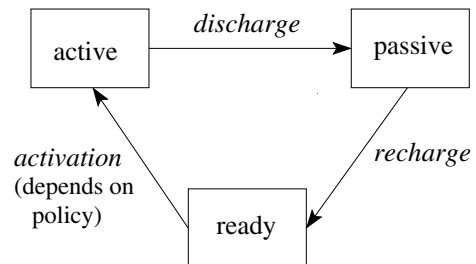


Fig. 1. States and transitions

overall system performance would typically improve (possibly with diminishing returns) with a more redundant deployment of sensors.

We consider the scenario mentioned above, where sensors have been deployed redundantly in the area of interest. We assume that sensor nodes involved in sensing get discharged after a certain duration of time, and need to be recharged till they can start sensing again. We consider the decision problem of when the recharged sensors should be *activated* (i.e., switched on) so as to maximize the long-term utility of the system.

B. Problem Formulation

We assume that sensors are energy-constrained, but rechargeable, and at any instant of time, each sensor could be in one of three states: i) *active* ii) *passive*, or iii) *ready*. In the active state, the sensors are powered on and are sensing. A sensor in the active state suffers a gradual depletion of battery energy, and enters the passive state when its battery is completely discharged. Sensors that are passive are powered off, and are simply recharging their batteries. When its battery is completely charged, the sensor enters the ready state. Sensors in the ready state do not participate in sensing, and wait to get activated. Figure 1 explains the three different states, and the transitions between them.

Let *discharge time* denote the time a sensor spends in the active state, and *recharge time* denote the time a sensor spends in the passive state. In a realistic sensing environment, the discharge and recharge times will depend on various random factors. Sensors can transmit information (resulting in energy usage) on the occurrence of “interesting” events, which may be generated according to a random process. Therefore in our system model, we assume that the discharge and recharge times are random, although we study the special case of deterministic discharge/recharge times as well.

Although a sensor can power itself off during the ready state, it has to wake up periodically and exchange messages with its neighbors to keep track of the system state in its neighborhood. Therefore, in reality, we would expect that energy will be drained even in the ready state, but probably at a fairly steady rate (possibly due to polling its neighbors to check out the system activation state). However, the energy discharge rate in the ready state can be expected to be much slower than the discharge rate in the active state.

We assume that the performance of the system is characterized by a continuous, non-decreasing, strictly concave function U satisfying $U(0) = 0$. More specifically, $U(n)$ represents the utility derived *per unit area, per unit time*, from n active sensors covering an area. Note that different sensors can be located at different points in the overall physical space of interest, and the coverage patterns of different nodes can be different. Therefore, the coverage areas of different sensors will typically be different. This implies that at any time, utilities in different parts of the area of interest can differ significantly from one another.

Note that the strict concavity assumption merely states the fact that the system has diminishing returns with respect to the number of active sensors. As an example of a practical utility function, consider the scenario where each sensor can detect an event with probability p_d . If the utility is defined as the probability that the sensing system is able to detect an event, then $U(n) = (1 - (1 - p_d)^n)$, where n is the number of sensors that are active. Note that this utility function is strictly concave, and satisfies $U(0) = 0$.

We are interested in maximizing the time-average utility of the system. Let \mathcal{A} denote a generic area element in the physical space of interest. Let $n_P(\mathcal{A}, t)$ denote the number of active sensors that cover area element \mathcal{A} at time t , under policy P . The time-average utility under policy P , is given by

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \int_{\mathcal{A}} U(n_P(\mathcal{A}, t)) dA dt. \quad (1)$$

In Euclidean coordinates system, $dA = dx dy$, and $n_P(\mathcal{A}, t) = n_P(x, y, t)$, in the above expression. The decision problem that we consider in this paper is that of finding P so that the objective function in (1) is maximized.

As mentioned before, our decision problem is that of determining how many sensors to activate at any time, from the set of ready sensors. Note that if we activate more sensors, we gain utility in the short time-scale. However, if the number of active sensors is already large, since the utility function exhibits diminishing returns, we may want to keep some of the ready sensors “in store” for future use. In fact, as we see later, the performance results for policies studied in this paper justifies this intuition.

C. Problem Challenges and Basic Approach

The stochastic nature of charging and discharging time of sensors makes the determination of optimal activation policies very hard in a general setting. Further, spatio-temporal correlations imply that at any point in time, the optimal activation policy for a sensor might depend on the history of the states of all

the sensors in the network. Although under specific cases the optimal policies may be formulated as semi-Markov decision problem, determining optimal policies can be computationally prohibitive. Since sensors are energy constrained we seek policies that can be implemented in a distributed manner with minimal information and computational overhead. Therefore, we focus on simple *threshold-based* policies (defined precisely later in the paper) and examine their performance.

To simplify the analysis and obtain fundamental performance insights, we examine the performance of threshold policies for a system of sensors whose coverage areas overlap completely with each other. In this case the objective function in equation (1) reduces to a single integral over the time domain. Under Markovian assumptions we derive tight bounds on the performance of threshold policies for two different lifetime correlation models of the sensor nodes. These performance bounds motivate our study of the performance of threshold based policies in a very general network setting, where coverage areas of different sensors could have partial, complete or no overlap with each other. Results from our extensive simulation studies show that even in this case, the performance of threshold activation policies is very promising.

D. Contribution of this Work

Our problem formulation is novel, and formalizes a practically important problem which has not yet received its due attention. As is the case in reality, we assume that the discharge and recharge times of the sensors are random. We consider two extreme correlation models of the discharge/recharge times of the different sensors: one in which these times are highly correlated, and the other in which these times are independent of one another. For the sake of analytical tractability, we first consider the case where sensor nodes have coverage areas that are completely overlapping. Assuming that the discharge/recharge times are exponential, we formulate the problem as a continuous-time Markov decision problem and provide a procedure for determining the optimal policy. Since the associated computation complexity is significant, we focus on the class of threshold decision policies. Threshold policies yield closed-form expressions, and the optimal threshold policy can be computed efficiently. We show that the time-average utility of the optimal threshold policy is within a factor of $\frac{3}{4}$ of the best possible performance, for both correlation models. Moreover, we show that correlation of the discharge/recharge times of the sensors degrades performance at all threshold values. Through numerical studies, we also show that the performance of the optimal threshold policy is very close to the best achievable performance. Subsequently, we extend our threshold based activation policy to a general network setting where the coverage areas of sensors can overlap only partially. The threshold policies derived for this scenario can be implemented by sensors in a distributed manner, based only on information about the local network state. Through extensive simulation studies, we show that the performance of our policy is very close to that of a globally optimal policy. Therefore, threshold policies allow us to obtain near-optimal solutions to our complex decision problem in an efficient manner.

E. Related Work

To the best of our knowledge, the problem of adaptive activation of rechargeable sensors in a dynamic sensing environment has not been addressed so far in the literature. We outline here some of the important work on energy-efficient MAC and adaptive wakeup of sensors, although these works only consider energy-constrained, but non-rechargeable sensors. A discussion on the importance of energy management in adhoc and sensors networks, along with a description of various performance objectives, is outlined in [12]. Energy-efficient MAC protocols are studied in [5], [6], [14]. The problem of minimizing power consumption during idle times is addressed in [4], [9], [13]. In [3], the authors use occupancy theory to analyze the effect of switching off idle nodes on the network lifetime. In [10], the effects of power conservation, coverage and cooperation on data dissemination is investigated for a particular data sharing architecture.

F. Paper Organization

The paper is structured as follows. In the next section, we consider a system of sensors with fully-overlapping coverage areas, and show that simple threshold policies can achieve near-optimal performance in this scenario. In Section III, we extend the threshold activation policy outlined in Section II to a more general scenario where sensor coverages may not overlap completely.

II. NODE ACTIVATION POLICIES WITH COMPLETE COVERAGE OVERLAP

In this section, we address our node activation decision problem for a system of sensors whose coverage areas overlap completely with each other.

A. System Model and Assumptions

We consider a system of N sensors covering the same area. We assume that the discharge and recharge times of each sensor are random variables.

Assumption 1: The discharge time and recharge time of any sensor are exponentially distributed with means $1/\mu_1$ and $1/\mu_2$ respectively. Moreover, $\mu_1 \geq \mu_2$.

Assumption 2: The energy level of a sensor does not change in the ready state.

The exponential model of the discharge/recharge times is assumed for analytical tractability. Moreover, the optimal policies under this assumption depend only on the number of sensors in the different states in the system, and not their exact energy levels. Without Markovian properties, the system can be very difficult to analyze, and implementing the optimal decision policies (if they can be obtained) would require more detailed system information and additional overhead. The assumption $\mu_1 \geq \mu_2$ is based on the observation that the recharging process in batteries is typically slower than the discharging process.

Assumption 2 basically states that a sensor remains in the fully charged state as long as it remains in the ready state. In reality, we would expect that energy will be drained even in the

ready state, but probably at a fairly steady rate, as discussed earlier. Since the energy discharge rate in the ready state can be expected to be much slower than the discharge rate in the active state, it is not considered in our current analysis.

B. Problem Statement

Let $n_P(t)$ denote the number of sensors in the active state at time t under policy P . Since the coverage areas of all sensors are completely overlapping, the optimization problem can be posed as that of finding a policy P that minimizes $\bar{U}(P)$, where $\bar{U}(P)$ is defined as

$$\bar{U}(P) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t U(n_P(t)) dt. \quad (2)$$

We assume that switching decisions can be taken at any instant of time. Clearly, these decisions would need to be taken only when the state of the overall system changes, i.e., the number of the sensors in the active, passive or ready states changes. In other words, these decisions need to be taken when some sensor makes a transition from the active to the passive state, or some sensor makes a transition from the passive to the ready state. It is worth noting here that although we will address our problem in this case from the perspective of a centralized decision maker, the decision policy that we develop can easily be implemented in a decentralized manner.

C. Sensor Lifetime Models

As we discuss later in more detail, the performance of decision policies depend considerably on how the discharge and recharge times of sensors are correlated. We consider two different correlation models of the discharge/recharge times of the different sensors:

i) *Independent Lifetime (IL) Model:* In this model, sensors are activated independently of one another, and all discharge and recharge times of all sensors are mutually independent.

ii) *Correlated Lifetime (CL) Model:* In this model, the discharge times of all sensors entering the active state at the same time is the same. Similarly, the recharge times of all sensors entering the passive state at the same time is the same. The discharge (recharge) times of sensors entering the active (passive) state at different times are independent of each other.

The two correlation models can be practically motivated in the following way. If the data transmission by a sensor (which is often the primary mode of energy expenditure) is independent of that of other active sensors, then the system is better represented by the IL model. However, in many scenarios, the sensors could perform data transmission collaboratively; in such a case, the CL model may be more appropriate. Note that these two models represent two extreme forms of correlation, and real-life situations can be expected to fall in between these two extremes.

Note that the optimal time-average utility (computed over all possible activation policies) could be different for the two correlation models. We denote optimal time-average utility for the IL and CL models as \bar{U}_I^* and \bar{U}_C^* , respectively.

D. Threshold Activation Policies

Note that the set of all possible activation policies can be very large, and the structure of these policies can be very complex. Therefore, determining the optimal activation policy for the IL and CL models, can be very difficult, and evaluating the optimal time-average utilities, \bar{U}_I^* and \bar{U}_C^* , can be computationally intensive. Therefore, we focus primarily on threshold activation policies. A threshold activation policy with parameter m , is characterized as follows: a ready sensor s is activated if the number of active sensors does not exceed m after s is activated; otherwise, s is kept in the ready state. In other words, a threshold policy with parameter m tries to maintain the number of active sensors as close to m as possible. Note that with such a policy, the number of active sensors can never exceed m , and there cannot be any ready sensors in the system when the number of active sensors is less than m . The time-average utility for threshold activation policy with parameter m , are denoted by $\bar{U}_{T,I}(m)$ and $\bar{U}_{T,C}(m)$, for the IL and CL models, respectively.

E. Analysis

In this section, we compare the performance of threshold activation policies with respect to the optimal activation policy. In the following, $\rho = \frac{\mu_1}{\mu_2} \geq 1$. For simplicity of exposition, we assume ρ is an integer, and N is divisible by $(\rho + 1)$, although our results can be generalized to the cases where these assumptions do not hold.

1) *Upper Bound on \bar{U}_I^* and \bar{U}_C^** : Since the optimal time-average utility is difficult to compute, we obtain an upper bound on it, and compare the performance of threshold policies with this bound.

Theorem 1: The optimal time-average utility for the two correlation models, \bar{U}_I^* and \bar{U}_C^* , are both upper-bounded by $U(\frac{N}{1+\rho})$, i.e.,

$$\bar{U}_I^* \leq U\left(\frac{N}{1+\rho}\right) \quad \text{and} \quad \bar{U}_C^* \leq U\left(\frac{N}{1+\rho}\right).$$

The proof of the above result involves concavity arguments and Jensen's Inequality [15] and the details are given in the Appendix. Theorem 1 implies that the time-average utility under *any* policy can not be greater than $U(\frac{N}{1+\rho})$. Further, the bound is achieved exactly when all the sensors have deterministic discharge and recharge times of lengths $1/\mu_1$ and $1/\mu_2$, respectively. With random discharge/recharge times, the bound may not be tight; however, as we show below, it is fairly good bound in our case.

Now we derive worst-case bounds on the performance of threshold policies with respect to the optimal policy, for the two correlation models.

2) *Threshold Activation Policies for the IL Model*: Consider a threshold activation policy with parameter $m \in \{1, 2, 3, \dots, N\}$. Let the time-average utility of the system for this threshold be denoted by $\bar{U}_{T,I}(m)$. Then $\bar{U}_{T,I}^*$, the optimal threshold-based time-average utility for the IL model, is defined as $\bar{U}_{T,I}^* = \max_{m=1}^N \bar{U}_{T,I}(m)$.

Next we state a result for the threshold policy with parameter N . Note that with a threshold of N , once a sensor

is completely recharged, it is immediately activated. In other words, no sensor is kept in the ready state.

Theorem 2: The time-average utility at threshold N for the IL model, $\bar{U}_{T,I}(N)$, is lower-bounded by $\frac{1}{2}U(\frac{N}{1+\rho})$, i.e.,

$$\bar{U}_{T,I}(N) \geq \frac{1}{2}U\left(\frac{N}{1+\rho}\right).$$

The proof of the above result involves concavity arguments and analytical approximations of the time average utility function. Details are provided in the Appendix. Theorem 2 implies that with a threshold of N , the performance of the system will be within 50% of the optimal performance over all policies. Our numerical studies also show that the time-average utility of the system with a threshold of N is usually quite close to the optimal. However, the optimal threshold could in general be much less than N . The best threshold policy can be found by finding the maximum of $\bar{U}_{T,I}(m)$ over all $m \in \{1, 2, 3, \dots, N\}$, using the expression in (5) listed in the Appendix. Theorem 2, in conjunction with Theorem 1, implies that the optimal threshold-based time-average utility, $\bar{U}_{T,I}^*$ satisfies $\bar{U}_{T,I}^* \geq \frac{1}{2}\bar{U}_I^*$.

It is possible to obtain a stronger bound on the performance of the optimal threshold policy for the IL model. The derivation of this bound uses results from the analysis of the CL model, that we discuss next.

3) *Threshold Activation Policies for the CL Model*: Consider a threshold activation policy with parameter m . We assume that N is a multiple of m . Initially, all N sensors are fully charged, and m of these are activated, and the remaining $N - m$ are in the ready state. It is easy to see that the sensors will become grouped into $c = N/m$ batches, each of size m , and always move through the different states in these batches. From our definition of a threshold activation policy, it follows that at most one batch can remain active at any time.

Let S_N denote the set of all factors of N , i.e., all integers which divide N . Then $\bar{U}_{T,C}^*$, the optimal threshold-based time-average utility for the CL model, is defined as $\bar{U}_{T,C}^* = \max_{m \in S_N} \bar{U}_{T,C}(m)$.

Next we state an important bound on $\bar{U}_{T,C}^*$.

Theorem 3: The optimal threshold-based time-average utility for the CL model, $\bar{U}_{T,C}^*$, is lower-bounded by $\frac{3}{4}U(\frac{N}{1+\rho})$, i.e.,

$$\bar{U}_{T,C}^* \geq \frac{3}{4}U\left(\frac{N}{1+\rho}\right).$$

The proof of Theorem 3 involves concavity arguments, and analytical approximations of the time average utility function (see Appendix for details). Theorem 3, together with Theorem 1, implies $\bar{U}_{T,C}^* \geq \frac{3}{4}\bar{U}_C^*$. Therefore, the performance of the best threshold policy is within a factor of $\frac{3}{4}$ of the optimal performance over all policies. The best threshold policy can be found by finding the maximum of $\bar{U}_{T,C}(m)$ over all $m \in S_N$, using the expression in (8) listed in the Appendix. As we describe later, our numerical results show that this maximum is typically attained at some intermediate value of m . From the proof of Theorem 3, it can be also shown that a threshold of $\frac{N}{1+\rho}$ achieves the bound of $\frac{3}{4}$.

4) *Comparison of IL and CL Models*: The following result states that for every threshold m , the performance of the IL model is at least as good as that for the CL model.

Theorem 4: For any $m \in S_N$, the time-average utility for the IL model, $\bar{U}_{T,I}(m)$, can be no less than the time-average utility under the CL model, $\bar{U}_{T,C}(m)$, i.e.,

$$\bar{U}_{T,I}(m) \geq \bar{U}_{T,C}(m).$$

The proof of Theorem 4 involves constructing equivalent queuing networks corresponding to the IL and CL models. We show that maximizing time average utility is equivalent to minimizing the mean waiting times at certain stations in these networks. Using prior results on the impact of pooled servers on network throughput and Little's law, we can prove the desired result. An outline of the proof is given in the appendix while the detailed proof is given in the technical report [8]. Theorem 4 implies that the presence of correlation amongst the discharge and recharge times of sensors in a batch degrades system performance.

Theorems 3 and 4 allow us to improve our earlier bound on the performance of the optimal threshold policy for the IL model.

Corollary 5: The optimal threshold-based time-average utility for the IL model, $\bar{U}_{T,I}^*$, is lower-bounded by $\frac{3}{4}U(\frac{N}{1+\rho})$, i.e.,

$$\bar{U}_{T,I}^* \geq \frac{3}{4}U\left(\frac{N}{1+\rho}\right).$$

Corollary 5, in conjunction with Theorem 1, implies $\bar{U}_{T,I}^* \geq \frac{3}{4}\bar{U}_I^*$. Therefore, for both the IL and CL models, the performance of the best threshold policy is within a factor of $\frac{3}{4}$ of the best achievable performance.

F. Numerical Results

In this section, we report results from numerical experiments on the performance of threshold policies for the IL and CL models under different parameter settings. For the utility function $U(n) = 1 - (1 - p_d)^n$, we conduct numerical experiments for different values of $p_d (= 0.1, 0.9)$, $N (= 16, 32, 48)$, and $\rho (= 3, 7, 15)$. To obtain the different values of ρ , we set $\mu_2 = 1$ and vary μ_1 . For each parameter setting, we compare the time-average utility of the system for different values of the threshold. The detailed results from the numerical experiments are described in [8].

1) *Performance under IL and CL models:* Figures 2 and 3 depict typical plots that describe the performance of the threshold policies in the presence of low (Figure 2) and high (Figure 3) probability of detection (p_d). Note that the figures show the time-average utilities $\bar{U}_{T,I}(m)$ and $\bar{U}_{T,C}(m)$ along with $U(\frac{N}{1+\rho})$, the upper bound on the maximum achievable time-average utility. Figures 2 and 3 indicates that for both the CL and IL models the time average utility is maximized at an intermediate value of the threshold. Further, the optimal threshold is distinct from $\frac{N}{1+\rho}$. For the CL model, when operating with a threshold greater than the optimal, the time-average utility decreases very rapidly with the threshold value. However, for the IL model the decrease is gradual and in many cases marginal. The rapid decrease in performance of the CL model for thresholds other than the optimal emphasizes the need to model and understand impact of correlation on system performance.




Fig. 2. Time-average utility for both models ($\mu_1 = 7, \mu_2 = 1, N = 16, p_d = 0.1$)




Fig. 3. Time-average utility for both models ($\mu_1 = 7, \mu_2 = 1, N = 16, p_d = 0.9$)

Tables I and II list the ratio of the time-average utility obtained at the optimal threshold ($\bar{U}_{T,I}^*$ or $\bar{U}_{T,C}^*$) to the lower bound of $\frac{3}{4}U(\frac{N}{1+\rho})$. Note that this ratio must lie between 1 and $\frac{4}{3}$. A value close to 1 indicates a tight lower bound, whereas a value close to $\frac{4}{3}$ indicates that performance of the optimal threshold policy is close to the best achievable performance.

Table I indicates that for low values of the probability of detection p_d , the time average utility obtained by the optimal threshold policy for the CL model is very close to the lower bound. However, performance of the optimal threshold policy for the IL model is fairly close to the maximum value $U(\frac{N}{1+\rho})$. Table II indicates that for high values of the probability of detection p_d , the time average utility obtained by the optimal threshold policy for both models is very close to the maximum value $U(\frac{N}{1+\rho})$, although the performance for the IL model is slightly better, as expected. In summary, in most cases, the optimal threshold policies yield performance that are very close to that maximum achievable performance. When this is not the case (for instance, for the CL model and low values of p_d), the performance is fairly close to our lower bound of $\frac{3}{4}U(\frac{N}{1+\rho})$. Further, the numerical experiments also indicate that our bounds are fairly robust to the choice of N and ρ .

2) *Distribution Independence:* Simulations were carried out for the IE model with recharge/discharge times randomly dis-

	N=16	N=32	N=48
$\rho=3$	1.29	1.30	1.31
$\rho=7$	1.28	1.29	1.29
$\rho=15$	1.27	1.28	1.28

(a) IL Model

	N=16	N=32	N=48
$\rho=3$	1.06	1.06	1.06
$\rho=7$	1.14	1.09	1.09
$\rho=15$	1.22	1.16	1.17

(b) CL Model

TABLE I

RATIO OF OPTIMAL THRESHOLD-BASED TIME AVERAGE UTILITY AND LOWER BOUND, FOR $p_d = 0.1$.

	N=16	N=32	N=48
$\rho=3$	1.33	1.33	1.33
$\rho=7$	1.24	1.33	1.33
$\rho=15$	1.14	1.25	1.32

(a) IL Model


	N=16	N=32	N=48
$\rho=3$	1.31	1.32	1.33
$\rho=7$	1.21	1.32	1.33
$\rho=15$	1.14	1.21	1.31

(b) CL Model

TABLE II

RATIO OF OPTIMAL THRESHOLD-BASED TIME AVERAGE UTILITY AND LOWER BOUND, FOR $p_d = 0.9$.

tributed under various different distributions viz Beta, Gamma, Uniform. The performance of threshold activation policies was measured by computing the time-average utility. It is observed that the performance of a threshold activation policy is independent of the distribution of the recharge and discharge times. This can be inferred by observing the performance of the threshold activation policies under the various distributions with same mean and variance, as shown in Figure 4.

Fig. 4. Time-average utility plot for various distributions with same mean ($\mu_1 = 7, \mu_2 = 1$) and variance ($= \text{mean}^2/3$)

The simulation results also indicate that the time average utility increases with decrease in the variance of the recharge and discharge times.

III. NODE ACTIVATION POLICIES WITH PARTIAL COVERAGE OVERLAP

In a realistic deployment scenario, nodes may be deployed at random, and therefore, nodes will typically cover different areas in the physical space of interest. In other words, the coverage areas of two sensors may overlap only partially, or may not overlap at all (i.e., be disjoint). In this section, we extend our threshold activation policy to this very general scenario.

As mentioned before, the case of partial coverage overlap is very difficult to model and analyze, even for the special class of threshold policies. In this section, therefore, we will try to develop a solution heuristically, based on the insights obtained for the complete coverage overlap scenario. We then show, through extensive simulations, that our solution yields a performance trend that is similar to that in the case with complete coverage overlap. In particular, we observe that the performance achieved still satisfies the three-fourth bound with respect to the upper bound on the optimum over all policies.

A. Distributed Node Activation Algorithm

To motivate our distributed activation algorithm, let us assume that a sensor i wants to maintain a utility of $U(m_i)$ per unit area per unit time in its coverage area. In other words, if the coverage area of the sensor is denoted by \mathcal{A}_i , then the sensor targets to derive a utility of $|\mathcal{A}_i|U(m_i)$ per unit time. When the sensor is in the ready state, then at any decision instant, the sensor computes the current utility per unit time in its coverage area. If the current utility is less than the targeted utility, then the node activates itself; otherwise, the node remains in the ready state until the next decision instant.

A sensor can compute the utility derived from its coverage area in the following manner. For a generic area element $A \in \mathcal{A}_i$, let $n(A, t)$ denote the number of sensors covering A at time t . Then the utility per unit time in the coverage area of node i is calculated as

$$\int_{\mathcal{A}_i} U(n(A, t)) dA. \quad (3)$$

Assume that node i can communicate with all nodes whose coverage areas overlap with its own coverage area. Then the sensor can periodically poll those neighbors to know their activation state. Assuming that the sensor i knows the coverage patterns of those neighbors, it can compute the targeted utility by evaluating the expression in (3). Therefore, the proposed algorithm can be realized in a distributed setting based only on local information.

Note that the algorithm is motivated by the threshold activation policy discussed in the previous section, and in the case of complete coverage overlap, it reduces to a distributed implementation of the threshold policy described earlier.

In practice the decision interval needs to be chosen carefully to ensure that not too much energy is wasted in the ready state by periodic wakeup and polling, while guaranteeing good performance.

The thresholds m_i can be defined globally or locally, and accordingly we have two variants of our policy:

- i) *Global threshold policy*: In this case, the $m_i = m \forall i$, where the fixed threshold m is chosen appropriately.
- ii) *Local threshold policy*: In this case, the m_i can be different for each i , depending on the local neighborhood of the individual sensor nodes.

In Section III-D, we comment on the appropriate choice of the local and global thresholds, to yield optimum performance. We can intuitively expect the local threshold policy to perform better, particularly in scenarios where there is a high spatial variance in the density of nodes in the deployment region. For the local threshold policy, nodes in areas with larger density can have a higher threshold, while nodes in a sparser region can set its threshold to a lower value. However, if the nodes are deployed more or less uniformly, then both these policies are observed to perform very well in simulations, although local threshold policy performs slightly better.

B. Upper Bound on Optimal Time-average Utility

Next we state an useful upper bound on the optimal time-average utility derived from a sensor network with partial coverage overlap. We assume that the mean discharge and recharge times are $\frac{1}{\mu_1}$ and $\frac{1}{\mu_2}$ respectively, $\rho = \frac{\mu_1}{\mu_2} \geq 1$. We do not make any assumption on the distribution of the discharge and recharge times. Let A denote a generic area element in the physical space of interest, and $N(A)$ denote the number of active sensors that cover area element A .

Corollary 6: The optimal time-average utility for a general network of sensors is upper-bounded by

$$\int_A U \left(\frac{N(A)}{1 + \rho} \right) dA . \quad (4)$$

The above result can be proved following the same line of analysis as in the proof of Theorem 1, and is therefore omitted here. Since the optimal policy is difficult to formulate and compute in this case, we will compare the performance of our algorithm with respect to this upper bound.

C. Sensor Lifetime Models

We consider five different lifetime models for this scenario. The first two models (independent and correlated lifetime models) are extensions of the IL and CL models considered for the complete coverage overlap case. The next two models (independent and correlated event-based lifetime models) are event-based. We assume that discharging and recharging depends on events that occur randomly in the deployment region (in a way described below). Finally, for the sake of comparison, we also consider a deterministic lifetime model, where the discharge and recharge times of each sensor are

fixed.

i) *Independent Lifetime Model*: The discharge and recharge times of the sensors are exponential *i.i.d* with means $1/\mu_1$ and $1/\mu_2$ respectively.

ii) *Correlated Lifetime Model*: The discharge and recharge times of the sensors are exponential with means $1/\mu_1$ and $1/\mu_2$ respectively. However, the discharge (recharge) times of all sensors entering the active (passive) state at the same time is the same. The discharge (recharge) times of sensors entering the active (passive) state at different times are independent of each other.

iii) *Independent Event-based Lifetime Model*: Events are assumed to occur randomly in the physical space of interest, and a sensor node gets discharged (by a fixed amount q) only when an event occurs within its coverage area. Events are assumed to occur according to a Poisson process, and are uniformly distributed in the area of interest. A sensor node, on activation, is assumed to have a total energy of Q units. Therefore, an active sensor gets fully discharged once Q/q events have occurred within its coverage area. The recharge process is modelled similar to the discharge process. A passive node gets fully recharged once a certain number of random “recharge events” have occurred in its coverage area. The mean inter-event times for discharging and recharging are chosen so that the mean discharge and recharge times of sensors are equal to $1/\mu_1$ and $1/\mu_2$ respectively.

iv) *Correlated Event-based Lifetime Model*: Here the network is divided into imaginary blocks of equal sizes. As in the case of the independent event-based lifetime model, events occur according to a poisson process, and are assumed to be uniformly distributed in the area of interest. However, an event occurring anywhere in the block affects all the sensors located in this block in a similar manner. This introduces spatial correlation between the discharge and recharge times of the sensors. The degree of spatial correlation depends on the sizes of the blocks. In this model too, the mean inter-event times are chosen so that the mean discharge and recharge times of nodes are equal to $1/\mu_1$ and $1/\mu_2$ respectively.

v) *Deterministic Lifetime Model*: The discharge and recharge times of the sensors are fixed and are equal to $1/\mu_1$ and $1/\mu_2$ respectively.




Fig. 5. Performance of different lifetime models with global thresholds




Fig. 6. Performance of different lifetime models with local thresholds




Fig. 8. Correlated Events model performance with local thresholds




Fig. 7. Correlated Events model performance with global thresholds

D. Simulation Results and Discussion

The performance of the node activation algorithm described above is evaluated using simulations for a wide range of parameters for both the cases of global and local thresholds. In the representative simulation results presented here, the simulation setup and the parameters used are as follows. A total of $N = 52$ sensors, each having a circular coverage pattern of radius 12 units, are thrown uniformly at random in an area of size 50×50 . With these parameters, the mean coverage of the network (\bar{N}), defined as the average number of sensors covering any point in the deployment region, is observed to be approximately 9.1. For the event-based lifetime models, we use $Q/q = 100$. For the correlated event based model, number of blocks is set to 4. The utility function used is given by $U(n) = 1 - (1 - p_d)^n$ where $p_d = 0.1$. Also the upper bound on the maximum achievable utility in this case is calculated to be 0.159511.

Figures 5 and 6 show the performance of the various models for global and local thresholds. Let us define α , the local threshold parameter, as $\alpha = m_i / (\frac{n_i}{1+\rho})$, where n_i is the number of sensors (including i) that cover the point where i is located. Note that in Figure 6, the time-average utility is plotted against this local threshold parameter α .

From Figure 5, we observe that with a fixed global threshold of $m = \bar{N}/(1+\rho)$ (≈ 2 in this case), the time average utility

is greater than three-fourth of the upper bound. Similarly, from Figure 6, we observe that with a local threshold of $m_i = n_i/(1+\rho)$ (the case of $\alpha = 1$), the time average utility is greater than three-fourth of the upper bound as well. For all of the event models considered, we see that this threshold value also achieves the close to best performance attained over all thresholds. Simulations performed for other network configurations yielded similar results. The performance for the deterministic model is close to the optimal for the threshold of $m_i = n_i/(1+\rho)$ and $m = \bar{N}/(1+\rho)$ for the local and global threshold policies respectively.

Note that the Figures 5 and 6 also show that the performance at higher thresholds drops significantly as the degree of spatial correlation in the sensor lifetimes increases. Figures 7 and 8 also demonstrate this fact more clearly. Note that the degree of spatial correlation increases with fewer number of blocks (i.e., with the increase in the block size). With the increase in spatial correlation, the performance drops significantly for the activation policy at higher thresholds.

REFERENCES

- [1] I. F. Akyildiz, W. Su, Y. Sankarasubramanian and E. Cayirci, "Wireless Sensor Networks: A Survey," *Computer Networks* 38, pp. 393-422, 2002.
- [2] Bertsekas, D.P., *Dynamic Programming and Optimal Control, Volume II*, Athena Scientific, Belmont, Massachusetts, 2001.
- [3] D. Blough and P. Santi, "Investigating Upper Bounds on Network Lifetime Extension for Cell-Based Energy Conservation Techniques in Stationary Ad Hoc Networks," Proc. of Mobicom 2002, Atlanta, Sept. 2002.
- [4] B. Chen, K. Jamieson, H. Balakrishnan and R. Morris, "Span: An Energy Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks," *Proceedings of ACM Mobicom 2001*, Rome, 2001.
- [5] I. Chlamtac and C. Petrioli and J. Redi, "Energy-conserving Access Protocols for Identification Networks," *IEEE/ACM Transactions on Networking*, 7(1):51-59, Feb. 1999.
- [6] A. Chokalingam and M. Zorzi, "Energy Efficiency of Media Access Protocols for Mobile Data Networks," *IEEE Transactions on Communications*, 46(11):1418-21, Nov. 1998.
- [7] Y. Dallery, K. E. Stecke, "On the Optimal Allocation of Servers and Workloads in Closed Queuing Networks" *Operations Research* 38(4):694-703, 1990.
- [8] K. Kar, A. Krishnamurthy, "Threshold Activation Policies in a Random Sensing Environment," *Technical Report*, Rensselaer Polytechnic Institute, July, 2004, www.ecse.rpi.edu/~koushik/sensor-threshold-techrep.pdf.

- [9] R. Kravets and P. Krishnan, "Power management techniques for mobile communications," *In Proceedings of The Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, October 1998.
- [10] M. Papadopoulou, H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," *ACM SIGMOBILE Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc) 2001*, Long Beach, CA, Oct. 2001.
- [11] J. G. Shanthikumar, D. D. Yao, "On Server Allocation in Multiple Center Manufacturing Systems" *Operations Research* 36(2):333-342, 1988.
- [12] S. Singh, M. Woo, C. S. Raghavendra, "Power Aware Routing In Mobile Ad-hoc Networks," *Proceedings of Mobicom '98*, Dallas, October 1998.
- [13] Y. Xu, J. Heidemann, D. Estrin, "Geography-informed Energy Conservation for Ad Hoc Routing," *Proceedings of Mobicom 2001*, Roime, July 2001.
- [14] W. Ye, J. Heidemann, and D. Estrin, "Medium Access Control with Coordinated, Adaptive Sleeping for Wireless Sensor Networks," *Technical Report ISI-TR-567*, USC/Information Sciences Institute, January, 2003.
- [15] R.L. Wheeden, A.Zygmund, *Measure and Integral, An Introduction to Real Analysis*, Marcel Dekker, Inc., NY, pp. 694-703, 1977.

APPENDIX I: PROOF OUTLINES OF THEOREMS

Proof Outline of Theorem 1: The proof of the above result involves concavity arguments and Jensen's Inequality [15]. Let f and p be measurable functions finite a.a. on \mathbb{R} . Suppose, that fp and p are integrable on \mathbb{R} , $p \geq 0$, and $\int p > 0$. If ϕ is convex in an interval containing the range of f , then Jensen's inequality states that:

$$\phi \left(\frac{\int_{\mathbb{R}} fp}{\int_{\mathbb{R}} p} \right) \leq \frac{\int_{\mathbb{R}} \phi(f)p}{\int_{\mathbb{R}} p}$$

Let $n(t)$ denote the number of sensors in the active state at time t . Since $U(\cdot)$ is concave, substituting $\phi = U(\cdot)$, $f = n(t)$ and $p = 1$ in the above, Jensen's Inequality implies that:

$$U \left(\frac{\int_0^T n(t)dt}{T} \right) \geq \frac{\int_0^T U(n(t))dt}{T}$$

Since, $U(\cdot)$ is continuous, we have:

$$\lim_{T \rightarrow \infty} U \left(\frac{\int_0^T n(t)dt}{T} \right) \geq \lim_{T \rightarrow \infty} \frac{\int_0^T U(n(t))dt}{T}$$

Define $\psi_i(t)$ such that $\psi_i(t) = 1$ if sensor i is in active state at time t and $\psi_i(t) = 0$ if sensor i is in passive state at time t . Then, continuity of $U(\cdot)$ also implies

$$\begin{aligned} \lim_{T \rightarrow \infty} U \left(\frac{\int_0^T n(t)dt}{T} \right) &= U \left(\lim_{T \rightarrow \infty} \frac{\int_0^T n(t)dt}{T} \right) \\ &= U \left(\lim_{T \rightarrow \infty} \frac{\int_0^T \sum_{i=1}^{i=N} \psi_i(t)dt}{T} \right) \end{aligned}$$

Since $\psi_i(t)$ is positive and bounded,

$$\begin{aligned} \lim_{T \rightarrow \infty} U \left(\frac{\int_0^T n(t)dt}{T} \right) &= U \left(\lim_{T \rightarrow \infty} \sum_{i=1}^{i=N} \frac{\int_0^T \psi_i(t)dt}{T} \right) \\ &= U \left(\sum_{i=1}^{i=N} \lim_{T \rightarrow \infty} \frac{\int_0^T \psi_i(t)dt}{T} \right) \end{aligned}$$

Further, since all sensors are identical, for any k

$$U \left(\lim_{T \rightarrow \infty} \frac{\int_0^T n(t)dt}{T} \right) = U \left(N \lim_{T \rightarrow \infty} \frac{\int_0^T \psi_k(t)dt}{T} \right)$$

Since the times each sensor spends in active and passive states are independent, with mean $\frac{1}{\mu_1}$ and $\frac{1}{\mu_2}$, we have

$$\frac{1}{1+\rho} \geq \lim_{T \rightarrow \infty} \frac{\int_0^T \psi_k(t)dt}{T}$$

where the equality holds if the sensor spends zero time in the ready state. Therefore we have

$$U \left(\frac{N}{1+\rho} \right) \geq U \left(\lim_{T \rightarrow \infty} \frac{\int_0^T n(t)dt}{T} \right)$$

This implies

$$\begin{aligned} U \left(\frac{N}{1+\rho} \right) &\geq U \left(\lim_{T \rightarrow \infty} \frac{\int_0^T n(t)dt}{T} \right) \\ &\geq \lim_{T \rightarrow \infty} \frac{\int_0^T U(n(t))dt}{T} \end{aligned}$$

The above theorem implies that the time-average utility under any policy can not be greater than $U(\frac{N}{1+\rho})$. In particular, the optimal time-average utility for the two correlation models, \bar{U}_I^* and \bar{U}_C^* , are both upper-bounded by $U(\frac{N}{1+\rho})$, thus proving Theorem 1. ■

Proof Outline of Theorem 2: Using steady state Markov chain analysis, the time-average utility of the system, $\bar{U}_{T,I}(m)$, can be computed as

$$\bar{U}_{T,I}(m) = \frac{\sum_{i=1}^N U(i)\alpha(i, m)}{\sum_{i=0}^N \alpha(i, m)}, \quad (5)$$

where $\alpha(i, m)$, $i = 1, 2, \dots, N$, are defined as

$$\alpha(i, m) = \begin{cases} \binom{N}{i} \rho^{-i} & \text{if } i \leq m, \\ \binom{N}{i} \frac{i! \rho^{-i}}{m! m^{i-m}} & \text{otherwise.} \end{cases} \quad (6)$$

Using the above expressions, time average utility obtained for a threshold of $m = N$ is:

$$\bar{U}_{T,I}(N) = \sum_{i=0}^N \frac{U(i) \binom{N}{i} (1/\rho^i)}{(1+1/\rho)^N}$$

We define $w = \frac{N}{\rho+1}$ and show that $\frac{\bar{U}_{T,I}(N)}{\bar{U}_{T,I}(w)} \geq \frac{1}{2}$. We have

$$\begin{aligned} \bar{U}_{T,I}(N) &= \sum_{i=0}^N \frac{U(i) \binom{N}{i}}{\rho^i (1+1/\rho)^N} \\ &= \sum_{i=1}^w \frac{U(i) \binom{N}{i}}{\rho^i (1+1/\rho)^N} + \sum_{i=w+1}^N \frac{U(i) \binom{N}{i}}{\rho^i (1+1/\rho)^N} \end{aligned}$$

But from the concavity of $U(\cdot)$ we have, $U(k) \geq (k/w)U(w)$ for $k \leq w$ and $U(k) \geq U(w)$ for $k > w$. Hence

$$\begin{aligned} \frac{\bar{U}_{T,I}(N)}{\bar{U}_{T,I}(w)} &\geq \sum_{i=1}^w \frac{(i/w) \binom{N}{i} (1/\rho^i)}{(1+1/\rho)^N} \\ &\quad + \sum_{i=w+1}^N \frac{\binom{N}{i} (1/\rho^i)}{(1+1/\rho)^N} \\ &\geq \frac{1}{2w} (w+1) \left[\sum_{i=1}^w \frac{\binom{N}{i} (1/\rho^i)}{(1+1/\rho)^N} \right] \\ &\quad + \sum_{i=w+1}^N \frac{\binom{N}{i} (1/\rho^i)}{(1+1/\rho)^N} \\ &\geq \frac{1}{2}. \end{aligned} \quad (7)$$

The result follows from (7) and Theorem 1. ■

Proof Outline of Theorem 3: Using steady-state Markov chain analysis, the time-average utility of the system, $\bar{U}_{T,C}(m)$, can be computed as

$$\bar{U}_{T,C}(m) = U(m) \left(1 - \frac{\rho^c}{c!} \sum_{i=0}^c \frac{\rho^i}{i!} \right). \quad (8)$$

where $c = N/m$.

To prove the lower bound on $\bar{U}_{T,C}(\hat{m})$, it is sufficient to show that there exists an \hat{m} , such that $\bar{U}_{T,C}(\hat{m}) = \frac{3}{4}U\left(\frac{N}{1+\rho}\right)$. In particular, we show that the result holds for $\hat{m} = \frac{N}{1+\rho}$, by considering two cases, namely $\rho = 1$ and $\rho \geq 2$.

Case 1 ($\rho = 1$):

Since $\rho = 1$, $\hat{m} = \frac{N}{1+\rho} = \frac{N}{2}$ and $c = \frac{N}{\hat{m}} = 2$. Hence,

$$\begin{aligned} \bar{U}_{T,C}(\hat{m}) &= \left(1 - \frac{\rho^2}{2!} \sum_{i=0}^2 \frac{\rho^i}{i!} \right) U\left(\frac{N}{1+\rho}\right) \\ &= \left(\frac{1+\rho}{1+\rho + \frac{\rho^2}{2}} \right) U\left(\frac{N}{1+\rho}\right) \\ &= 0.8U\left(\frac{N}{1+\rho}\right) \\ &\geq \frac{3}{4}U\left(\frac{N}{1+\rho}\right) \end{aligned}$$

Case 2 ($\rho \geq 2$):

Since $\rho \geq 2$, $\hat{m} = \frac{N}{1+\rho} \leq \frac{N}{3}$ and $c = \frac{N}{\hat{m}} \geq 3$. Hence,

$$\begin{aligned} \bar{U}_{T,C}(\hat{m}) &= \left(1 - \frac{\rho^c}{c!} \sum_{i=0}^c \frac{\rho^i}{i!} \right) U\left(\frac{N}{1+\rho}\right) \\ &\geq \left(1 - \frac{\rho^c}{\sum_{i=c-3}^c \frac{\rho^i}{i!}} \right) U\left(\frac{N}{1+\rho}\right) \\ &\geq \left(1 - \frac{1}{4 + \frac{2}{\rho} - \frac{1}{\rho^2}} \right) U\left(\frac{N}{1+\rho}\right) \\ &\geq \left(1 - \frac{1}{4} \right) U\left(\frac{N}{1+\rho}\right) \end{aligned}$$

$$\geq \frac{3}{4}U\left(\frac{N}{1+\rho}\right). \quad \blacksquare$$

Proof Outline of Theorem 4: Figure 9 provides a queuing network representation of the IL model operating with a threshold of m . It is easily seen that the Markov chain for this queuing network is identical to that of the IL model described earlier. Figure 10 provides a queuing network representation

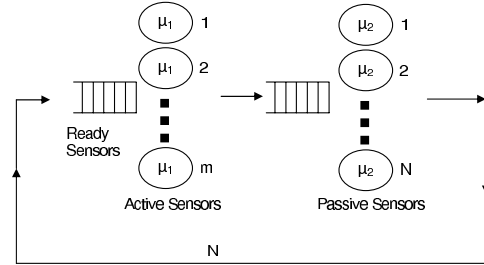


Fig. 9. Queuing network representation for the IL model

of the CL model operating with a threshold of m . Again, it can be seen that the Markov chain for this queuing network is identical to that of the CL model described earlier. To

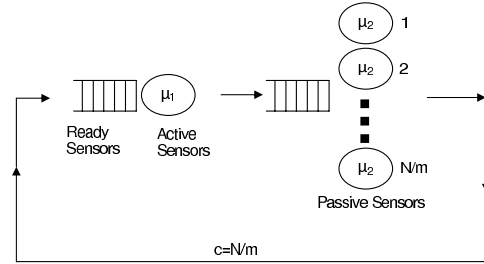


Fig. 10. Queuing network representation for the CL model

show that $\bar{U}_{T,I}(m) \geq \bar{U}_{T,C}(m)$ we construct the multi-class queuing network model of an intermediate system shown in Figure 11. There are m classes of customers with c customers in each class. Each class of customer visits one of the exponential m servers in station 1 (active sensors) and one of the N exponential servers in station 2 (passive sensors). Note that the behavior of each class of customers in the network shown in Figure 11 is identical to that of the batches in the network representation of the CL model. Further, like the network representation of the IL model this network has N customers. However, the m servers in station 1 of the network in Figure 11 are dedicated – one to each class of customers. Comparing networks in Figures 10 and 11 we note that the probability of at least one of the m servers is busy (in Figure 11) is at least as much (if not greater) than the probability that station 1 (in Figure 10) is busy. Further, comparing networks in Figures 9 and 11 we note that the probability that all m servers are busy is at least as much (if not greater) for the network shown in Figure 9. This observation follows from the result shown in [7] and [11] that state that pooling of servers result in improved station utilization and throughput. These two observations imply that $\bar{U}_{T,I}(m) \geq \bar{U}_{T,C}(m)$. ■

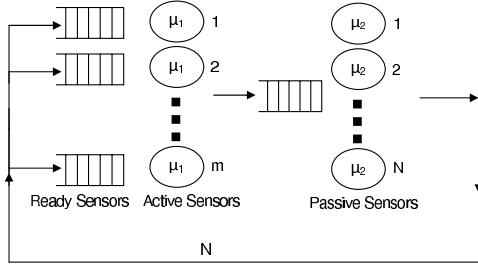


Fig. 11. Queuing network model of intermediate system

APPENDIX II: OPTIMAL POLICY COMPUTATION IN THE COMPLETE COVERAGE OVERLAP CASE

In this section we formulate the problem of determining the optimal activation policy for the case of complete coverage overlap. We describe a solution procedure and discuss the associated computational complexity and motivate the need to study threshold activation policies.

We represent the system state $x \in X$ by the parameter tuple (a_x, s_x) , where a_x denotes the number of available sensors (active+ready) and s_x denotes the number of active sensors. Note that, since we assume that there are N sensors, $a_x = 0, \dots, N$ and $s_x = 0, \dots, a_x$. Since charging and discharging lifetimes are assumed to be exponentially distributed, the system in state x transitions to state y after an exponentially distributed time with probability $p_{x,y}(u)$ that depends on the policy u executed. The objective is to find the optimal policy \mathcal{U}^* that maximizes the time average utility given by

$$\bar{U} = \lim_{T \rightarrow \infty} \frac{1}{T} E \int_0^T g(x(t), u(t)) dt. \quad (9)$$

where $g(x, u) = 1 - (1 - p_d)^n$ is the time utility gained per unit time with n active sensors when the system is in state x .

Note that under these assumptions, the problem of finding the optimal policy becomes a continuous time Markov decision problem where the objective is to maximize the average utility per stage. Properties of the state space and transition probabilities enable us to restrict the search space for the optimal policy. First, we note that the problem formulation implies a finite space. Second, it can be shown that the Markov chain resulting from every stationary policy would be irreducible. Then it follows from Bertsekas [2](page 277) that there exists a stationary policy $u(x)^*$ for $x \in X$ that yields the same utility as the optimal policy \mathcal{U}^* . Further, under a stationary policy, the stochastic process $X(t)$ forms a regenerative process with finite expected cycle length and therefore the time average utility can be represented in an equivalent and more convenient form given by

$$\bar{U} = \lim_{N \rightarrow \infty} \frac{1}{E(t_N)} E \int_0^{t_N} g(x(t), u(t)) dt. \quad (10)$$

We define $G(x, u)$ to be the single stage expected utility and $\bar{\tau}_x(u)$ to be the expected transition time corresponding to state x and policy u . Then by formulating the corresponding

stochastic shortest path problem [2] it can be shown that the problem of finding the optimal stationary policy is equivalent to finding the $u^*(x)$ for $x \in X$ such that

$$h^*(x) = \min_{u \in \mathcal{U}(x)} \left[G(x, u) - \lambda^* \bar{\tau}_x(u) + \sum_{y \in X} p_{x,y}(u) h^*(y) \right] \quad (11)$$

The above set of equations can be solved for the vector $h(x)^*$ for $x \in X$ and the scalar λ^* , the optimal average utility $\bar{U}^*(x)$ for all states $x \in X$ using techniques such as value iteration, policy iteration or linear programming. The specific values of $\bar{\tau}_x(u)$, and $p_{x,y}(u)$ used in the solution procedure would depend on the particular policy u executed in state x . To illustrate the procedure, consider a particular stationary policy, $u_\alpha(x)^*$ for $x \in X$, where

$$u_\alpha(x) = f_\alpha(a_x) \quad \text{for} \quad 0 \leq f_\alpha(a_x) \leq a_x. \quad (12)$$

Corresponding to this policy, assuming sensor lifetimes are independent, we have

$$\begin{aligned} \bar{\tau}_x(u_\alpha(x)) &= \frac{1}{f_\alpha(x)\mu_1 + (N - a_x)\mu_2} \\ p_{x,y}(u_\alpha(x)) &= \frac{f_\alpha(x)\mu_1}{f_\alpha(x)\mu_1 + (N - a_x)\mu_2} \\ &\quad \text{for } (y) = (a_x - 1, f_\alpha(a_x - 1)) \\ &= \frac{(N - a_x)\mu_2}{f_\alpha(x)\mu_1 + (N - a_x)\mu_2} \\ &\quad \text{for } (y) = (a_x + 1, f_\alpha(a_x + 1)) \\ &= 0 \quad \text{otherwise} \end{aligned}$$

Note that threshold policies and activation based policies are subsets of the particular stationary policy chosen above. It is not necessary that the particular stationary policy $u_\alpha(x)^*$ for $x \in X$, be optimal. Using policy iteration, we revise the choice of the policy till we attain the optimal policy. First, corresponding to the current policy choice, we find a $h_\alpha(x)$ and λ_α satisfying

$$h_\alpha(x) + \lambda_\alpha \bar{\tau}_x(u_\alpha(x)) = G(x, u_\alpha(x)) + \sum_{y \in X} p_{x,y}(u_\alpha(x)) h_\alpha(y) \quad (13)$$

for $x \in X$. Then we find a $u_{\alpha+1}(x)$ such that

$$u_{\alpha+1}(x) = \arg \min_{u \in \mathcal{U}(x)} \left[G(x, u) + \sum_{y \in X} p_{x,y}(u) h_\alpha(y) \right] \quad (14)$$

We update $h_\alpha(x)$ and λ_α till $h_{\alpha+1}(x) = h_\alpha(x)$ and $\lambda_{\alpha+1} = \lambda_\alpha$. This terminates the policy iteration algorithm. It is evident from the above procedure that significant computations would be necessary in the improvement step of the policy iteration procedure. The set of stationary policies although finite is significantly large. Therefore, we restrict our research to a set of simpler policies such as the threshold policies and examine their performance in comparison to the optimal policies.