

Optimizing Resource Constrained Distributed Collaborative Sensing

Michael Stecklein, Hasan Burhan Beytur, Gustavo de Veciana and Haris Vikalo
Department of Electrical and Computer Engineering
The University of Texas at Austin

Abstract—Collaborative sensing of spatio-temporal events/processes is at the basis of many applications including e.g., spectrum and environmental monitoring, and self-driving cars. A system leveraging spatially distributed possibly airborne sensing nodes can in principle deliver better coverage as well as possibly redundant views of the observed processes. This paper focuses on modeling, characterising and quantifying the benefits of optimal sensor activation/scanning policies in resource constrained settings, e.g., constraints tied to energy expenditures or the scanning capabilities of nodes. Under a natural model for the process being observed we show that a periodic sensor activation policy is optimal, and characterize the relative phases of such policies via an optimization problem capturing knowledge of the sensor geometry, sensor coverage sets, and spatio-temporal intensity and event durations. Numerical and simulation results for simple different sensor geometries exhibit how performance depends on the underlying processes. We also study the gap between optimal and randomized policies and how it scales with the density of sensors and resource constraints.

I. INTRODUCTION

The increasing prevalence of terrestrial and airborne nodes with the ability of both sense and communicate is playing a central role in enabling a broad set of applications such as self-driving cars, positioning, environmental, city and RF spectrum monitoring [1]–[3]. For many applications, either because of the cost of deployment and maintenance or due to application-related requirements the number of nodes, battery capacity and/or energy harvesting capabilities, and computational power of nodes may be limited. Under such limitations understanding what is possible through optimized collaborative operation of distributed sensing nodes is both a relevant and interesting problem. In this paper we consider a basic scenario where spatially distributed nodes are enabled to collaboratively sense spatio-temporal activity in an environment. We focus on the fundamental problem exploring to what degree optimized sensor activation policies can minimize the number of missed events under a resource constraint. As we shall see optimal activation policies can improve performance up to 25% with respect to say random activation policies depending on event duration distributions and energy limitations, but simple randomization may be quite effective in other scenarios.

There has been extensive interest and research towards developing effective algorithms and systems to monitor and detect spatio-temporal events and communicate such information [1]–[3]. Further to exploit the correlated sensor measurements to minimize the network traffic and energy efficiency, compressed sensing based data gathering techniques are proposed in [4]–[7]. Although the techniques proposed in these works decrease in-network traffic and improve coverage, they do not discuss how sensor nodes should be activated. However, it is shown that under certain conditions nontrivial sensor activation patterns can significantly improve performance and energy efficiency. For example in [8], it is shown that optimizing the waiting time between two samples with respect to the network statistics can significantly reduce the “age” of the available information while decreasing unnecessary network traffic. Another class of

problems considered are those associated with scalability in the context of large-scale sensor networks which might leverage wireless relaying, see e.g., [9]. In this paper the authors propose a method based on an aggregation/compression and employing mobile sinks to smooth out inhomogeneities in traffic/energy across the network. The optimization of sensor activation when nodes have energy harvesting capability is studied in [10]. In this work, the authors model the sampling problem for a *single* node as a Markov Decision Process and characterize the optimal policy. This work also exhibits various simulation results for the case with multiple sources using a heuristic activation policy. Our work specifically explores settings where sensors might have overlapping coverage sets.

The primary motivation and goal of this paper is to explore how one might optimally coordinate sensor activation so as to maximize the number of events that are observed while meeting energy constraints. To that end we propose and analyze a generic model where Poisson distributed spatio-temporal events having random duration and are collaboratively sensed by a set of nodes with overlapping coverage sets. In this setting we provide a natural characterization of the optimal sensing policies as being periodic and exhibit an optimization problem and solution which permits one to identify an optimal policy given prior knowledge of the sensor geometry, coverage sets, and spatio-temporal intensity and duration statistics of the events. The importance of such results is to provide a baseline for what the performance limits would be.

Additionally, we show a bound on the performance gap between such randomized and the optimal sensing activation policies. It follows that for dense sensing configurations with substantial coverage overlap this gap decreases in the overlap. This shows that the combination of sensor density and periodic activation with randomized phases not only leads to near optimal sampling, but also, robustness to underlying geometry and systems statistics. Finally we include some representative numerical examples exhibiting the theoretical results and insights developed in the paper.

RF spectrum monitoring has been studied for a long time and continues to gain importance [11]–[13]. Some recent works, e.g., [11], [12], propose techniques for wideband spectrum monitoring using inexpensive hardware such as USRP Radios. The framework proposed in our work can capture the joint optimization of distributed sensing nodes’ activity either under energy constraints, where nodes cannot be “on” all the time, and/or modified to capture constraints on sensing node capabilities, e.g., their sweep rate and bandwidth resolution. The results in this paper can be used to study the benefits of coordinating sensor activation to maximize coverage in frequency-time-space and improve the accuracy of a spectrum monitoring applications by coordinating sensor activation.

The paper is organized as follows. In the Section II we describe our system model. Section III provides a result on the existence of optimal collaborative sensing policies which are periodic, followed by a characterization of the performance

of the optimal periodic policies. Section IV develops results for periodic sensing policies, and in particular establishes a bound on the performance gap between random and optimal phased policies. We exhibit numerical and simulation results in Section V, and conclude the paper with Section VI.

II. SYSTEM MODEL

We consider a set of n sensors $S = \{1, \dots, n\}$. Each sensor $s \in S$ has an associated 2-D coverage region $C^s \subset \mathbb{R}^2$, and the collection of all sensors' coverage regions is denoted $\mathcal{C} = \{C^s \mid s \in S\}$. We denote the union of all coverage regions by $R = \bigcup_{s \in S} C^s$. The coverage region of sensors can model the region that the sensor can observe and the sensor modality, e.g., beam forming in a given direction.

We shall let the set $S(\mathbf{x}) \subseteq S$ denote the sensors that cover the location $\mathbf{x} \in \mathbb{R}^2$, i.e., $S(\mathbf{x}) = \{s \in S \mid \mathbf{x} \in C^s\}$. Note that the number of possible unique subsets of $S(\mathbf{x})$ depends on the geometry of the coverage sets \mathcal{C} , and there are a finite number of such disjoint subsets $n_{\mathcal{P}}$. Note also that $n_{\mathcal{P}}$ is necessarily less than the number of possible subsets of S , i.e., $n_{\mathcal{P}} < 2^n$. We denote by P_j regions which are observed by the same subset of sensors. In other words, for any two locations \mathbf{x}, \mathbf{x}' such that $S(\mathbf{x}) = S(\mathbf{x}')$, there is a j such that $\mathbf{x}, \mathbf{x}' \in P_j$. By definition of P_j , we have that $P_j \cap P_k = \emptyset$ for $j \neq k$. Thus, $\mathcal{P} = \{P_j \mid j = 1, \dots, n_{\mathcal{P}}\}$ forms a partition induced by the coverage sets \mathcal{C} of R . We let S_j be the unique subset of sensors that observe partition P_j , i.e. $S_j = S(\mathbf{x})$ for some $\mathbf{x} \in P_j$, and $n_j \triangleq |S_j|$. An example of a set of coverage sets and the associated partition is provided in Fig. 1.




Fig. 1. An example of coverages for $n = 3$ sensors. The defined quantities in this example are: $\mathcal{C} = \{C^1, C^2, C^3\}$, $n_{\mathcal{P}} = 4$, $\mathcal{P} = \{P_1, P_2, P_3, P_4\}$, $S_1 = \{1\}$, $S_2 = \{1, 3\}$, $S_3 = \{3\}$, $S_4 = \{2\}$.

Each sensor $s \in S$ has an associated sequence of (possibly random) sampling times, $T^s = (T_i^s)_i$, and each sampling period has a fixed duration of length δ . We shall assume that during any sampling interval $(T_i^s, T_i^s + \delta]$, the sensor s observes all activity within its coverage region C^s . The union of all of its sampling intervals is denoted $\mathcal{T}^s = \bigcup_i (T_i^s, T_i^s + \delta]$.

We will constrain sampling times such that any two samples cannot be closer than τ in time, i.e., $T_{i+1}^s - T_i^s \geq \tau \forall s \forall i$ almost surely. Here, the parameter τ denotes a constraint on a minimum period between samples, and is such that $\tau > \delta$. This constraint can be viewed as limiting the maximum sampling frequency or eventually tied to a constraint on the maximum power expenditure of a sensor.

Our model is quite general and can be easily modified to capture other settings such as constraints on the sweep rate and bandwidth resolution of an RF spectrum monitoring node, see e.g., [11]. We only need to define τ and δ , such that

$$\tau = \frac{f_{\text{end}} - f_{\text{start}}}{k}, \quad \delta = \frac{f_B}{k}$$

where f_{start} and f_{end} are the end frequencies within which the chirp signal varies, k denotes the sweep rate indicating how

fast the sensor sweeps the spectrum, and f_B is the bandwidth of the baseband signal sampled by the ADC of the receiver.

We model activity/events in R as a time-stationary marked point process $\Gamma = (\mathbf{E}_i)_i$ where each event $\mathbf{E}_i = (A_i, D_i, \mathbf{X}_i)$ is characterized by three components: its arrival time A_i , its duration D_i , and its location $\mathbf{X}_i \in \mathbb{R}^2$. The arrivals $(A_i)_i$ follow a stationary point process with intensity λ such that the expected number of arrivals in an interval of length ε is given by $\lambda\varepsilon$. We assume that locations $(\mathbf{X}_i)_i$ are i.i.d with a 2-D spatial probability density function $f_{\mathbf{X}}(\mathbf{x})$ on R . Similarly, the durations $(D_i)_i$ are i.i.d. with cumulative distribution function F_D . Note under this model activity/events need not be spatially homogeneous.

A policy π schedules sensor activation times subject to the aforementioned constraints. We will let $T^{s,\pi}$ denote the sequence of (possibly random) sampling times of sensor s under policy π , and similarly $\mathcal{T}^{s,\pi}$ is the union of sampling intervals of sensor s under policy π . All policies and their sampling sequences are assumed to be independent of the event process Γ . Ideally a policy should coordinate the activation of sensors in order to minimize the number of unobserved, or ‘‘missed’’, events. An event \mathbf{E}_i is missed if there is no sensor which covers the event’s location and is active between its arrival and departure, i.e., $\nexists s \in S$ s.t. $\mathbf{X}_i \in C^s$ and $(A_i, A_i + D_i] \cap \mathcal{T}^s \neq \emptyset$. To quantify the performance of a sampling policy, we define several quantities.

The set of events located in partition P_j that are missed under policy π is given by,

$$\mathcal{M}_j^\pi \triangleq \{(A_i, D_i, \mathbf{X}_i) \mid \mathbf{X}_i \in P_j \text{ and } (A_i, A_i + D_i] \cap \mathcal{T}^{s,\pi} = \emptyset \forall s \in S_j\}.$$

We denote the set of missed events in partition P_j which are initiated in the time interval $(0, t]$ by,

$$\mathcal{M}_j^\pi(0, t] \triangleq \{(A_i, \mathbf{X}_i, D_i) \in \mathcal{M}_j^\pi \mid A_i \in (0, t]\}.$$

Then, the number of missed events located in partition P_j and initiated in an interval $(0, t]$ is given by,

$$M_j^\pi(0, t] \triangleq |\mathcal{M}_j^\pi(0, t]|.$$

Counting such misses across all partitions, we have,

$$M^\pi(0, t] \triangleq \sum_{j=1}^{n_{\mathcal{P}}} M_j^\pi(0, t].$$

Let Π denote the set of all stationary sensor activation policies that satisfy the sampling constraints. An initial objective might be to determine a stationary policy $\pi^* \in \Pi$ that minimizes the expected number of misses over any interval, i.e.,

$$\pi^* \in \underset{\pi \in \Pi}{\text{argmin}} \mathbb{E}[M^\pi(t, t + d)] \quad \forall t \in \mathbb{R}, d > 0.$$

III. CHARACTERIZING OPTIMAL POLICIES

In the sequel we refer to ‘‘ τ -periodic policies’’ as the subset of stationary sampling policies, $\Pi^\tau \subset \Pi$, which for each sensor is periodic with the minimal period τ allowed by the sampling constraint.

Theorem 1 (Existence of an Optimal Periodic Policy): For a given stationary event process Γ there exists a stationary optimal policy $\pi^* \in \Pi$ which is τ -periodic, i.e., $\pi^* \in \Pi^\tau$, and minimizes the expected number of misses on any interval.

This theorem can be shown using a coupling argument, and showing that there exists a τ -periodic stationary policy performing as well as or better than any stationary policy. The detailed proof is given in Appendix A of [14].

Theorem 1 asserts the existence of a τ -periodic policy which is optimal but does not characterize such optimal policies. To that end, we start by considering deterministic τ -periodic policies. These policies are parameterized by a vector of phases $\phi = (\phi_s \mid 0 \leq \phi_s \leq \tau, s \in S)$ within $[0, \tau]$ representing the relative shift of the τ -periodic sampling times of each sensor—we denote the associated sampling policy by π^ϕ and note that under policy π^ϕ sensor s will sample deterministically at times $t^{s, \pi^\phi} = (\tau i + \phi_s)_i$.

Each phase vector ϕ has an associated ordering of the n sensors based on their sampling times in $[0, \tau]$ and conversely each ordering specifies a set of phase vectors ϕ consistent with that ordering. We introduce notation allowing us to work with such orderings. We define a class Ω of functions $\omega : \{1, \dots, n\} \rightarrow S$ which capture an ordering on the set of sensors through a one-to-one mapping between indices and sensors. More explicitly, $\omega(i)$ denotes the i^{th} sensor in the ordering that ω imposes on the sensors in S . Thus, the ordering that ω induces over S is $(\omega(1), \dots, \omega(n))$ and a phase vector ϕ consistent with that ordering will satisfy $\phi_{\omega(1)} \leq \dots \leq \phi_{\omega(n)}$.

Furthermore for a subset $S_j \subset S$ of cardinality n_j we define $\omega^{(j)} : \{1, \dots, n_j\} \rightarrow S_j$ as the ordering that ω induces on the sensors in S_j . Namely, $\omega^{(j)}(i)$ is the i^{th} sensor in the ordering of S_j under ω , and for any sensors $s, r \in S_j$, $\omega^{-1}(s) \leq \omega^{-1}(r)$ implies $(\omega^{(j)})^{-1}(s) \leq (\omega^{(j)})^{-1}(r)$.

Now, consider a vector of phases ϕ and a corresponding ordering ω and policy π^ϕ induced by ϕ . We evaluate the performance of π^ϕ over $(0, \tau]$ and rewrite the objective function:

$$\mathbb{E}[M^{\pi^\phi}(0, \tau)] = \mathbb{E}\left[\sum_{j=1}^{n_p} M_j^{\pi^\phi}(0, \tau)\right] = \sum_{j=1}^{n_p} \mathbb{E}[M_j^{\pi^\phi}(0, \tau)].$$

Dividing the interval $(0, \tau]$ into subintervals of length ε yields,

$$\mathbb{E}[M^{\pi^\phi}(0, \tau)] = \sum_{j=1}^{n_p} \sum_{k=0}^{\tau/\varepsilon-1} \mathbb{E}[M_j^{\pi^\phi}(k\varepsilon, (k+1)\varepsilon)].$$

For infinitesimally small ε , the expected number of misses in partition j over an interval of width ε is equal to the expected number of event arrivals in partition j times the probability the sampling policy misses an event.

Under our activity model, the expected number of event arrivals in partition j in an interval of width ε is $\lambda\varepsilon w_j$, where w_j is the probability of an event arrived in partition j , i.e., $w_j = \mathbb{P}(\mathbf{X}_i \in P_j) = \iint_{P_j} f_X(\mathbf{x}) d\mathbf{x}$.

In order to miss an event i , its arrival time A_i must be after the active interval of the previous sample of its partition, and the event's departure time $A_i + D_i$ must be before the next sample of its partition. At a time t in partition j , we denote the time of the previous sample by $s_{\text{prev}}^{j, \phi}(t)$ and that of the next sample by $s_{\text{next}}^{j, \phi}(t)$. Then, the probability of missing an event that arrives at time $k\varepsilon$ is $F_D(s_{\text{next}}^{j, \phi}(k\varepsilon) - k\varepsilon) \mathbb{1}(k\varepsilon > s_{\text{prev}}^{j, \phi}(k\varepsilon) + \delta)$. Fig. 2 exhibits an example to help visualize this for a partition having two sampling times $t_0^{s_1}$ and $t_0^{s_2}$ in the interval $(0, \tau]$ associated with sensors s_1 and s_2 respectively. Note that the phase of the first sensor is less than the phase of the second sensor, such that $\phi_{s_1} < \phi_{s_2}$. In the figure, the event $\mathbf{E} = (A, D, \mathbf{X})$ arrives in an interval of width ε . Although it departs before the next sample time, $A + D < s_{\text{next}}^{j, \phi}(A) = t_0^{s_2}$ it arrived before the end of the previous sample interval, $A < s_{\text{prev}}^{j, \phi}(A) + \delta = t_0^{s_1} + \delta$ so it will not be missed.




Fig. 2. An example of an arrival in an interval of ε width being observed by the previous sample.

Substituting this result into the equation for $\mathbb{E}[M^{\pi^\phi}(0, \tau)]$ and letting ε go to zero, we have

$$\begin{aligned} &= \sum_{j=1}^{n_p} \sum_{k=0}^{\tau/\varepsilon-1} \lambda\varepsilon w_j F_D(s_{\text{next}}^{j, \phi}(k\varepsilon) - k\varepsilon) \mathbb{1}(k\varepsilon > s_{\text{prev}}^{j, \phi}(k\varepsilon) + \delta). \\ &= \sum_{j=1}^{n_p} \lambda w_j \int_0^\tau F_D(s_{\text{next}}^{j, \phi}(u) - u) \mathbb{1}(u > s_{\text{prev}}^{j, \phi}(u) + \delta) du. \end{aligned}$$

Thus for partition j , we define the intervals between its associated sampling times, or “gaps”, as the sequence $(g_k^{j, \phi} : k = 1, \dots, n_j)$ where,

$$g_k^{j, \phi} = \begin{cases} \phi_{\omega^{(j)}(k+1)} - \phi_{\omega^{(j)}(k)}, & \text{for } k < n_j, \\ \tau - \phi_{\omega^{(j)}(n_j)} + \phi_{\omega^{(j)}(1)}, & \text{for } k = n_j. \end{cases}$$

Note that $\sum_{k=1}^{n_j} g_k^{j, \phi} = \tau$. Then, we can equivalently write,

$$\mathbb{E}[M^{\pi^\phi}(0, \tau)] = \sum_{j=1}^{n_p} \lambda w_j \sum_{k=1}^{n_j} \int_\delta^{g_k^{j, \phi}} F_D(g_k^{j, \phi} - u) du.$$

To understand this equivalence, we use Fig. 2 as an example. The interval $(0, \tau]$ can be decomposed into the intervals $t_0^{s_2} - t_0^{s_1}$ and $\tau - t_0^{s_2} + t_0^{s_1}$. Between samples $t_0^{s_1}$ and $t_0^{s_2}$, the probability of a miss is zero for the first δ , and then it is $F_D(g_1^{j, \phi} - u)$, (i.e. the probability that the event's duration is longer than the remainder of the gap), for $u \in (t_0^{s_1} + \delta, t_0^{s_2}]$, where $g_1^{j, \phi} = t_0^{s_2} - t_0^{s_1}$ in this example. Notice that since F_D is a C.D.F., it is a non-decreasing function, so the integral over F_D is a convex function of ϕ . Thus, $\mathbb{E}[M^{\pi^\phi}(0, \tau)]$ can be written as a sum of convex functions, and it is convex.

Next, let us consider a particular fixed ordering ω . A phase vector ϕ satisfies the ordering ω if $\phi_{\omega(1)} \leq \dots \leq \phi_{\omega(n)}$. Suppose we wish to find the phase vector that minimizes $\mathbb{E}[M^{\pi^\phi}(0, \tau)]$ amongst all phase vectors that satisfy a given ordering ω . This is achieved by solving the optimization problem:

$$\begin{aligned} \min_{\phi} \quad & \sum_{j=1}^{n_p} \lambda w_j \sum_{k=1}^{n_j} \int_\delta^{g_k^{j, \phi}} F_D(g_k^{j, \phi} - u) du \\ \text{s.t.} \quad & 0 = \phi_{\omega(1)} \leq \phi_{\omega(2)} \leq \dots \leq \phi_{\omega(n)} \leq \tau, \end{aligned} \quad (1)$$

$$\text{where } g_k^{j, \phi} = \begin{cases} \phi_{\omega^{(j)}(k+1)} - \phi_{\omega^{(j)}(k)}, & \text{for } k < n_j, \\ \tau - \phi_{\omega^{(j)}(n_j)} + \phi_{\omega^{(j)}(1)}, & \text{for } k = n_j. \end{cases}$$

Since the objective is convex, and the constraint on ϕ is convex, this optimization problem is convex.

Observe that $\phi_{\omega(1)}$ is pinned to zero in order to prevent shifted versions of ϕ . Shifted policies are redundant since they result in circular shifts over the interval $(0, \tau]$, and hence they don't change the value of the objective function. Finally, we propose a construction of an optimal τ -periodic policy. We construct this stationary minimally periodic policy, $\pi^{(s.m.p.)}$, as follows. We denote by $\Omega^{(1)} \subset \Omega$ the subset of orderings of sensors for which the first sensor is sensor 1, i.e. $\Omega^{(1)} = \{\omega \in \Omega \mid \omega(1) = 1\}$. For each ordering $\omega \in \Omega^{(1)}$, evaluate the convex optimization problem for that ω . Select the ordering ω^* that achieves the smallest minimum, and let its argmin be ϕ^* . (Note

that these may be non-unique). Shift the deterministic policy parameterized by ϕ^* , $\pi^{(\phi^*)}$, by a uniformly random value $\Phi \in (0, \tau]$. This randomly shifted policy is $\pi^{(s.m.p.)}$.

Theorem 2 (A stationary optimal periodic policy): The $\pi^{(s.m.p.)}$ policy is an optimal stationary periodic policy.

Proof: By Theorem 1, there exists an optimal τ -periodic policy. Therefore, it is sufficient to show that $\pi^{(s.m.p.)}$ is optimal amongst this subset of policies. By stationarity of policy $\pi^{(s.m.p.)}$, it achieves the same performance over the interval $(0, \tau]$ as it does over any interval. Therefore, it is further sufficient to show that the deterministic policy $\pi^{(\phi^*)}$ used in the construction of $\pi^{(s.m.p.)}$ is optimal over $(0, \tau]$ amongst all possible parameters ϕ . Since the construction of $\pi^{(s.m.p.)}$ finds ϕ^* by minimizing the objective over all possible non-redundant parameters ϕ , (searching over $\Omega^{(1)}$ instead of Ω removes only redundant parameters), we can conclude that $\pi^{(s.m.p.)}$ is optimal. ■

IV. CHARACTERIZING RANDOMIZED SENSING POLICIES

In certain settings, computation of the optimal policy may be intractable, e.g., when the number of sensors is large, or computational resources are small, or the system's properties change frequently. In such scenarios, a simpler policy is desirable. This motivates the use of a randomized policy, in which the sensors' sampling phases are random and independent of each other. Below, we formalize this policy and bound its performance loss in both a general and a dense setting.

We define the random policy, $\pi^{(rand.)}$, as a sampling policy whose random sampling times for each sensor s are given by:

$$T_i^{s, \pi^{(rand.)}} = \Phi_s + i\tau \quad \forall s \in S,$$

where $\Phi_s \forall s \in S$ are independent and identically distributed random variables uniform over $(0, \tau]$. Note that each sensor's sampling times are independent of all other sensors, and each sensor samples every τ .

The best performing possible realization of the set of random variables Φ_s will perfectly align the samples of all sensors with those of an optimal policy $\pi^{(s.m.p.)}$. This is possible since the possible realizations of $\pi^{(rand.)}$ span all τ -periodic policies.

Conversely, the worst performing realization of $\pi^{(rand.)}$ is when all samples are synchronized, i.e., when $\Phi_s = \Phi_r$ almost surely $\forall s, r \in S$, since regions with overlapping sensor coverage are the sampled simultaneously. We can bound the performance loss of the random policy $\pi^{(rand.)}$ by bounding the loss of the worst case realization of the random policy.

For ease of notation, we introduce a few new definitions. We define the worst case synchronized realization of the random policy as $\pi^{(sync.)}$, which has sampling times given by:

$$T_i^{s, \pi^{(sync.)}} = \Phi_1 + i\tau \quad \forall s \in S,$$

where Φ_1 is uniformly random over $(0, \tau]$. We also introduce the following function:

$$\mathcal{F}_D^\delta(x) = \begin{cases} \int_\delta^x F_D(x-u)du, & \text{for } x > \delta, \\ 0, & \text{for } x \leq \delta. \end{cases}$$

Note that our objective function corresponding to the mean number of missed events can be written as a weighted sum of these functions. Lastly, we define the loss of a policy π with respect to the optimal policy $\pi^{(s.m.p.)}$ as:

$$\mathcal{L}(\pi) = \mathbb{E}[M^\pi(0, \tau)] - \mathbb{E}[M^{\pi^{(s.m.p.)}}(0, \tau)].$$

Theorem 3 (Performance loss bounds for random policy): The loss of the random policy $\pi^{(rand.)}$ with respect to an

optimal policy $\pi^{(s.m.p.)}$ is bounded by that achieved by synchronizing sensors phases $\pi^{(sync.)}$, which itself can be bounded as follows:

$$\mathcal{L}(\pi^{(rand.)}) \leq \mathcal{L}(\pi^{(sync.)}) \leq \lambda \mathcal{F}_D^\delta(\tau) - \lambda \mathbb{E}[N] \mathcal{F}_D^\delta\left(\frac{\tau}{\mathbb{E}[N]}\right),$$

where N is an RV such that $\mathbb{P}(N = n_j) = w_j$, for $j = 1, \dots, n_p$, where w_j correspond to the probabilities of a given event being located in each partition. Hence, $\mathbb{E}[N]$ is the weighted average of the number of sensors observing each set of the partition.

For the proof, we refer the reader to Appendix B in [14].

Furthermore, we characterize a setting for which we can find a tighter bound on the loss. We call this setting "d-dense" if all locations $\mathbf{x} \in R$ are observed by at least $d \geq 1$ sensors. This also implies that for all j we have $n_j \geq d$.

In d -dense settings for which $d \gg 1$, realizations of the ordered phases of the sensor samples of the random policy converge towards an even spacing between samples for any given partition. A visualization of such an even spacing is shown in Figure 3. For a partition j with $n_j = 4$, even spacing within a period τ would result in a spacing of $\frac{\tau}{n_j+1} = \frac{\tau}{5}$ between each sample and the edges of the period. Note that across the boundaries of periods, this results in a gap of $\frac{2\tau}{n_j+1}$ between samples.




Fig. 3. An example of samples in a partition j over one period τ that are evenly spaced with respect to each other and to the borders of the period.

We can see that the random policy converges to even spacing as d goes to infinity if we consider the order statistics of $(\Phi_s \forall s \in S_j)$ for a particular partition j . Let $\Phi_{(1)}^j, \dots, \Phi_{(n_j)}^j$ denote their order statistics, and let $G_{(k)}^j = \Phi_{(k)}^j - \Phi_{(k-1)}^j$. Without loss of generality, let $\tau = 1$. Since $\Phi_s \sim \text{Uniform}(0, \tau = 1]$, then $\Phi_{(k)}^j \sim \text{Beta}(k, n_j + 1 - k)$ and $G_{(k)}^j \sim \text{Beta}(1, n_j)$. Then, we have that the expected size of the gaps between sensors' samples is $\mathbb{E}[G_{(k)}^j] = \frac{1}{n_j+1}$. The variance of $G_{(k)}^j$ is $O(\frac{1}{n_j^2})$, which means that as d becomes large, n_j becomes large, so the distribution of the random gap lengths concentrates at its mean $\frac{1}{n_j+1}$. Thus, the random policy converges to even spacing.

Using this intuition, for d -dense settings with a large finite d , we can derive a tighter bound on the loss of performance of the random policy $\pi^{(rand.)}$, $\mathcal{L}(\pi^{(rand.)})$, by considering the random gaps $G_{(k)}^j$ which are concentrating on even spacings.

Theorem 4 (Bounds on a random policy for dense settings): For a d -dense setting, the loss of the random policy $\pi^{(rand.)}$ with respect to the optimal policy $\pi^{(s.m.p.)}$ is bounded as follows:

$$\mathcal{L}(\pi^{(rand.)}) \leq \frac{\lambda \tau C}{d} = O\left(\frac{1}{d}\right),$$

$$C = f_D\left(\frac{1}{d+1}\right) + 2 + \sum_{p=3}^{\infty} \frac{p\tau^{p-2}}{p!(p-2)} \mathcal{F}_D^\delta\left(\frac{1}{d+1}\right).$$

The proof has been relegated to Appendix C in [14].

It is worth considering the case where $\frac{1}{d+1} < \delta$. As d becomes larger, there is an increasing probability that the random policy will achieve a perfect coverage, i.e., the union of all sampling intervals be $(0, \tau]$. If $\frac{1}{d+1} < \delta$, the constant in the bound in Theorem 4 will reduce to $C = f_D(\frac{1}{d+1})$, since $\mathcal{F}_D^\delta(\frac{1}{d+1}) =$

0. This tighter bound reflects the probability that the random policy will achieve perfect coverage.

V. NUMERICAL RESULTS

To get a sense of the quantitative performance of optimized sensor activation policies we considered a variety of simple scenarios starting with a single sensor, with two sensors with overlapping coverage, and subsequently with 5 sensors with increasingly overlapping coverage. To examine the effect of a power constraint, we consider a simple model where each sensor must operate under a power constraint ρ given by $\delta/\tau = \rho$, where $1/\tau$ corresponds to the sampling frequency and δ the sensing, “on” time. This corresponds to a simple proxy for a constraint on energy expenditures on sensing devices.

The idea here is to build some intuition regarding the impact of various system characteristics, e.g., the event durations, as well as the impact that increasing overlap amongst sensors’ coverage sets and increasing power constraint ρ will have on the performance of an optimal sensing versus a random policy.

A. Single Sensor

We began by evaluating the simplest non-trivial setting: a single sensor. For a variety of event durations F_D with a mean duration of 1, we computed the mean miss rate of the optimal policy, i.e. $\frac{1}{\tau} \mathbb{E}[M^{\pi^{(s,m,p)}}(0, \tau)]$. We plotted the miss rates against the sampling frequency $\frac{1}{\tau}$, when δ is set to be zero. The resulting plots for uniform, exponential, constant, and Pareto distributions are shown in Figure 4.




Fig. 4. Mean miss rate versus sampling frequency $\frac{1}{\tau}$ for uniform, exponential, constant, Pareto and Weibull distributions, all with a mean of 1.

We can see that for an increase in sampling frequency, the returns on performance are diminishing. Also, the rate of diminishing returns are very dependent on the event duration distribution F_D . This implies that in order to properly balance power consumption with performance, it is necessary to have prior knowledge of the event duration distribution. For example, doubling the sampling frequency from 2 to 4 for the constant distribution in Figure 4 provides no benefit, while doing so for the Pareto distribution decreases the mean miss rate significantly.

Figure 4 also provides insight into the maximum possible benefit that can be achieved by leveraging overlapping sensors. For example, let us consider a setting where sensors sample at a frequency of 2 and events arrive with exponential interarrivals. Then, we expect a mean miss rate of about .3 for any regions covered by only a single sensor. Any regions which are overlapped by multiple sensors can also perform no worse than .3, but no better than 0, so the marginal improvement from overlapping sensors is bounded.

B. Two Sensors

We next considered settings with two sensors and characterized the benefits of coordinating their sampling times.

We began by comparing the mean miss rate of the optimal policy against the random policy for varying sampling frequencies $1/\tau$. Figure 5a shows these plots for a setting where event durations are uniformly distributed. They assume that half of each sensor’s coverage overlaps with the other’s. Similar to the single sensor case, the observations are assumed to be instantaneous, i.e., $\delta = 0$. These plots show the mean miss rates for: the optimal policy, the average realization of the random policy, the worst possible case of the random policy, (i.e. synchronized), and the lower quartile of the random policy.




Fig. 5. Mean miss rate versus sampling frequency $\frac{1}{\tau}$ (a, left), power constraint $\rho = \delta/\tau$ (b, middle) and percent overlap of the sensors’ coverages (c, right) of the optimal, random, worst case random, and lower quartile of random policies, for a uniform event duration.

The plots show that the marginal benefit of the optimal policy over the random policy is again dependent on the event duration distribution as well as the sampling frequency. In both plots, we see the largest increase in performance in the low and middle sampling frequencies, and we see negligible improvement in very low and high sampling frequencies.

The “Optimal” and “Worst Case” lines in the plots provide empirical bounds on the performance on the random policy, since they represent the best and worst performing possible realizations. The optimal policy takes full advantage of the overlap of the two sensors’ coverage regions, while the worst case synchronized realization takes no advantage of it.

We can observe the influence of coverage overlap on the performance bounds in Figure 5c. We vary the fraction of the coverage regions overlap between the two sensors from none to fully overlapping, and we plot the mean miss rates. As expected, the optimal policy performs much better than the worst case synchronized policy for greater overlap. Interestingly, the random policy tends to perform slightly closer to the optimal policy than to the worst case for greater overlaps.




Fig. 6. Relative performance versus power constraint $\rho = \delta/\tau$ of random policy over optimal policy, for different event duration.

By letting the observation duration vary according to the relation $\delta = \tau\rho$, we observe the impact of the power constrained sensor nodes in Figure 5b – for the case where event durations are uniformly distributed. In this setting, the sampling

frequency is $1/\tau = 2$, and the sensors' coverage regions overlap by about 50%. As the power constraint increases, i.e., the duration of each observation increases, the mean miss rates for all schemes decrease. In Figure 6, the relative performance of random policy with respect to the optimal case under various event duration distributions is shown. These plots exhibit the mean miss rate ratio of a randomized sensing policy versus the optimal policy. We see that the marginal benefit of using the optimal policy changes depending on the power constraint. If the power is not constrained, random policy can perform as well as the optimal policy. Otherwise, the optimal policy can significantly improve the performance depending on the event duration distribution. Although employing the optimal policy requires some computational effort, it enables us to maximize performance while possibly expending reducing sensors' energy expenditures. Additionally, in the setting with two sensors, the distribution for the duration of events has a limited impact on the relative performance of optimal policy, except for the case where durations are uniformly distributed.

Overall, the simulations reveal that although there are certain settings for which the optimal policy does perform noticeably better than the random policy, for the majority of the settings the marginal benefit of the optimal policy over the random policy is small or even negligible. Furthermore, the random policy tends to perform much better than the worst case when there is a lot of overlap in coverage, e.g., in dense settings. This further justifies the use of the simpler random policy instead of the computationally expensive optimal policy in some settings, such as those that are dense.

C. Set of sensors with increasing coverage sets

The units are arbitrary, for simplicity assume that distance is in "miles" and time is in "seconds". 5 of the coverage regions are evenly spaced on a circle of radius 1.5 miles, and the 6th coverage region is at the center of the circle. The radii of the coverage regions takes values between 0.5 and 6.0 miles.




Fig. 7. Partitions and coverage redundancy with increasing coverage radius.




Fig. 8. Mean miss rates for optimal and random policies with increasing coverage radius.

The minimum sampling period τ is 1 second. The event durations were independent and uniformly distributed on $[0, \tau]$. The sampling duration δ was set to 0 (i.e. instantaneous samples). The events arrival process was Poisson with an overall rate of $\lambda = 1$ sec, and homogeneous in the coverage area. We examined the performance of the optimal sensing strategy versus the randomized one, as we increased the radius of disc shaped coverage areas associated with each sensor. The upper row of Figure 7 exhibits the partitions induced as one increases the coverage radius of the sensors. The lower row of the same, shows the redundancy in coverage at various locations as one increases the coverage radius. The lighter the shade of gray the higher the redundancy at that location. Figure 8 exhibits

the performance in terms of the miss rate of the optimal and random policies. We have also added the performance of the lower quartile for the random policy to show that a randomized sensing policy could lead to a substantially worse performance if the realization of the phases were poorly aligned.

VI. CONCLUSIONS

In this paper we considered a fairly generic framework for a distributed sensing platform wherein sensors are deployed to observe stochastic events which have limited durations. With some care one can develop a good understanding of how to optimize the sensing policy subject to a simple proxy constraint on sensor's activation/energy.

- Optimizing the sensing policy provides benefits, which may be quite dramatic depending on the power constraint, for settings with very limited sensor resources and with overlapping coverages.
- Randomizing sensor phases of τ -periodic sensor activation may provide adequate robust performance, particularly in dynamic settings where optimizing may be costly.
- Context dependent coordination of joint sensor activation to provide "slices" of network activity that deliver features (e.g., for localization of nodes) may also benefit from coordinated optimal sensing policies.

VII. ACKNOWLEDGEMENT

This research is supported by the US Army Futures Command. The contents of this document represent the views of the authors and are not necessarily the official views of, or endorsed by, the US Government, Department of Defense, Department of the Army or US Army Futures Command.

The work of G. De Veciana and H. Vikalo was also partially supported NSF Award ECC-1809327.

REFERENCES

- [1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer Networks*, 2002.
- [2] P. Zhang, I. Nevat, G. W. Peters, G. Xiao, and H.-P. Tan, "Event Detection in Wireless Sensor Networks in Random Spatial Sensors Deployments," *IEEE Transactions on Signal Processing*, Nov. 2015.
- [3] H. N. Pham, Y. Zhang, P. E. Engelstad, T. Skeie, and F. Eliassen, "Optimal cooperative spectrum sensing in cognitive sensor networks," in *Wireless Communications and Mobile Computing*, Germany, 2009.
- [4] C. Luo, F. Wu, J. Sun, and C. W. Chen, "Compressive data gathering for large-scale wireless sensor networks," in *Mobile Computing and Networking*, New York, NY, USA, 2009.
- [5] V. Singh, V. K. Singh, and M. Kumar, "In-network data processing based on compressed sensing in WSN: A survey," *Wireless Personal Communications*, 2017.
- [6] V. K. Singh, M. Kumar, and S. Verma, "Node Scheduling and Compressed Sampling for Event Reporting in WSNs," *IEEE Transactions on Network Science and Engineering*, Jul. 2019.
- [7] Z. Charbiwala, Y. Kim, S. Zahedi, J. Friedman, and M. B. Srivastava, "Energy efficient sampling for event detection in wireless sensor networks," in *Low Power Electronics and Design*, San Francisco, USA, 2009.
- [8] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksall, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Transactions on Information Theory*, 2017.
- [9] C.-S. Choi, F. Baccelli, and G. de Veciana, "Densification leveraging mobility: An IoT architecture based on mesh networking and vehicles," in *Mobile Ad Hoc Networking and Computing*, New York, USA, 2018.
- [10] Z. Ren, P. Cheng, J. Chen, D. K. Yau, and Y. Sun, "Dynamic Activation Policies for Event Capture with Rechargeable Sensors," in *Distributed Computing Systems*, Jun. 2012.
- [11] Y. Guddeti, R. Subbaraman, M. Khazraee, A. Schulman, and D. Bharadia, "SweepSense: Sensing 5 GHz in 5 milliseconds with low-cost radios," in *Networked Systems Design and Implementation*, Boston, MA, Feb. 2019.
- [12] L. Shi, P. Bahl, and D. Katabi, "Beyond sensing: Multi-GHz realtime spectrum analytics," in *Symposium on Networked Systems Design and Implementation*, Oakland, CA, May 2015.
- [13] A. A. Cheema and S. Salous, "Digital FMCW for ultrawideband spectrum sensing," *Radio Science*, 2016.
- [14] M. Stecklein, H. Beytur, G. de Veciana, and H. Vikalo, "Optimizing resource constrained distributed collaborative sensing (extended)." [Online]. Available: <http://users.ece.utexas.edu/~gustavo/papers/SBD21.pdf>