Dynamic Activation Policies for Event Capture with Rechargeable Sensors

Zhu Ren* Peng Cheng* Jiming Chen* David K.Y. Yau†‡ Youxian Sun*
*State Key Lab. of Industrial Control Technology, Zhejiang University, Hangzhou 310027, China
†Department of Computer Science, Purdue University, West Lafayette, IN, USA
‡Advanced Digital Sciences Center, Singapore
zhuren.zju@gmail.com, {pcheng, jmchen, yxsun}@iipc.zju.edu.cn, yau@cs.purdue.edu

***Abstract*—We consider the problem of event capture by a rechargeable sensor network. We assume that the events of interest follow a renewal process whose event inter-arrival times are drawn from a general probability distribution, and that a stochastic recharge process is used to provide energy for the sensors' operation. Dynamics of the event and recharge processes make the optimal sensor activation problem highly challenging. In this paper we first consider the single-sensor problem. Using dynamic control theory, we consider a full-information model in which, independent of its activation schedule, the sensor will know whether an event has occurred in the last time slot or not. In this case, the problem is framed as a Markov decision process (MDP), and we develop a simple and optimal *greedy* policy for the solution. We then further consider a partial-information model where the sensor knows about the occurrence of an event only when it is active. This problem falls into the class of partially observable Markov decision processes (POMDP). Since the POMDP's optimal policy has exponential computational complexity and is intrinsically hard to solve, we propose an efficient heuristic *clustering* policy and evaluate its performance. Finally, our solutions are extended to handle a network setting in which multiple sensors collaborate to capture the events. We provide extensive simulation results to evaluate the performance of our solutions.**

***Keywords*-rechargeable sensors; event capture; dynamic activation; Markov decision process.**

I. Introduction

In sensor network design, the scheduling of sensors to maximize the capture of events of interest is an important problem. As manual maintenance of networks is expensive in many situations, perpetual unattended operation of the networks is highly desirable and will require the sensors to be self-sufficient in their energy use. The requirement is possible to attain with rechargeable sensors, which are capable of harvesting external energy from the environment, e.g., solar energy, kinetic energy from movements of carriers, and energy from radio waves [1], [2], [3].

However, energy harvesting is generally a gradual process, and its rate is usually insufficient to support uninterrupted operation of the sensors [4]. Hence, it is necessary to duty cycle the sensors to improve energy efficiency, and the design of optimal activation policies for the sensors that can maximize system performance under energy constraints becomes a crucial problem.

When the distribution of event inter-arrival times has *memory*, the activation of sensors to capture the events should adapt to the current state of the system. In particular, it may be suboptimal to duty cycle the sensors according to any fixed periodic schedule, but on-line information can be exploited in the activation to improve system performance. On the other hand, since energy-constrained sensors have limited processing capacity, an adaptation mechanism that is simple enough for efficient implementation is important. In this paper, we consider the problem that one or more rechargeable sensors are deployed randomly within an area to monitor certain points of interest (PoI), where interesting events may appear one after another. We aim to exploit memory in the event process to derive optimal or practically effective, yet simple, strategies to activate the sensors. Our goal is to maximize the quality of monitoring (QoM) (i.e., event capture performance) under energy constraints by a given probabilistic recharge process. The simplicity of our algorithms makes them feasible to apply by resource-constrained wireless sensor nodes.

We assume that the events at each PoI to be detected are random, but they can be described by a general renewal process. We will use slotted time where the duration of a time slot is fixed and given. We propose an activation policy under which, at the beginning of each time slot, the sensor is chosen to be active with a certain computed probability (and chosen to be inactive otherwise). Our performance metric is the *instantaneous* capture probability of events, i.e., if an event to be detected by sensor s arrives in time slot t, we would like to maximize the probability that s is chosen to be active at the beginning of t so that the event can be captured without delay. Instantaneous capture may allow, for example, immediate handling of events that may cause harm if they remain unidentified, or it may allow important but transient information to be learned about the events before the information is lost. Henceforth, we will use "event capture" to mean "instantaneous event capture" where there is no confusion.

Denote by u the average rate of energy use under a designed activation policy. We assume that each sensor has an energy bucket (a "battery") of size K (in energy units). Energy consumed from the bucket by the sensor can be replenished by a stochastic recharge process of average rate e (in power units). We use the result that when K is large enough, the sensor will absorb bursts in the energy use and recharge processes in that if $e \geq u$, then the sensor will not run out of energy w.p. 1. We show that the event capture problem falls within scope of Markov decision processes

(MDP), a standard tool for decision making; e.g., Zhao *et al.* [5] have used this tool to optimize cognitive radio communication. Although MDP has been used to detect Markov events by Jaggi *et al.* [6], their method does not apply for general renewal events. The renewal process is more challenging than Markov events, in that the next state is not only determined by the immediately preceding state, but possibly the whole prior history. Hence, the state space of a renewal process may increase exponentially with time. Particularly, in a partial-information model that we will consider, owing to the "curse of dimensionality," optimizing the activation policy is computationally prohibitive in general [7], [8]. Broadly speaking, it has been proved [7] that a general problem with a finite state space is PSPACE-complete, which is even less likely to have polynomial-time solutions than NP-complete problems.

Our contributions in this paper are as follows:

- Under a *full-information* assumption, we show how on-line information and the memory property of general renewal processes can be exploited to maximize event capture by energy-constrained sensors. The optimization is formulated as a Markov decision process (MDP) over the (possibly infinite) set of per-slot activation probabilities. We solve the optimization problem as a linear program under given constraints. We develop a simple and asymptotically optimal greedy activation policy which approaches optimality as K becomes large.
- Under a *partial-information* assumption, we formulate our problem as a partially observable Markov decision process (POMDP). Constructing a complete information state space, we demonstrate that the optimal policy has exponential computational complexity, which is intractable. Exploiting the characteristics of event arrivals, we then develop an efficient and practically effective heuristic clustering policy that outperforms alternative, plausible periodic and aggressive policies.
- For a sensor network in which multiple sensors may be employed to monitor the same events, we propose simple and practical coordination strategies to extend the single-sensor policies to a collaborative setting.

The remainder of the paper is organized as follows. Section II discusses related work. In Section III, we give preliminaries and a formal definition of our problem. In Section IV-A, we formulate the optimization of a single sensor under full information, as a Markov decision process, and construct a greedy activation policy as a solution. For a partial-information model, Section IV-B formulates the single-sensor optimization problem as a partially observable Markov decision process. Section V considers the case of multiple sensors. Section VI presents extensive simulation results to verify our analysis and illustrate the performance of the proposed algorithms.

II. Related Work

Event capture in sensor networks can be framed as area-, point-, or barrier-coverage problems, which have been studied in [9], [10]. For a dense network, algorithms for maximizing the network lifetime while guaranteeing area coverage have been proposed in [11]. Network coverage algorithms generally aim to guarantee worst-case performance and as such do not exploit event dynamics in the system design.

On the other hand, it is widely accepted that real-world events can be modeled as stochastic processes. Poisson arrivals are in general accurate for the arrivals of many independent events, such as emitted particles from a radioactive source [12] or arrivals of phone calls to a switching center [13]. More generally, the distribution of times between events may have memory, in which case their arrivals can be modeled as general renewal processes. The Pareto distribution has been shown to be an accurate model for workload arrivals in a communication network [14], and the Weibull distribution has found applications in modeling failures in reliability engineering, channel fading in wireless communications, and wind speeds in weather forecasting. Yau *et al.* [15] have exploited the renewal processes of events to optimally duty cycle sensors for maximum information capture. Their focus is on how different forms of utility functions for the information capture can impact the system design. He *et al.* [16] generalize the analysis to a dense network of sensors, and show how different coordination strategies between the sensors can maximize energy efficiency. In both [15] and [16], the optimization is over fixed periodic schedules of the sensors and as such does not exploit on-line information in the control. We show how such dynamic on-line information can be exploited to further improve system performance.

For rechargeable sensors, Seyedi and Sikdar [17] have developed adaptive transmission strategies for body sensor networks that exploit information about both the events of interest and the energy recharge process. Fan *et al.* [18] have proposed centralized and asynchronous distributed algorithms for increasing the steady-state throughput of data extraction from sensor nodes while guaranteeing that no nodes will use up their energy. Jaggi *et al.* [6] and Li *et al.* [19] assume that events happen according to a two-state Markov chain. They design activation policies for one sensor, or two sensors monitoring two independent events, based on whether an event occurred in the last time slot or not. Their methods do not apply for general renewal processes, which have more memory than Markov chains. On the other hand, our work focuses on general renewal processes, where the event inter-arrival times can be drawn from an arbitrary probability distribution. We admit a much wider solution space than the binary state differentiation in their problems. Furthermore, we provide solutions for a network of N rechargeable sensors. The collaborative approach solves the problem that the recharge process of a single sensor may be too slow for practical performance.

III. Problem Formulation and Basic Approach

A. Problem formulation

We define a *renewal process* of event arrivals. In this process, stochastic events occur one after another at a PoI, and the inter-arrival times between successive events are i.i.d. The renewal process can model a wide range of real-world events. For example, the Poisson process works well for the arrivals of many independent events; the Weibull distribution applies for channel fading in wireless communications; and the Pareto distribution applies for certain real-world network and computing workloads. Let random variable X denote the time interval between two successive events, and its **pdf** and **cdf** be given by $f(x)$ and $F(x)$, respectively, where $F(x) = Pr(X < x) = \int_{t=0}^{x} f(t)dt$. Since $f(x)$ can be an arbitrary distribution, it can be either continuous or discrete, and the renewal process is quite general. We assume that time is divided into slots of a fixed duration, and it is not possible for more than one events to occur in the same slot. If an event occurs in the time interval $(t-1, t]$, we say that it occurs in slot t. We focus on a certain PoI, and assume that an event occurs there in slot 0.

We consider a deployment of N rechargeable sensors around the PoI. The goal of the sensors is to capture any events appearing at the PoI. Furthermore, it is important to capture the events as soon as they arrive. Immediate capture will allow harmful events in the real world (e.g., leaks in a water system) to be handled immediately before they can do significant damage.

We assume that each rechargeable sensor owns an energy bucket of capacity K (in energy units). In each time slot, environmental sources will recharge the sensor with e_t units of energy, where $e_t \geq 0$ is a random variable with expectation $\mathbb{E}(e_t) = e > 0$ and variance σ^2 ($\sigma \ll K$), but the exact recharge process is unknown. Assume that each sensor consumes δ_1 energy for sensing if it is active in a time slot, and that it consumes zero energy if it is inactive or asleep. Further assume that the sensor consumes an additional δ_2 ($\delta_2 \geq \delta_1$) energy in this time slot if it captures an event, so that the sensor will take an activation decision only when it has at least $\delta_1 + \delta_2$ energy. For clarity of exposition, in this paper we use "in slot t" to mean the time interval $(t-1, t]$. We assume that the recharge process is completed at the very beginning of each slot, and that the sensor makes a decision to activate or not after the recharge.

The state of sensor s in time slot t can be characterized by four components $(B_t^s, A_t^s, V_t, O_t^s)$. $B_t \in \{0, 1, \ldots, K\}$ is the sensor's battery level at the beginning of slot t; $A_t \in \{0, 1\}$ is the sensor's state in slot t, which equals 1 if the sensor is active in the slot and equals 0 otherwise; $V_t \in \{0, 1\}$ is an indicator function of event occurrence in slot t, which equals 1 if an event occurred in slot t and equals 0 otherwise; O_t is the sensor's observation in slot t, where $O_t = 1$ if $V_t \times A_t = 1$, $O_t = 0$ if $V_t = 0$ and $A_t = 1$, and $O_t = \phi$ if $A_t = 0$. So, $O_t = 1$ means that the event was captured by the sensor in slot t.

Let the total time interval of the sensors' deployment be $(0, T]$, where T is extremely large. Then $\sum_{t=1}^{T} V_t$ is the total number of event occurrences during $(0, T]$. Given an activation policy π for the N sensors in the deployment, we define the system's quality of monitoring (QoM) as the event capture probability

$$U_K(\pi) = \lim_{T \to \infty} \frac{\sum_{t=1}^{T} V_t \times max_s\{A_t^s\}}{\sum_{t=1}^{T} V_t}. \quad (1)$$

Here $max_s\{A_t^s\}$ means that even if more than one sensors detected the event in slot t, we only count the capture once. As the average recharge rate e of a sensor should be high enough to support the discharge rate, our goal in this paper is to find an efficient activation policy for the sensors to maximize the capture probability $U_K(\pi)$, i.e. find $\pi^* = \arg\max_{\pi} U_K(\pi)$ such that π is *energy balanced*.

A summary of the notations in this paper is given in Table I.

e	Each sensor's average recharge rate (energy/slot).
K	Each sensor's battery capacity.
X	Random variable denoting the time interval between two consecutive events. It follows a distribution with **pdf** $f(x)$ and **cdf** $F(x)$.
$\pi(e)$	An activation policy with average discharge rate e.
$\Pi^{EB}(e)$	The set of all activation policies with average discharge rate e, i.e. $\{\pi(e)\}$.
$U(\pi(e))$	Ideal QoM (i.e., event capture probability) of policy $\pi(e)$ under the *energy assumption*.
$U_K(\pi(e))$	Practical QoM if policy $\pi(e)$ is employed by sensors each with battery capacity K.
$\pi_{FI}^*(e)$	Optimal activation policy with average discharge rate e under full information. It is defined as a vector $\{c_i\}$, where c_i is the probability that sensor will take activation action a_1 in state h_i, i.e., $Pr(a_1 \vert h_i)$.
$\pi_{PI}^*(e)$	Optimal activation policy with discharge rate e under partial information.
$\pi_{PI}'(e)$	Heuristic activation policy with discharge rate e under partial information. It is defined as a vector $\{c_i\}$.
h_i	Event state under full information, $i \in \mathbb{Z}^+$.
$\mathcal{H}$	The complete state space $\{h_i\}$ under full information.
$\mathcal{F}(\hat{\mathcal{F}})$	Complete (modified) state space under partial information.
$H_t(F_t)$	Event state in slot t under full/partial information, $H_t \in \mathcal{H}$, $F_t \in \mathcal{F}$.

Table I: Notations.

B. Basic Approach

In this paper we consider two different information scenarios: *full information* and *partial information*. In the full information case, we assume that if a sensor is not active in a time slot, it will not be able to capture necessary details about an event that occurred in the time slot. However, the sensor will still know that the event has occurred after the fact, i.e., at the end of the time slot, presumably because the event left a mark that is easy to identify. In the partial information case, the sensor will be aware of an event's occurrence only when the sensor was active in the slot concerned. The design objective for our optimal activation policy can be formulated as a Markov control problem with an average reward criterion [20]. By doing so, we can analyze the full-information and partial-information problems using a Markov decision process. In the next sections, we will first construct full- and partial-information policies for the *single-sensor* problem, where there is only one sensor monitoring the PoI (i.e., $N = 1$). Then, we will extend the policies to accommodate the case of multiple sensors.

IV. Single-Sensor Problems

In this section, we assume that the PoI is monitored by one sensor only. To be complete, the state of a sensor in the MDP should include the sensor's energy level, an indicator function for the event's occurrence, and the sensor's activation history, i.e., $\{(B_t, V_t, A_t)\}$ (here, we omit the superscript s, since only one sensor is involved). In practice, however, it is hard to keep track of the dynamic energy level accurately. Furthermore, determining the transition probabilities using B_t will require more statistical information about the recharge process e_t than we will assume. To simplify the problem, we will develop our activation policies based on the following approach:

1) First, assume that in every slot, the residual energy $B_t \geq \delta_1 + \delta_2$. We call this the *energy assumption*.
2) Neglect B_t and e_t. Let $\Pi^{EB}(e)$ be the set of all policies which are energy balanced, i.e., the long-term rate of energy consumption by a sensor under the policy is equal to the average recharge rate e. Under the energy assumption, we exploit known event dynamics to find the optimal policy $\pi^*(e) \in \Pi^{EB}(e)$ that maximizes the probability of event capture.
3) Furthermore, evaluate the practical performance of policy $\pi^*(e)$ without the energy consumption.

Our approach is extended from that in [6]. Note that, under the energy assumption, $\pi(e) \in \Pi^{EB}(e)$ can be successfully executed by the sensor. In reality, however, the execution of policy $\pi(e)$ by a particular sensor will be subject to sufficient residual energy B_t at all time. Although $\pi(e)$ is energy balanced in the long term, it is possible that for some certain slots, the policy $\pi(e)$ prescribes a decision to activate, but the sensor does not currently have enough energy. We will show that, under certain conditions, the probability that such an insufficient energy condition occurs in a slot is small, and the gap between real performance and the theoretical result is very small when K is large enough, so that $\pi^*(e)$ can apply.

A. Full information model

In this subsection, we focus on the full information model. This model applies in certain real-life settings such as handling leaks in a water distribution network – a leak may cause significant damage if it is not captured upon arrival, but the damage may have lasting effects to enable easy after-the-fact detection (e.g., water stains left with damaged objects).

As we discussed above, we will first use the energy assumption and neglect B_t. Based on that, we develop an optimal policy $\pi^*_{FI}(e) \in \Pi^{EB}(e)$ and assume that $\pi^*_{FI}(e)$ can be successfully executed by the sensor. Two

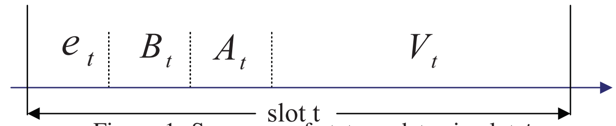

Figure 1: Sequence of state updates in slot t.

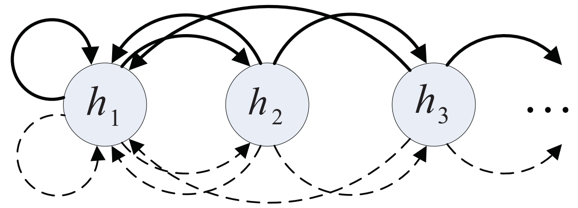

Figure 2: The transition process between states $\{h_i\}$ in the MDP.

important probabilities α_i and β_i will be used throughout our exposition:

$$\alpha_i = Pr(i-1 < X \leq i) = F(i) - F(i-1), \tag{2}$$

$$\beta_i = Pr(X \leq i | X > i-1) = \frac{\alpha_i}{1 - F(i-1)}, \tag{3}$$

where α_i is the probability that, after an event renewal in slot 0, the first subsequent event occurs in slot i, and β_i is the conditional probability that the first event did not occur during $(0, i-1]$, but that it will occur in the next slot i.

1) An MDP formulation: For ease of exposition, we fix the sequence of state updates in each time slot as shown in Figure 1. Specifically, first the recharge process e_t is completed and the sensor's energy level is updated; then the sensor makes an activation decision A_t according to its state, and finally the event V_t, if it occurs, will do so after the decision. So, in slot $t+1$, $B_{t+1} = B_t - A_t(\delta_1 + V_t\delta_2) + e_{t+1}$. We now develop a constrained MDP with average criterion.

State space $\mathcal{H}$: If the sensor is at the beginning of slot t and the latest event occurred in slot $t-i$ ($i \geq 1$), we say that the sensor is in state h_i. Hence h_i means that the event has not occurred for the last $i-1$ slots, and h_1 means that an event occurred in the previous slot. At any time t, the sensor needs to count only the distance i: $min\ \{i | V_{t-i} = 1,\ i \geq 1\}$. This distance is independent of how many times the sensor has activated before slot t. Furthermore, under *full information*, the sensor always knows the value of i, i.e., the sensor can identify its state h_i. If the sensor takes an activation action, it will capture the event w.p. β_i. We get a set $\mathcal{H} = \{h_i, i \in \mathbb{Z}^+\}$ over i, and $\mathcal{H}$ forms the MDP's state space.

Transition probabilities and the reward function: Let $\mathcal{A}$ be the sensor's action space containing two elements: active (denoted by a_1) and inactive (denoted by a_2). We can obtain the transition process within $\mathcal{H}$ as shown in Figure 2, where the solid lines represent transitions based on a_1 and the dotted lines represent those based on a_2. The transition process is a countable-state Markov chain. The system's transition probability from state $H_t \in \mathcal{H}$ in slot t to state $H_{t+1} \in \mathcal{H}$ in slot $t+1$, given the decision $A_t \in \mathcal{A}$, is given as $p(H_{t+1}|H_t, A_t)$. The reward, $r(H_{t+1}|H_t, A_t)$, gained by the sensor in the transition has value 1 if the sensor captured an event, and has value 0 otherwise.

Objective Function: Our goal now is to compute the optimal policy $\pi^*_{FI}(e) \in \Pi^{EB}(e)$ which achieves value $\sup_{\pi_{FI}\in\Pi^{EB}(e)} \lim_{T\to\infty} \frac{1}{T}\mathbb{E}^{\pi_{FI}}\left\{\sum_{t=1}^{T} r(H_t, A_t)\right\}$. Note that the Markov chain defined by an activation policy is uni-chain. Hence it is sufficient to consider only stationary policies. In the next subsection, we will analyze the structure of $\pi^*_{FI}(e)$ and generate an activation policy that can be implemented with minimal information. The policy will also give us insights on how the sensor's energy should be allocated to maximize the event capture probability.

2) Structure of the optimal policy: Let $M(T)$ denote the expected number of events occurring within $(0, T]$, and $\mu = \sum_{i=1}^{\infty} i\alpha_i$ be the expectation of the inter-arrival times. We assume that the sensor activates itself with probability c_i in every state h_i, i.e., $c_i = Pr(a_1|h_i)$, and set $\pi_{FI}(e) = (c_1, c_2, \cdots)$ as a stationary policy. Given the recharge process, the sensor parameters, and the event process, the problem is to find the optimal vector $\pi^*_{FI}(e)$ such that the event capture probability is maximized without violating the energy balance.

Consider a specific event inter-arrival time X. The expected number of activations during X, denoted by $n(\pi_{FI}(e))$, is

$$\begin{aligned} n(\pi_{FI}(e)) &= \mathbb{E}[\mathbb{E}(\sum_{j=1}^{i} c_j)|X=i] = \sum_{i=1}^{\infty}\alpha_i\left(\sum_{j=1}^{i} c_j\right) \\ &= \sum_{i=1}^{\infty} c_i[1-F(i-1)] \end{aligned} \tag{4}$$

where we have used the identity that $\sum_{i=1}^{\infty}\alpha_i = 1$. Let the sensor's operation interval be $(0, T]$, where T is extremely large. Then the amounts of energy received and consumed during $(0, T]$ are, respectively,

$$E_{in} \triangleq Te + K, E_{out} \triangleq M(T)[\delta_1 n(\pi_{FI}(e)) + \delta_2 \sum_{i=1}^{\infty}\alpha_i c_i].$$

In the long run, it is necessary to ensure that $E_{in} = E_{out}$ so that the sensor has sufficient energy to sustain operations, which gives

$$M(T)[\delta_1 n(\pi_{FI}(e)) + \delta_2 \sum_{i=1}^{\infty}\alpha_i c_i] = Te + K. \tag{5}$$

Dividing (5) on both sides by T and taking a limit, we obtain

$$\lim_{T\to\infty}\frac{M(T)}{T}[\delta_1 n(\pi_{FI}(e)) + \delta_2 \sum_{i=1}^{\infty}\alpha_i c_i] = e + \lim_{T\to\infty}\frac{K}{T}$$

$$\sum_{i=1}^{\infty}[\delta_1(1-F(i-1)) + \delta_2\alpha_i]c_i = e\mu, \tag{6}$$

where the result $\lim_{T\to\infty} M(T)/T = 1/\mu$ is used.

Therefore, under condition (6), the optimal decision vector $\pi^*_{FI}(e)$ can be obtained by solving the following linear program:

$$Max \qquad U(\pi_{FI}(e)) = \sum_{i=1}^{\infty}\alpha_i c_i \tag{7}$$

$$\begin{aligned} s.t. \qquad & \sum_{i=1}^{\infty}\xi_i c_i = e\mu \\ & 0 \le c_i \le 1, \end{aligned} \tag{8}$$

where $\xi_i \triangleq \delta_1(1-F(i-1)) + \delta_2\alpha_i$ from (6).

Note that the linear program may have an infinite number of variables, and is not solvable in general. A truncation approach could be used to obtain a good approximation. However, it is quite useful to investigate the conceptual structure of the optimal policy. To do so, let us look at a simple example. Suppose we have a simple event process with two per-slot conditional probabilities of event occurrence, namely $\beta_1 = 0.6$ and $\beta_2 = 1$, which means that the event process **pdf** is $\alpha_1 = 0.6$, $\alpha_2 = 0.4$. Suppose that 800 events appear consecutively. If the sensor always activates in slot 1, it will activate 800 times and is expected to capture 480 events. On the other hand, if the sensor always activates in slot 2, it is expected to activate 320 times and capture 320 events. Thus if there is not enough energy for more than 320 activations, it is always more beneficial to allocate all the energy to slot 2 (the efficiency is 100%). But if there is surplus energy, we should then allocate it to slot 1 (where the efficiency is 60%). More generally, we have the following result.

Theorem 1: If the per-slot conditional probabilities of event occurrences are increasing in the slot number, i.e., $\beta_1 \le \beta_2 \le \cdots \le \beta_i \le \cdots$, and there exists an index k such that $\sum_{i=k+2}^{\infty}\xi_i < e\mu \le \sum_{i=k+1}^{\infty}\xi_i$, then the optimal policy can be specified as

$$\pi^*_{FI}(e) = (0, \cdots, 0, c_{k+1}, 1, 1, \cdots 1, \cdots), \tag{9}$$

where $c_{k+1} = \frac{1}{\xi_{k+1}}\left[e\mu - \sum_{i=k+2}^{\infty}\xi_i\right]$. The optimal capture probability achieved by policy $\pi^*_{FI}(e)$, under the energy assumption, is $U(\pi^*_{FI}(e)) = 1 - F(k+1) + c_{k+1}\alpha_{k+1}$.

Proof: See Appendix A. ■

It can be seen that in this optimal policy, the recharge energy is first allocated to the slot having the highest conditional probability β_i. If there is surplus energy remaining afterwards, then it is allocated to the slot having the next highest conditional probability, and so on. When $e = \delta_1 + \delta_2/\mu$, all the elements in (9) is 1. The sensor can always take action a_1. The decision rule corresponds to a "greedy" policy in which activation energy is allocated to those slots with the highest conditional probabilities of event occurrences. Under the energy assumption, the policy $\pi^*_{FI}(e)$ can always execute successfully, and the sensor will achieve the optimal event capture probability $U(\pi^*_{FI}(e))$.

Remark 1: It is easy to generalize Theorem 1 in concept to the case in which the β_i's are not monotonically increasing. This is accomplished by first sorting the β_i's in increasing order. Then, we update α_i and ξ_i according to the sorted β_i's. Theorem 1 now applies readily.

In the development so far, we have ignored the sensor's changing energy levels B_t. Policy $\pi^*_{FI}(e)$ and its optimal capture probability $U(\pi^*_{FI}(e))$ are obtained under the energy assumption. We now prove that, even when changes in B_t are considered (i.e., without the energy assumption), the policy $\pi^*_{FI}(e)$ is asymptotically optimal with respect to the energy bucket size K. Let $U_K(\pi^*_{FI}(e))$ be the capture probability under policy $\pi^*_{FI}(e)$ when it is applied by a sensor with an energy bucket of capacity K. We have this result: $\lim_{K\to\infty} U_K(\pi^*_{FI}(e)) = U(\pi^*_{FI}(e))$. The result holds because Theorem 1 implies that

$$\lim_{T\to\infty} \mathbb{E}^{\pi^*_{FI}(e)} \left\{ \frac{\sum_{t=1}^{T} V_t \times A_t}{\sum_{t=1}^{T} V_t} \right\} = U(\pi^*_{FI}(e)),$$

where A_t is determined by $\pi^*_{FI}(e)$. Comparing the sensor energy to a server queue, we can view the energy process as job arrivals and departures in a $G/G/1$ queueing system. Under policy $\pi^*_{FI}(e)$, the service rate is equal to the arrival rate, and so the fraction of time the server is idle is an infinitesimally small value ϵ [21]. It follows that

$$\begin{aligned} & U_K(\pi^*_{FI}(e)) \\ &= \lim_{T\to\infty} \mathbb{E}^{\pi^*_{FI}(e)} \left\{ \frac{\sum_{t=1}^{T} V_t \times (1-\epsilon) A_t}{\sum_{t=1}^{T} V_t} \right\} \\ &= (1-\epsilon) U(\pi^*_{FI}(e)). \end{aligned}$$

As $\lim_{K\to\infty}(1-\epsilon) = 1$, we get $\lim_{K\to\infty} U_K(\pi^*_{FI}(e)) = U(\pi^*_{FI}(e))$.

Remark 2: The time-average utility $U_K(\pi^*_{FI}(e))$ converges to the optimal value as K becomes large, which is independent of the energy discharge and recharge processes. It is clear that for the energy bucket not to overflow, K has to be large enough. This means that we need a large buffer to absorb the energy bursts. An infinite battery size is physically impossible, however. In this paper, we will quantify the sufficiency and impact of practical K values by simulations in Section VI.

B. Partial information model

In certain practical scenarios, the sensor cannot determine if an event has occurred or not in a slot unless it is active in the slot. In this case, the sensor may have observations $O_t = \phi$. We call this the partial-information model, and solve it by Markov control theory in this section. Our discussion is structured as follows. In Section IV-B1, we formulate the problem as a partially observable MDP (POMDP), and analyze the structure of the POMDP optimal policy $\pi^*_{PI}(e)$ and its computational complexity. In Section IV-B2, we provide an efficient heuristic clustering policy $\pi'_{PI}(e)$ inspired by a main practical property of $\pi^*_{FI}(e)$.

1) POMDP formulation: In the full-information case, when an event occurs in a time slot, the system will move to state h_1 whether the sensor activated or not in the time slot. In the partial-information case, however, the internal state of the underlying process is unknown by default. The sensor's knowledge of the internal state is totally based on its past observations and decision history. The problem can be formulated as a POMDP, but its solution is more complex than the full-information counterpart.

Typically, the key to finding truly optimal policies in POMDP is to cast the problem as an equivalent completely observable MDP. Assume that the sensor is at the beginning of slot i, and the sensor captured an event in slot 0 but did not capture any events during slots 1 to $i-1$. If the sensor has only captured an event in slot 0, then the sensor's full history of observations is:

$$f_i = (O_0 = 1, O_1, O_2, \cdots, O_{i-1}), \qquad (10)$$

where $O_0 = 1$ means that the sensor observed an event in slot 0, and $O_j \in \{0, \phi\}$, $j = 1, \dots, i-1$. If the sensor was not active for k of the $i-1$ slots, then there are 2^k different possible sequences of event occurrences (1 means an event occurred, and 0 means otherwise) that are consistent with the sensor's observations, each of which is denoted by $f_{i,j}$, $j = 1, \dots, 2^k$. An example for $i = 3$ and $k = 2$ is illustrated as: $f_{3,1} = (1,0,0)$, $f_{3,2} = (1,0,1)$, $f_{3,3} = (1,1,0)$, $f_{3,4} = (1,1,1)$. Here we notice the extra complexity of a renewal process over a Markov chain: In a Markov chain, if the sensor knows about the event occurrence in slot $i-1$, the probability of event occurrence in slot i is completely defined, whereas in a renewal process, the probability depends on all the observed *and* unobserved event occurrences since the last renewal defined by the last observed event.

Let $\mathcal{F}_t = \{f_{i,k}\}_{i\le t}$ be the set of possible past event states in slot t. Since the events have the renewal property, $\mathcal{F}_t$ can represent the sensor's total information set. The POMDP's state space is then $\mathcal{F} = \bigcup_{t>0} \mathcal{F}_t$, and $\mathcal{F}$ is sufficient to determine the sensor's best actions. Assume that the information set in slot t is $F_t \in \mathcal{F}_t$. Associating the transition probabilities $p(F_{t+1}|F_t, A_t)$ with the achieved reward $r(F_t, A_t)$, the POMDP can be transformed into an equivalent MDP with an infinite state space $\mathcal{F}$. Since this MDP is uni-chain, an optimal policy $\pi^*_{PI}(e)$ can be constructed.

2) Heuristic clustering policy: Although $\pi^*_{PI}(e)$ can be obtained in principle by an equivalent MDP, finding it is intrinsically hard in practice. The reason is that the dimension of the information set $\mathcal{F}_t$ grows exponentially with time t, which makes the problem intractable. Even good truncation methods are quite difficult to obtain. Furthermore, since F_t may be quite complex, it is expensive to calculate the transition probabilities $p(F_{t+1}|F_t, A_t)$. In fact, because of this "curse of dimensionality," the computational problem is even less likely to have polynomial-time solutions than NP-complete problems [7], [8]. We are thus motivated to find an effective heuristic solution in practice.

The structure of $\pi^*_{FI}(e)$ tells us that it is advisable to spend energy in slots where events are most likely to occur. We call the region where the events are most likely to occur the *hot region*. The sensor should focus its attention in this region.

Further, for an arbitrary renewal process characterized by the distribution of event inter-arrival times X, we can calculate the probability of event occurrence in every slot i as β_i. We believe that for many renewal processes, the β_i's concentrate in a hot region. (An important exception is the Poisson process, whose β_i's are constant.) Therefore, in order to maximize the capture probability, we propose the following clustering policy

$$\begin{aligned}\pi_{PI} &= (0, \cdots, 0, c_{n_1}, 1, 1, \cdots, 1, c_{n_2}, 0, \cdots, 0, \\ &\quad c_{n_3}, \text{ aggressive activation policy}) \qquad (11)\end{aligned}$$

where c_{n_1},c_{n_2},$c_{n_3} \in [0, 1]$, $n_1 \le n_2 \le n_3$, and the "aggressive policy" is one in which the sensor always activates whenever it has sufficient energy to do so. This heuristic policy implements the following ideas:

- The slots in the range $n_1 \le i \le n_2$ allow the sensor to take action a_1 with high priority within the hot region.
- It is possible for a partial-information sensor to miss an event even if it was active in the hot region (cf. the full-information model in which event occurrences are always known). To recover from the missed information, we use $c_i, i \ge n_3$ to activate the sensor aggressively (i.e., activate whenever there is sufficient energy to do so) until a new event is captured. A captured event will renew the sensor's operation schedule and therefore provide recovery. We call the slots $n_3, \ldots, +\infty$ the *recovery region*.
- Since the energy recharge rate is not enough to sustain continuous activation of the sensor, the sensor must sleep in less critical time slots to accumulate enough energy for energy-intensive operation in the hot region. The slots with $c_i = 0$ form what we call the *cooling region*.

Optimization of the sensor activation policy consists in determining suitable values for n_1, n_2, and n_3 in order to divide the sensor's operation into the cooling, hot, and recovery regions.

Under policy π_{PI}, if the sensor is at the beginning of slot t and the latest *captured* event occurred in slot $t - i$, we say that the sensor is in state f_i. The state f_i is different from the state h_i in the full-information case. It means that the sensor has not captured an event for $i - 1$ slots, and f_1 means that an event was captured in the last slot. Let $\hat{\mathcal{F}} = \bigcup f_i$ denote a new state space. Whenever an event is captured, the sensor "renews" to state f_1 automatically. Let $\mathbf{P} = (p(f_j|f_i))$ denote the transition matrix within $\hat{\mathcal{F}}$. When policy (11) is employed, we have

$$p(f_j|f_i) = \begin{cases} c_i\beta_i, & j = 1, \\ 1 - c_i\beta_i, & j = i + 1, \\ 0, & otherwise, \end{cases}$$

where the conditional probabilities β_i are given in Appendix B. The term $p(f_1|f_i)$ is the probability that the sensor will capture an event in state f_i. Next we calculate the optimal policy under (11), denoted by $\pi'_{PI}(e)$.

Let $Y = (y_1, y_2, \cdots)$ denote the stationary distribution of $\hat{\mathcal{F}}$. Then Y can be obtained by $Y\mathbf{P} = Y$ and $\sum y_i = 1$. Let t_1 and t_2 be two successive slots where the sensor captured an event (note that the sensor may miss some events between t_1 and t_2). In this case, $1/y_1$ equals the expected value of $t_2 - t_1$, and so the capture probability under policy (11) is $U(\pi_{PI}) = y_1\mu$. The expected energy consumed in one slot is $E_{out}(\pi_{PI}) = \sum y_i c_i(\delta_1 + \beta_i\delta_2)$. Therefore, our goal is to find the optimal c_{n_1}, c_{n_2}, and c_{n_3} that maximize $\{U(\pi_{PI})|E_{out}(\pi_{PI}) \le e\}$. A truncated method for dynamic programming can be used: increase n_3 gradually and enumerate n_1 and n_2 in $[1, n_3]$ until $max\ \{U(\pi_{PI})|E_{out}(\pi_{PI}) \le e\}$ cannot be further increased. Furthermore, from π_{PI}'s structure, we know that n_3's increase will be bounded. Otherwise, the sensor would sleep forever from n_2 (n_2 is likewise bounded, except if the sensor has infinite initial energy or $e \ge \delta_1 + \delta_2/\mu$).

The above heuristic policy tries to exploit knowledge of the conditional probabilities β_i's. At the same time, it implements safeguard to recover from any missed information. The heuristic policy is simple to implement and its decisions are almost deterministic. The activation policy needs significantly less memory to store than the truly optimal policy. In particular, it can be implemented by a resource-constrained sensor using local state only.

In the simulations in Section VI, we will show that the clustering policy can significantly outperform two plausible alternative policies: the periodic policy π_{PE} and aggressive policy π_{AG}. In the periodic policy, the sensor activates itself for θ_1 slots every $\theta_2 \ge \theta_1$ slots. In the aggressive policy, the sensor activates itself whenever $B_t \ge \delta_1 + \delta_2$. Also, we remark that $\pi'_{PI}(e)$ is a coarse method. A finer approximation for the optimal solution $\pi^*_{PI}(e)$ can be obtained by augmenting the decision type as follows: introduce transition points $c_{n_4}, c_{n_5}, \ldots$, etc., after c_{n_3}. This way, we will get progressively more detailed policies which converge to $\pi^*_{PI}(e)$. The tradeoff is that more finely controlled policies will be more costly to implement. Further, recall that the policy $\pi'_{PI}(e)$ is developed under the energy assumption. Similar to the result in Section IV-A, $\pi'_{PI}(e)$ satisfies the asymptotic property: $U_K(\pi'_{PI}(e))$ approaches $U(\pi'_{PI}(e))$ as K approaches infinity.

V. Multi-Sensor Problem

The energy recharge rate of a single sensor may be limited and not enough to achieve the required QoM of practical applications. To overcome the problem, multiple rechargeable sensors may be deployed to monitor events at the same PoI to improve the QoM. In this section, we assume that $N > 1$ identical sensors are used to monitor a PoI, and the QoM is defined as the aggregate event capture probability given in (1).

A simple way to use N sensors is to let them work independently without any coordination or information exchange. Specifically, each sensor will follow an activation policy proposed in the previous sections according to its own information state. Without coordination, however, the sensors are prone to activating at the same time slots and duplicate each other's efforts in the monitoring. Hence, our objective is to use a coordination protocol among the sensors in order to avoid redundant activations.

To do so, we assign the sensors to time slots in a renewal period in a round robin manner. For example, for two sensors 1 and 2, we let 1 take charge of the odd time slots and 2 of the even ones. We now design an activation policy for each sensor under the full-information and partial-information models, respectively.

A. Full information model

In the full-information model, we assume that each sensor knows about the occurrence of an event in a time slot whether it was active or not in that slot. We can therefore treat all the sensors as one logical "big" sensor, and derive an activation policy for this big sensor with energy recharge rate $N \times e$. We call this activation policy under full information M-FI, which is described as follows:

- **Step 1** Given δ_1, δ_2, T, N, e, and the greedy policy $\pi^*_{FI}(Ne)$ by Theorem 1, where Ne is the aggregate recharge rate of the N sensors, index each sensor from 1 to N. Assume that an event occurred in slot 0, and initialize the system with $t = 1$.
- **Step 2** In slot t, for $t = kN + s$ ($k \in \mathbb{Z}^+$, $1 \le s \le N$), sensor s takes charge of the slot and calculates the event state H_t. *(All the other sensors are inactive in this slot.)*. Note that $H_t = h_{t-i}$, where i is the slot in which the latest event occurred. All the sensors know i under full information.
- **Step 3** Within each of its assigned slots, sensor s decides to activate according to policy $\pi^*_{FI}(Ne)$, based on the state H_t.
- **Step 4** $t \leftarrow t + 1$; go to Step 2.

We use a simple example with two sensors to illustrate the collaborative sensor schedule. We assume that the greedy policy obtained by Theorem 1 in this example is

$$\pi^*_{FI}(2e) = (0, 0, 1, 1, 1, \ldots). \tag{12}$$

A trace of the activation schedule for 7 initial slots is given below (recall that an initial event has occurred in slot 0).

slot t:	1	2	3	4	5	6	7
sensor s:	1	2	1	2	1	2	1
event occurrence V_t:	0	0	0	1	0	1	0
event state H_t:	h_1	h_2	h_3	h_4	h_1	h_2	h_1
sensor 1's action A^1_t:	a_2	I	a_1	I	a_2	I	a_2
sensor 2's action A^2_t:	I	a_2	I	a_1	I	a_2	I

In the table, the symbol *I* means that the sensor is not responsible for the time slot and thus stays inactive. In slot 1, where sensor $s = 1$ takes charge, by policy (12), since $c_1 = 0$, sensor 1 is inactive. In slot 2, sensor 2 becomes responsible while sensor 1 keeps inactive. Since $H_2 = h_2$ and $c_2 = 0$, sensor 2 decides to stay inactive in this slot. In slot 3, sensor 1 becomes responsible again. It decides to activate (since $c_3 = 1$), but an event did not occur in this time slot. In slot 4, sensor 2 is responsible and decides to activate (since $c_4 = 1$) and captures an event. The collaborative sensor schedule gets renewed in slot 5. Note that under full information, the sensors always know about the occurrences of past events, independent of whether they were active or not. Therefore, in slot 7, the event state H_7 still gets updated to h_1, although the event in slot 6 was not captured by any sensors.

It should be noted that load balancing is another important issue for multi-sensor activation policies. This is because in the single-sensor case, the sensor can store extra energy gained in its low-duty regions to cover the energy deficits of high-duty regions. With multiple sensors, if their energy use is imbalanced due to uneven division of responsibilities, it will not be physically possible to transfer the extra energy stored in one sensor to cover the needs of another. The simple M-FI policy cannot guarantee load balancing in general. For example, if $\beta_1 = 0$ and $\beta_2 = 1$ and there are two sensors, one of the two sensors will be doing all the work. This problem can be mitigated by letting the sensors round robin over time slots with $\beta_i > 0$ only. With this change, then under the energy assumption, the sensors will achieve load balancing over a working window of T slots when T is large enough. Practically, the simulation results in Section VI show that M-FI can achieve good load balancing for a number of "natural" event distributions such as Weibull and Pareto.

B. Partial information model

Under partial information, we assume that the sensors know about the occurrence of an event only when at least one of them was active in the time slot of the event. Specifically, a sensor that captured an event will report it to the sink node, which will then notify all the sensors by broadcasting over a low-power channel, so that the energy consumed for receiving the notification is negligible.

We now give an M-PI policy for multi-sensor activation under partial information. Note that in M-PI, the value of i in Step 2 of the M-FI policy is not known exactly, and we use the heuristic clustering policy in Section IV-B to guide the operations of the sensors. The M-PI policy is stated as follows:

- **Step 1** Calculate $\pi'_{PI}(Ne)$ according to the heuristic clustering policy in Section IV-B, where Ne is the aggregate recharge rate of the sensors. Assume that an event occurred in slot 0, and initialize the system with $t = 1$.
- **Step 2** In slot t, for $t = kN + s$ ($k \in \mathbb{Z}^+$, $1 \le s \le N$), sensor s takes charge and calculates the event state F_t. Note that $F_t = f_{t-i}$, where i is the latest slot in which an event was captured.
- **Step 3** Within each of its assigned slots, sensor s decides to activate according to policy $\pi'_{PI}(Ne)$ based on state F_t.
- **Step 4** $t \leftarrow t + 1$; go to Step 2.

VI. SIMULATIONS

We evaluate the proposed activation policies using numerical results for a wide range of parameters, for both cases of single and multiple sensors. We use $\delta_1 = 1$ and $\delta_2 = 6$ (in energy units) as the energy parameters of sensing. We measure the event capture probabilities achieved over a long working duration of $T = 10^6$ time slots. Simulations are carried out for both the full- and partial-information models, with varying energy bucket size K and the energy recharge process e_t following one of three different models. Each reported experiment uses the same distribution of event inter-arrival times X, which is either Weibull $W(\eta_1, \eta_2)$ or Pareto $P(\gamma_1, \gamma_2)$. Specifically, the **pdf**s are, respectively,

$$f(x) = \frac{\eta_2}{\eta_1}(\frac{x}{\eta_1})^{\eta_2-1}exp(-(\frac{x}{\eta_1})^{\eta_2}), x > 0, \eta_1, \eta_2 > 0.$$

$$f(x) = \frac{\gamma_1 \gamma_2^{\gamma_1}}{x^{\gamma_1+1}}, x \geq \gamma_2, \gamma_1, \gamma_2 > 0.$$

A. Single-sensor experiments

1) Asymptotic property: In this part, we will show that, when the battery capacity K increases, the capture probabilities $\pi^*_{FI}(e)$ and $\pi'_{PI}(e)$ achieved in practice will approach their respective optimal values derived under the energy assumption.

Let the average recharge rate be $e = 0.5$, and the event inter-arrival times X follow the Weibull distribution $W(40, 3)$. The simulation results for the three different recharge processes, namely Bernoulli, Periodic, and Uniform, are shown in Fig. 3. For the Bernoulli process, the environment recharges the battery with $c = 1$ unit of energy w.p. $q = 0.5$ in each slot. In the Periodic process, the battery is recharged with 5 units of energy in every 10 slots. And the Uniform process recharges the battery with 0.5 unit of energy in every slot.

For varying energy bucket capacity K, Fig. 3(a) shows the achieved performance $U_K(\pi^*_{FI}(e))$ of the optimal policy $\pi^*_{FI}(e)$ under full information, whereas Fig. 3(b) shows the corresponding achieved performance of policy $\pi'_{PI}(e)$ under partial information. We have proved that computing the exact $\pi^*_{PI}(e)$ is computationally hard. We now present results achieved by the clustering policy $\pi'_{PI}(e)$ and verify their asymptotic property. It can be observed that both $U_K(\pi^*_{FI}(e))$ and $U_K(\pi'_{PI}(e))$ converge to their theoretical asymptotic optimal as K increases, which agrees with our analysis, and they appear to do so quickly. Note that the convergence is independent of the recharge process, which shows the robustness of the proposed policy to detailed statistics of the recharging. In practical applications, we may need to consider only the event process and the average recharge rate e, and then we can achieve good performance for a practical K that is large enough.

2) Comparison between different policies: We assume that the energy bucket size is constant, given by $K = 1000$ (in energy units), and provide the sensor with $K/2$ units of initial energy at time 0. We assume that the recharge process is Bernoulli, where we fix $q = 0.5$ and increase c in a number of experiments. In each experiment, therefore, the energy recharge rate is $e = q \times c$. We compare our proposed policy $\pi'_{PI}(e)$ with three alternative policies, namely the aggressive policy π_{AG}, a periodic policy π_{PE}, and a policy π_{EBCW} similar to the one proposed in [6].

We first compare the policy $\pi'_{PI}(e)$, under partial information, with the aggressive and periodic policies. Figs. 4(a) and 4(b) depict the fraction of events captured by the three policies for the two different event types $W(40, 3)$ and $P(2, 10)$, respectively. For the periodic policy, which activates the sensor for θ_1 time slots every θ_2 slots, we fix $\theta_1 = 3$ and calculate $\theta_2(e) = \frac{\theta_1 \delta_1}{e} + \frac{\theta_1 \delta_2}{e\mu}$, where μ is the expectation of the event inter-arrival times to achieve energy balance. From the results, notice that although the clustering policy $\pi'_{PI}(e)$ is coarse-grained, it outperforms either the aggressive or periodic policy.

Next, we compare the clustering policy with a result in [6]. In [6], events driven by a two-state Markov process are analyzed. They assume that the events follow a Markov chain with parameters $a, b > 0.5$, where $a = Pr(1|1)$ and $b = Pr(0|0)$. If we let X be the distance of the current time slot from the last time slot where an event was last captured, their Markov chain can be transformed into a renewal process in our formulation, and we can solve their problem by our policy. However, our policy $\pi'_{PI}(e)$ can be applied to any transition matrix, not limited by $a, b > 0.5$. From Figure 5, it can be seen that when $a, b > 0.5$, our result is the same as [6]. For other values of a, b, $\pi'_{PI}(e)$ outperforms π_{EBCW}.

B. Multi-sensor experiments

Let the recharge process be the same for all the sensors, which is a Bernoulli process with $q = 0.1$ and varying c. The battery size is set as $K = 1000$. Figs. 6(a) and (b) compare the capture probabilities of M-FI and M-PI with those of the aggressive policy π_{AG} and periodic policy π_{PE}, respectively. In a set of experiments, we vary the number of sensors N and the recharge energy c per recharge. In using π_{AG} for multiple sensors, we assign time slots to each sensor as described in Section V. Then each sensor applies the aggressive policy within their own assigned time slots. For policy π_{PE}, we assume that its parameters θ_1 and θ_2 are such that the policy is energy balanced. With multiple sensors, each sensor will take turns to be in charge of θ_2 consecutive time slots, within which the sensor applies the periodic policy. For example, if $N = 2$, $\theta_1 = 3$, and $\theta_2 = 5$, sensor 1 will be in charge of the time slots $t = \{10 * n + i\}$, for $i = 1, \ldots, 5$ and $n = 0, 1, \ldots$, and sensor 2 will be in charge of the remaining slots. In this way, each sensor will be energy balanced.

The results show that both M-PI and M-FI outperform the multi-sensor aggressive and periodic polices over the range of parameters in these experiments. The capture probability of M-PI approaches that of M-FI when either the number of sensors N or the energy per recharge c becomes large enough. It can also be seen that the capture probability obtained by either π_{AG} and π_{PE} increases, almost linearly.

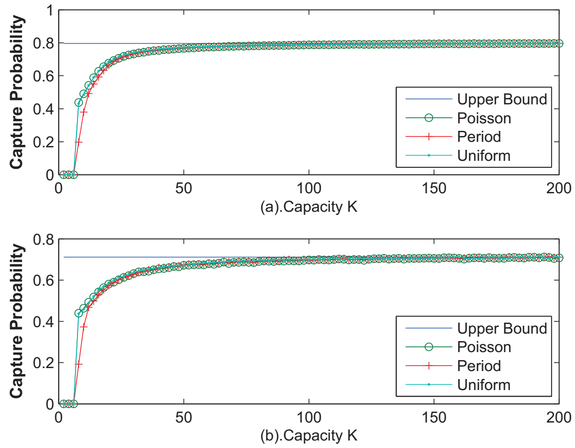

Figure 3: Achieved event capture probabilities of policies π^*_{FI} and π'_{PI}. The event inter-arrival times follow $X \sim W(40, 3)$.

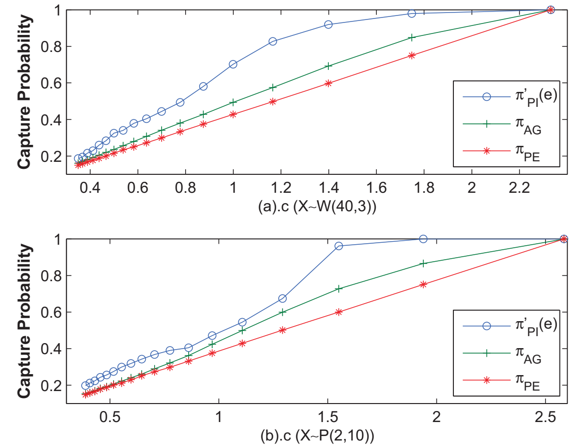

Figure 4: Comparison of policy $\pi'_{PI}(e)$ with other two policies π_{AG} and π_{PE}. $K = 1000$ and $q = 0.5$.

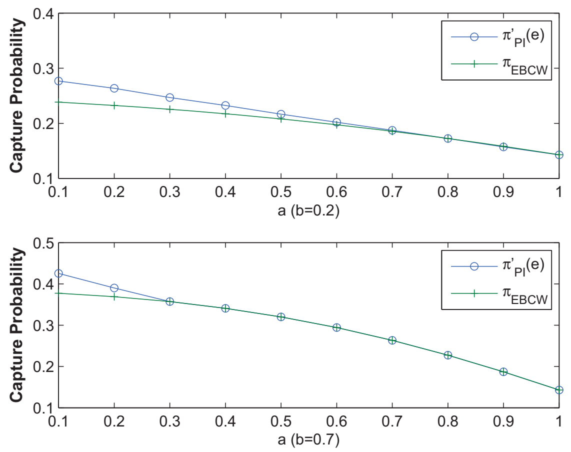

Figure 5: Comparison of policy $\pi'_{PI}(e)$ with π_{EBCW}, where the event follows a Makrov chain, and the recharge process is Bernoulli with q=0.5 and c=2. K=1000.

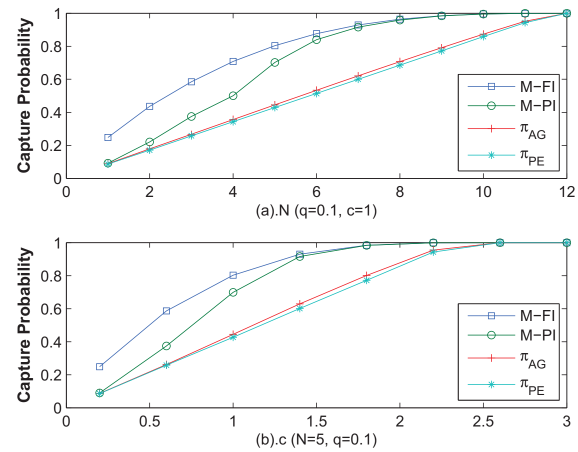

Figure 6: Achieved QoM for increasing N and c. The event inter-arrival times follow $X \sim W(40, 3)$. $K = 1000$.

In comparison, M-FI and M-PI approach 100% capture probability significantly faster, which shows the usefulness of exploiting event memory in the activation.

VII. Conclusion

We have considered the optimization of sensor activation policies for maximum event capture, under the energy constraints of a sensor recharge process. We first solved the case of full information, where a greedy policy exploits memory in the event process to optimize performance. We further analyzed a partial-information model, and showed that finding the optimal solution in this case is computationally intractable. We then proposed a heuristic clustering policy for partial information, and showed that the policy is simple to implement and can outperform alternative aggressive and periodic policies. Lastly, we considered the case of multiple sensors. Simulation results showed that the proposed round robin strategies of assigning responsible sensors to time slots are effective.

Acknowledgments

Research was supported in part by NSFC under grant numbers 61028007 and 61004060, and in part by National High Technology Research and Development Program of China under grant number 2011AA04010. D. Yau was additionally supported in part by U.S. NSF under grant number CNS-0964086.

References

[1] A. Kansal, J. Hsu, S. Zahedi, and M. Srivastava, "Power management in energy harvesting sensor networks," *ACM Transactions on Embedded Computing Systems*, vol. 6, 2007.

[2] J. A. Paradiso and T. Starner, "Energy scavenging for mobile and wireless electronics," *IEEE Pervasive Computing*, vol. 4, pp. 18–27, 2005.

[3] J. Wenck, J. Collier, J. Siebert, and R. Amirtharajah, "Scaling self-timed systems powered by mechanical vibration energy harvesting," *ACM Journal on Emerging Technologies in Computing Systems*, vol. 6, Article 5, 2010.

[4] V. Pryyma, D. Turgut, and L. Boloni, "Active time scheduling for rechargeable sensor networks," *Computer Networks*, vol. 54, pp. 631–640, 2010.

[5] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: A POMDP framework," *IEEE Journal on Selected Areas in Communications*, vol. 25, pp. 589 – 600, 2007.

[6] N. Jaggi, K. Kar, and A. Krishnamurthy, "Rechargeable sensor activation under temporally correlated events," *Wireless Networks*, vol. 15, pp. 619–635, 2009.

[7] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of Markov decision processes," *Mathematics of Operations Research*, vol. 12, pp. 441–450, 1987.

[8] J. N. Tsitsiklis, "Computational complexity in Markov decision theory," *HERMIS-An International Journal of Computer Mathematics and its Applications*, vol. 9, pp. 45–54, 2007.

[9] H. Gupta, S. Das, and Q. Gu, "Connected sensor cover: self-organization of sensor networks for efficient query execution," In *Proc. of ACM MobiHoc*, 2003.

[10] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. B. Srivastava, "Coverage problems in wireless ad-hoc sensor networks," In *Proc. of IEEE INFOCOM*, 2001.

[11] H. Zhang and J. C. Hou, "Maintaining sensing coverage and connectivity in large sensor networks," *Wireless Ad-hoc and Sensor Networks*, vol. 1, pp. 89–124, 2005.

[12] R. E. Lapp and H. L. Andrews, *Nuclear Radiation Physics.* Prentice Hall, 1948.

[13] P. Reynolds, *Call Center Staffing.* The Call Center School Press, 2003.

[14] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic," In *Proc. of ACM SIGCOMM*, 1993.

[15] D. Yau, N. Yip, C. Ma, N. Rao, and M. Shankar, "Quality of monitoring of stochastic events by periodic and proportional-share scheduling of sensor coverage," In *Proc. of ACM CoNext*, 2008.

[16] S. He, J. Chen, D. K. Y. Yau, H. Shao, and Y. Sun, "Energy-efficient capture of stochastic events by global- and local-periodic network coverage," In *Proc. of ACM MobiHoc*, 2009.

[17] A. Seyedi and B. Sikdar, "Energy efficient transmission strategies for body sensor networks with energy harvesting," *IEEE Transactions on Communications*, vol. 58, pp. 2116–2126, 2010.

[18] K. Fan, Z. Zheng, and P. Sinha, "Steady and fair rate allocation for rechargeable sensors in perpetual sensor networks," In *Proc. of ACM Sensys*, 2008.

[19] H. Li, N. Jaggi, and B. Sikda, "Relay scheduling for cooperative communications in sensor networks with energy harvesting," *IEEE Transactions on Wireless Communications*, vol. 10, pp. 2918–2928, 2011.

[20] D. P. Bertsekas, *Dynamic Programming and Optimal Control.* Athena Scientific, 2000.

[21] R. Wolff, *Stochastic Modeling and the Theory of Queues.* Prentice Hall, 1989.

[22] E. P. C. Kao, *An Introduction to Stochastic Processes.* Duxbury Press, 1997.

APPENDIX A.

Proof: [Proof of Theorem 1] Let $\pi^*_{FI}(e) = (c_i)$ be an optimal solution. We need to prove that if $\beta_i \leq \beta_j$, then $c_i \leq c_j$.

Assume to the contrary that $c_i > c_j$ in the optimal solution $\pi^*_{FI}(e)$. We may then design a new policy $\pi'_{FI}(e) = (\cdots, c_i - \Delta, \cdots, c_j + x, \cdots)$, where $\Delta, x > 0$. Using (8), we have

$$\Delta[\delta_1(1 - F(i-1)) + \delta_2\alpha_i] = x[\delta_1(1 - F(j-1)) + \delta_2\alpha_j].$$

The difference between $U(\pi^*_{FI}(e))$ and $U(\pi'_{FI}(e))$ is

$$\begin{aligned} & U(\pi^*_{FI}(e)) - U(\pi'_{FI}(e)) \\ = \; & \alpha_i c_i + \alpha_j c_j - \alpha_i(c_i - \Delta) - \alpha_j(c_j + x) \\ = \; & \Delta\alpha_i - \alpha_j \frac{\Delta[\delta_1(1 - F(i-1)) + \delta_2\alpha_i]}{\delta_1(1 - F(j-1)) + \delta_2\alpha_j} \\ = \; & \frac{\Delta\delta_1(\beta_i - \beta_j)}{\delta_1(1 - F(j-1)) + \delta_2\alpha_j} \\ < \; & 0. \end{aligned}$$

This indicates that $\pi^*_{FI}(e)$ cannot be an optimal solution, which is a contradiction. ■

APPENDIX B.

Proof: Let S_i denote the time epoch (i.e., renewal period) where the i-th event occurred. Define the number of events occurring in $(0, t)$ as $N(t)$. Let $\Psi(t) = S_{N(t)+1} - t$ and $G_t(x) = Pr(\Psi(t) \leq x)$. We can obtain the distribution of $\Psi(t)$ as $G_t(x) = F(t+x) - \int_0^t (1 - F(t + x - y))m(y)dy$, $x \geq 0$, where $m(y) = \sum_{n=1}^{\infty} f_n(y)$ and $f_n(y)$ is the n-th fold convolution of distribution $f(x)$ [22]. Assume in a sample execution, the sensor's decisions determined by policy (11) are $(0, \cdots, 0, c_{t_1} = 1, 1, \cdots, c_{t_2} = 1, 0, 0, \cdots, 0, c_{t_3} = 1, 1, \cdots)$. Note that t_1 will be equal to n_1 or $n_1 + 1$, and similarly for t_2 and t_3. With $g_t(x) = G'_t(x)$, $\hat{G}_t(x) = G_t(x) - G_t(x-1)$ and $\bar{G}_t(x) = 1 - G_t(x)$, the capture probability in state f_i is $\beta_i =$

$$\begin{cases} G_{i-1}(1), & i \leq t_1; \\ \hat{G}_{t_1-1}(i - t_1 + 1)/\bar{G}_{t_1-1}(i - t_1), & t_1 + 1 \leq i \leq t_2; \\ \;[\int_{t_2-t_1+1}^{t_3-t_1} g_{t_1-1}(s)G_{i-t_1-s}(1)ds + \hat{G}_{t_1-1}(i - t_1 + 1)] & \\ /\bar{G}_{t_1-1}(t_2 - t_1 + 1), & t_2 + 1 \leq i \leq t_3; \\ \;[\int_{t_2-t_1+1}^{t_3-t_1} g_{t_1-1}(s)\frac{\hat{G}_{t_3-t_1-s}(i-t_3+1)}{\bar{G}_{t_3-t_1-s}(i-t_3)}ds + \bar{G}_{t_1-1}(t_3 - t_1 + 1) & \\ \frac{\hat{G}_{t_1-1}(i-t_1+1)}{\bar{G}_{t_1-1}(i-t_1)}]/\bar{G}_{t_1-1}(t_2 - t_1 + 1), & i \geq t_3 + 1. \end{cases}$$

■