Joint Sensor Duty Cycle Scheduling with Coverage Guarantee

Kin Sum Liu*   Jie Gao *   Shan Lin †   Hua Huang †   Brent Schiller †

ABSTRACT

Using optical sensors for indoor monitoring has been widely adopted in many smart building applications. An important design problem in this space is to explore the tradeoff between energy consumption and coverage quality. While it is important that the sensors achieve full coverage (i.e., every interesting target point can be monitored by at least one sensors), it is often a waste of energy to keep sensors on all the time as events are typically stochastic and rare and most of the time the sensors are on idle monitoring. In this paper we design efficient sensor duty cycles to ensure that any target point of interest is still covered sufficiently frequently while only a subset of sensors are kept on at any time slot. We denote by the maximum dark length for each target point p as the maximum duration in which p is covered at least once. We formulate two optimization problems: the min max dark length scheduling and the min average dark length scheduling. For both versions we provide efficient, practical algorithms with provable approximation guarantee. The two algorithms have been tested on two real testbed scenarios to evaluate its efficiency and coverage quality.

CCS Concepts

•**Networks** → *Sensor networks;* •**Theory of computation** → *Scheduling algorithms;*

Keywords

Sensor coverage, duty cycle scheduling, energy efficiency

1. INTRODUCTION

Optical networks have been widely deployed in modern buildings for various indoor monitoring applications, such

*Computer Science Department, Stony Brook University. Email: {kiliu, jgao}@cs.stonybrook.edu
†Electrical and Computer Engineering Department, Stony Brook University. Email: {shan.x.lin, hua.huang, brent.schiller}@stonybrook.edu

as motion detection using passive infrared sensors [1], safety surveillance using camera networks [2], window system control using visible light sensors [3], and indoor localization using infrared systems [4]. By the nature of the directionality of signals used in these applications, it is often a nontrivial problem to deploy the sensors to obtain full coverage of the domain. Once the sensor locations are determined, it is often assumed that all sensors are kept on all the time. Keeping all sensors on may consume a significant amount of energy especially when the events happen in a stochastic manner. Most of sensing energy is wasted on idle monitoring, which is inefficient and sometimes unnecessary. There is a need to reduce energy consumption on idle sensing as long as the application requirements are met.

For most of the optical sensing applications such as surveillance and target localization, full coverage is important so that every location of the area can possibly be monitored. Fortunately, these applications may not require every location in the domain to be covered *all the time*, because 1) the physical events that appear in the domain usually last for a certain period of time. For example, an intruder usually passes through a camera's view in a few seconds; the shadow cast by human movement lasts for hundreds of milliseconds; 2) these events can be collaboratively detected by a network of sensors. One location can be covered by multiple sensors. Although an individual sensor can miss an event, another sensor may still catch it. Therefore, our goal is to design efficient optical sensing schedules to turn sensors on and off, achieving a good tradeoff between detection delay and sensing energy consumption.

In this paper we formulate the sensor duty cycle scheduling problem in a general manner. We assume a domain $\mathcal{D}$ and a set of n sensors that collectively cover it. We do not make any special assumptions on sensing coverage range. Instead each node g is associated with a set $P(g)$ representing the region of coverage. This allows the results in this paper to be applied in the most general setting, incorporating a variety of sensing models (omni-directional models or non-isotropic ones such as line-of-sight constraints). We assume that time is slotted and at each time slot, k out of n sensors are turned on. k is a parameter that can be used to tune the expected energy consumption level. The problem is to decide which set of sensors to turn on in each slot and in what order such that each target point is covered sufficiently frequently. The set of target points can be either discrete or continuous.

We discussed two different evaluation metrics on the quality of coverage. The first version, the min max dark length

scheduling, minimizes the value C such that in every consecutive C slots every target is covered at least once. This provides worst case guarantee on all target points. We show that this problem in general is NP-hard, but if the set of sensor nodes $\mathcal{V}$ is a minimal set – meaning every node covers a unique target that is not covered by any other sensors (i.e., every sensor is needed for full coverage), then round robin on this set $\mathcal{V}$ is optimal.

In the second evaluation metric, we try to minimize the weighted average of the max dark length for each target in the domain $\mathcal{D}$. This has the freedom of being not tailored by the worst case scenario but rather measures average coverage quality. The weights can also be used to provide differentiation on the importance of different targets. For this problem we provide, again, an easy and intuitive algorithm with an $O(1)$-approximation when $k = \Omega(\log n / \log \log n)$ and $O(\log n / \log \log n)$-approximation otherwise. The algorithm is basically round robin on a random permutation in which there are $\tau(g)$ copies of each sensor g. The parameter $\tau(g)$ for each $g \in \mathcal{V}$ can be solved by convex optimization.

In addition to providing theoretical guarantees of the algorithms, we have also conducted both simulations and testbed experiments to evaluate our algorithms. Under different deployment topology and schedules, we can achieve a significant amount of energy saving on sensing with a modest decay of coverage quality.

The contributions of this work are listed as follows,

- To the best of our knowledge, this is the first work to investigate general sensor scheduling of optical sensor networks with theoretical guarantees. The algorithms are sufficiently simple and can be implemented in a real system.

- With extensive simulations and system experiments on an infrared localization system and a camera network, our algorithms can reduce energy consumption to about 1/3 with a maximum detection delay of 5 seconds.

In the rest of the paper we first survey related work on sensor duty cycle scheduling. Then we present our theoretical analysis and algorithm design under two optimization metrics. We report the simulation results and testbed evaluation last.

2. RELATED WORK

There are three major lines of research relevant to this work: scheduling algorithms in sensor networks, scheduling algorithms in camera networks, and coverage algorithms.

Scheduling Algorithms in Sensor Networks. Scheduling algorithms are studied in a wide range of computing systems, such as wireless networks [5, 6, 7, 8, 9, 10, 11, 12, 13], sensor systems [14, 15, 16, 17, 18, 19, 20, 21, 20, 22, 23, 24, 25, 26], and energy systems [27, 28]. Among these algorithms, duty cycle based scheduling algorithms have been widely adopted in communication [29, 30, 31] and energy management [32, 33] in low power systems. Various duty cycled sensing algorithms have also been developed with different detection and energy saving objectives. In [34], authors designed algorithms to adjust a node's duty-cycle according to the current amount of residual energy. In [35], a probabilistic scheduling of the duty cycles is designed to

ensure event detection probabilities. These works usually assume that a sensor is switched on-and-off once every certain period of time, which is easy to implement. What has been missing in such work is the lack of coordination among different sensors. In particular, when multiple sensors have significantly overlapping sensing ranges it would be more efficient to consider a joint optimization of multiple sensors at the same time, to further reduce energy consumption while maintaining the same level of coverage quality. This is also the main aspect that our work differs. The challenge of our problem is to consider the collaborative sensing on the target areas at a network scale with monitoring performance guarantees and reduced energy consumption.

Scheduling Algorithms in Camera Networks. There are a few papers that discuss specifically scheduling and control in optical sensor or camera networks. Authors of [36, 37] presented collaborative sensing algorithms using the field of view sensing model. In [38, 39, 40, 41], authors designed algorithms to determine the locations of cameras given different deployment requirements. In [42], authors formulate an optimization problem that aims to maximize the duration of coverage for the most critical part of the monitored region. Our work falls in this category as well, although we do not limit our algorithm to be just for camera scheduling. We allow a general definition of sensing ranges and in addition provide worst case guarantees for coverage quality. This complements and enriches existing work in this domain.

Coverage Algorithms. There are a large body of research literature on coverage algorithms in wireless sensor networks. Some of these works consider the sensor selection and scheduling problems to meet specific coverage goals [43]. In [44], an analytical framework for object detection in sensor networks is described, which allows mathematical analysis on average-case object detection quality in random and synchronized sensing schedules. Authors of [45] investigated the problem of sensor selection for minimizing error covariance using mobile sensors. In [46], authors presented a centralized and a distributed algorithm to select the minimal number of nodes to monitor p-percent of the monitored area. These works usually assume circular sensing range, which is not suitable for optical sensor network or camera networks discussed in this paper.

3. SCHEDULING PROBLEM STATEMENT

We define the scheduling problem in a generic manner. Given a domain $\mathcal{D}$ to be monitored and a set of n sensor nodes $\mathcal{V}$, we define by $P(g)$ the *region of coverage* by a sensor node $g \in \mathcal{V}$. Without loss of generality we can assume that the union of coverage ranges by all nodes in $\mathcal{V}$ is $\mathcal{D}$. $\mathcal{D}$ can be either a continuous domain (for example when the entire room needs to be monitored) or a discrete set of target locations of interest. We would like to design an efficient duty-cycle schedule of the sensor nodes such that $\mathcal{D}$ is still well covered and the total amount of time that the sensors stay on/active is reduced. This problem explores the trade-off between coverage quality and energy consumption in the design space.

In our definition, we assume that time is slotted and at any slot, only k nodes are active. k is a parameter that controls the amount of energy consumption expected in the system and is assumed to be given. We will need to choose which set of nodes to turn on and in what order to ensure

that the coverage quality remains high. Since not all sensors are on all the time, a point $p \in \mathcal{D}$ is not necessarily covered at all times. Thus we say p is *lighted up* if any of the active nodes at this slot covers p; and *dark* otherwise. Naturally we would like the dark period to be as short as possible. Rigorously, we define for each point $p \in \mathcal{D}$ the *max dark length* $T(p)$ such that p is lighted up at least in some slot for *every* contiguous $T(p)$ time slots. It is exactly one plus the maximum length of any contiguous dark period. An example of $T(p)$ is shown in Figure 1.
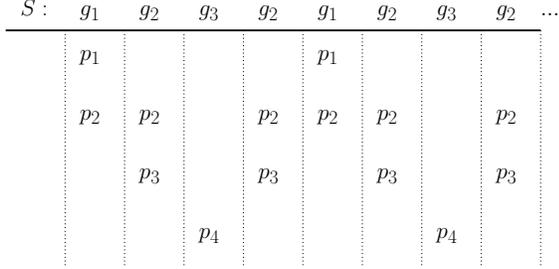
Figure 1. An example of four target points p_1, p_2, p_3, p_4 and three sensor nodes g_1, g_2, g_3. $P(g_1) = \{p_1, p_2\}$, $P(g_2) = \{p_2, p_3\}$, $P(g_3) = \{p_4\}$. $k = 1$. Suppose we have a schedule S that repeatedly turns on g_1, g_2, g_3, g_2 in this order. With this schedule the max dark length for the four targets are the following: $T(p_1) = 4, T(p_2) = 2, T(p_3) = 2, T(p_4) = 4$.

Now we formulate two research problems.

1. **Min Max Dark Length Scheduling.** Find a schedule S as an ordered list in which S_i is the set of nodes staying on in the ith time slot, $|S_i| \leq k$ such that the maximum dark length for any point $p \in \mathcal{D}$ is minimized:

$$\min \max_{p \in \mathcal{D}} T(p).$$

2. **Min Average Dark Length Scheduling.** Find a schedule S as an ordered list in which S_i is the set of nodes staying on in the ith time slot, $|S_i| \leq k$ such that the weighted average of maximum dark length for all points $p \in \mathcal{D}$ is minimized:

$$\min \sum_{p \in \mathcal{D}} w(p) \cdot T(p).$$

Here $w(p)$ is a weight parameter and can be set higher for locations that are more important in applications.

In the first problem we minimize the maximum dark length for all points in the domain, achieving a worst case guarantee. In the second problem formulation, we try to achieve a good average and further we can use the weight $w(\cdot)$ to allow differentiation of different target points in $\mathcal{D}$. In the following we will discuss the two versions separately.

4. MIN MAX DARK LENGTH SCHEDULING

The intuitive explanation for this objective function is that we want to cover every point regularly and we want to minimize the maximum time period that some point is uncovered, when only k nodes can be active during any time slot. Define C as the objective function $\max_{p \in D} T(p)$. First

we look at an intuitive schedule, by simply taking every k nodes in an arbitrary sequence of the n nodes in a Round Robin manner. That is, the first k nodes stay active in slot 1, the next k nodes stay active in slot 2 and so on. When n is not a multiple of k, the last time slot is filled up by arbitrarily other nodes. This schedule with period of $\lceil \frac{n}{k} \rceil$ is then repeated indefinitely.

Theorem 4.1. *Given a domain $\mathcal{D}$ and a set of n nodes $\mathcal{V}$, at any time only k nodes are operating simultaneously, a round robin scheduling has max dark length at most $\lceil \frac{n}{k} \rceil$.*

PROOF. For every $\lceil \frac{n}{k} \rceil$ slots in an arbitrary sequence of the n nodes, each sensor node appears at least once in the schedule. Since all nodes together cover $\mathcal{D}$, every point of $\mathcal{D}$ is covered at least once in one period. Therefore, $\forall p \in \mathcal{D}, T(p) \leq \lceil \frac{n}{k} \rceil$. $\square$

Given a set of node $\mathcal{V}$, we say $\mathcal{V}$ is a *minimal cover*, if each node of $\mathcal{V}$ covers a unique target point (which is not covered by any other nodes of $\mathcal{V}$). In other words, every node in $\mathcal{V}$ is crucial and cannot be omitted if we want $\mathcal{D}$ to be covered. In this case, the round robin scheduling is optimal. To show that we first state a useful property of an optimal schedule.
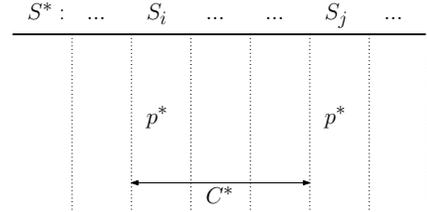
Figure 2. When each node covers a unique target, the round robin scheduling is optimal.

Suppose S^* is the optimal schedule and p^* be a point that $T(p^*) = C^*$. In this optimal schedule the set of active nodes in the ith slot is S_i. Assume that the optimal value C^* is achieved by p^* covered in S_i and S_j with $j > i$ but not in any slot in between. $C^* = j - i$. See Figure 2. For any other points $p' \in \mathcal{D} \setminus p^*$, $T(p') \leq T(p^*)$. So p' must be covered at least once between the ith and $(j-1)$th slot. Therefore, the union of the nodes that are active from the ith slot to the $(j-1)$th slot, $\cup_{\ell=i}^{j-1} S_\ell$, covers all points of $\mathcal{D}$. We denote the schedule $\bar{S} = S_i, S_{i+1}, \cdots, S_{j-1}$ a *crucial* sub-schedule of S^*.

Theorem 4.2. *If $\mathcal{V}$ is a minimal cover of $\mathcal{D}$, the Round Robin schedule is optimal.*

PROOF. We just need to show that the optimal scheduling achieves dark length no less than $\lceil \frac{n}{k} \rceil$. Suppose S^* is the optimal schedule and suppose $C^* < \lceil \frac{n}{k} \rceil$. Therefore, the total number of nodes that are active in the crucial sub-schedule is $|\cup_{\ell=i}^{j-1} S_\ell| \leq kC^* < n$. Thus at least one sensor node (say node x) is not active in this period of time. The target node p that is uniquely covered by x is not covered. This means p must have dark length $T(p) > C^*$, which is a contradiction that C^* is the optimal value. $\square$

When $\mathcal{V}$ is not a minimal cover, the challenge is to find a minimum number of nodes in $\mathcal{V}$ that provide full coverage of the domain $\mathcal{D}$, denoted by the *minimum cover*, which by itself is a classical set cover problem and thus is NP-hard.

But once this minimum set cover solution is given, doing a periodic round robin scheduling is the optimal.

Theorem 4.3 (*NP*-hardness). *Given a domain $\mathcal{D}$, a set of n nodes $\mathcal{V}$ among which k nodes operate simultaneously, we have two statements:*

1. *Repeating a Round Robin schedule on a minimum cover is the optimal schedule that minimizes the max dark length of all points of $\mathcal{D}$.*

2. *But finding this optimal schedule is NP-hard.*

PROOF. To prove the first statement, suppose S^* is an optimal schedule. Recall that the nodes that are active in a crucial subschedule cover all points of $\mathcal{D}$. This means that C^* is no smaller than $\lceil n^*/k \rceil$, where n^* is the size of the *minimum* cover. Since a round robin schedule using a minimum cover achieves max dark length of $\lceil n^*/k \rceil$, by Theorem 4.1, the optimality follows.

For the second claim, we use a reduction from the Set Cover problem. Given an instance of the Set Cover Problem $(\mathcal{U}, \mathcal{S})$, $\mathcal{U}$ is the universe and $\mathcal{S}$ contains sets $\{G_i\}$, where $G_i \subset \mathcal{U}$. We take $\mathcal{D}$ as the universe, and create a set of nodes where node i covers the subset $G_i \in \mathcal{S}$. We also set $k = 1$ in this reduction – if this special case of $k = 1$ is NP-hard it is obvious that the general problem is NP-hard as well. Our goal is to show that if we can solve the max dark length problem then we find the optimal set cover. For that, we will use two steps.

1. Claim 1: Any optimal schedule S^* can be turned into a periodic one.

2. Claim 2: For a periodic optimal schedule, the nodes that ever appear in the schedule is a minimum cover of $\mathcal{D}$.

Therefore, any algorithm that solves the max dark length problem basically solves the Set Cover problem.

For Claim 1, using the same setting as before, we can simply repeat the crucial subschedule $\tilde{S} = S_i S_{i+1} ... S_{j-1}$ indefinitely which achieves the same optimal value.

For Claim 2, recall that the union of the nodes that are active in the crucial subschedule covers all points of $\mathcal{D}$. Since $k = 1$ this must be a minimum cover — otherwise we take round robin on a minimum cover and we get a contradiction on optimality. □

To summarize the discussion on this metric, if $\mathcal{V}$ is a minimal cover of $\mathcal{D}$, we can find the optimal schedule immediately by repeating the Round Robin sub-schedule. Otherwise, we need to find a minimum cover of $\mathcal{D}$ to obtain the optimal schedule which is NP-hard, but has an $O(\log n)$ approximation by a greedy algorithm.

5. MIN AVERAGE DARK LENGTH SCHEDULING

While the previous version provides the worst case guarantee of coverage quality for any point in $\mathcal{D}$, it is possible that a small portion of $\mathcal{D}$ controls the worst case value. For example, suppose we have three nodes g_1, g_2, g_3. g_1 covers only one target p_1; g_2 covers one target p_2; and g_3 covers four targets p_3, p_4, p_5, p_6. In the optimal min max schedule we will rotate between these three nodes and each target

has a dark length of 3. But in the min average optimization we would favor g_3 (i.e., schedule it more frequently) as its coverage range is greater. To see that, just verify that the average dark length for repeating the schedule g_1, g_3, g_2, g_3 is $8/3 \leq 3$. The min average optimization is able to adjust to the differences in coverage range and in addition we can use the weight function $w(\cdot)$ to provide better coverage quality for targets of high importance. We assume that the weights are integer values.

Before we present our algorithm we first analyze properties of the optimal solution.

Theorem 5.1 (Periodicity). *There is a schedule for min average dark length problem which is periodic and is a 2-approximation.*

PROOF. We define $A^* = \sum_{p \in \mathcal{D}} w(p) \cdot T(p)$ as the optimal value. We also define $A(t)$ as the weighted average of the schedule in the first t time slots, i.e., $A(t) = \sum_{p \in \mathcal{D}} w(p) T_t(p)$, where $T_t(p)$ is the maximum dark length in the first t slots for target point p. A^* and $A(t)$ for any t are integers. In fact the sequence $A(t)$ when t goes to infinity will converge to A^*. For any target point p we notice that $T_t(p)$ is a non-decreasing function when t goes to infinity. Thus $A(t)$ is also non-decreasing. For a non-decreasing sequence to converge to a fixed value A^*, it means that for any $\varepsilon > 0$ there is a value T such that for any $t \geq T$, $|A(t) - A^*| \leq \varepsilon$. If we take $\varepsilon = 1/2$ we know that $A(t) = A^*$ for any $t \geq T$, since both $A(t)$ and A^* are integers. By the same argument we can say that $T_t(p) = T(p)$ for any $t \geq T$.

Now suppose the set of sensors active in the tth slot is denoted as S_t. We take the following sequence and repeat indefinitely:

$$S_1, S_2, \cdots S_{T-1}, S_T, S_{T-1}, \cdots, S_2.$$

The new schedule S' is a periodic schedule. Now we argue that for each target point $p \in \mathcal{D}$, $T'(p) \leq 2T(p)$, where $T'(p)$ is the maximum dark length in the new schedule S'.

Recall that for each point p, $T(p)$ is realized in the first T slots. We take $j \leq T$ to be the highest value such that S_j covers p, and $i \geq 1$ to be the smallest value such that S_i covers p. Clearly, $T - j \leq T(p)$; $i - 1 \leq T(p)$. Therefore, consider the dark gap formed by

$$S_{j+1}, S_{j+2}, \cdots, S_{T-1}, S_T, S_{T-1}, \cdots, S_{j+1}.$$

The dark length is no greater than $2T(p)$. Similarly, for the dark gap formed by

$$S_{i-1}, S_{i-2}, \cdots, S_2, S_1, S_2, \cdots S_{i-1}$$

the dark length is no greater than $2T(p)$. This proves the claim that the new schedule S' is a 2-approximation. □

Motivated by the above theorem, we will only consider periodic schedules. The best periodic schedule is a 2-approximation of the optimal. In the following we first look at a special case when the schedule is formed by taking a permutation of the n nodes and performing round-robin on the sequence. Then we generalize the results to a periodic schedule without the permutation requirement.

5.1 Scheduling with a Permutation

In this part we first look at a special case, when the schedule is determined by doing round robin on a permutation π of the nodes. We can define the optimal schedule with this

additional restriction as the one that minimizes the objective function $\sum_{p \in \mathcal{D}} w(p) \cdot T(p)$. We denote the permutation used in the optimal schedule as π^*.

$G(p) \subseteq V$, $|V| = 16$ $\ell = |G(p)| = 6$ bins
$G(p) = \{g_1, g_2, g_3, g_4, g_5, g_6\}$ $m = |V \setminus G(p)| = 10$ balls

Assume $G(p)$ appears in the permutation with this order $g_1, g_3, g_4, g_6, g_5, g_2$
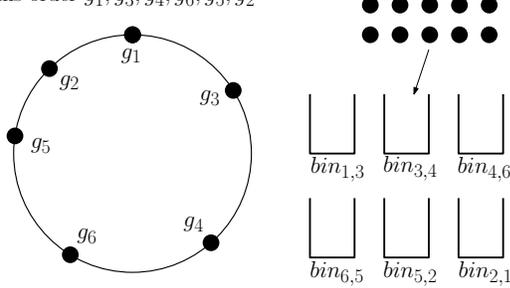
Figure 3. Repeat a random permutation of nodes in $\mathcal{V}$.

Consider a point $p \in \mathcal{D}$, define $G(p)$ as the collection of nodes of $\mathcal{V}$ that cover p. Therefore the schedule with a permutation π can be represented as placing the nodes in the order of π on a circle. See Figure 3. The nodes of $G(p)$ are spread out along the circle. Therefore the maximum dark length $T(p)$ is precisely one plus the maximum gap between these nodes of $G(p)$ on the circle. Suppose $|G(p)| = h(p)$. Thus the nodes that are not in $G(p)$ need to be placed in the gaps between the nodes of $G(p)$ on the circle. This gives a trivial lower bound on $T(p)$ for any permutation π:

$$T_\pi(p) \geq \lceil 1 + \frac{|\mathcal{V} \setminus G(p)|}{k \cdot |G(p)|} \rceil = 1 + \lceil \frac{n - h(p)}{k \cdot h(p)} \rceil = B(p).$$

Therefore,

$$A_\pi = \sum_{p \in D} w(p) T_\pi(p) \geq \sum_{p \in D} w(p) \cdot [1 + \lceil \frac{n - h(p)}{k \cdot h(p)} \rceil] = B.$$

The lower bound on the right hand side of the above inequality is usually impossible to achieve. But we show that by using a *random* permutation π we can actually get a good approximation to the lower bound B.

We now compute the expectation of the objective function with respect to running round robin on a random permutation using linearity of expectation:

$$\mathbf{E}[A_\pi] = \mathbf{E}[\sum_{p \in \mathcal{D}} w(p) T_\pi(p)] = \sum_{p \in \mathcal{D}} w(p) \cdot \mathbf{E}[T_\pi(p)]$$

To analyze $\mathbf{E}[T_\pi(p)]$ we will relate to the classical 'Balls and Bins' problem. For a particular target $p \in \mathcal{D}$, let $\ell = |G(p)|$ and $m = |\mathcal{V} \setminus G(p)|$. $B(p) \geq 1 + m/(k\ell)$. Then $T_\pi(p)$ under a random permutation is related to the maximum load of throwing m balls into ℓ bins independently and randomly. Figure 3 illustrates this connection.

The maximum load in the balls and bins game has been studied a lot in the literature [47, 48]. We will basically state the results.

- If $m = \Omega(\ell \log \ell)$, then $\mathbf{E}[T_\pi(p)] \leq 1 + O(1) \cdot m/(k\ell) = 1 + O(1) \cdot B(p)$, where $B(p)$ is the best possible (assuming balls are uniformly put in the bins).

- If $m = o(\ell \log \ell)$, $\mathbf{E}[T_\pi(p)] \leq (1 + O(1))\frac{\log \ell}{k \log \log \ell} \leq O(1)\frac{\log n}{k \log \log n}$.

If we take $k \geq \log n / \log \log n$, then from the above results we see that $\mathbf{E}[T(p)] \leq O(1) \cdot B(p)$. Otherwise, $\mathbf{E}[T_\pi(p)] \leq O(\log n / \log \log n) \cdot B(p)$.

Theorem 5.2 (Random Permutation). *Given n nodes $\mathcal{V}$ that cover a domain $\mathcal{D}$, the solution of doing round robin of a random permutation π gives a α-approximation for the min average dark length scheduling problem with additional restriction that only round robin of a permutation is used. Here $\alpha = O(1)$ if $k \geq \log n / \log \log n$, and $\alpha = O(\log n / \log \log n)$ otherwise.*

5.2 Approximation Algorithm

The observation from the previous section can be used to provide an approximation algorithm for the general case. We consider the optimal *periodic* schedule. By Theorem 5.1 if we have a β approximation of the best periodic schedule then it is a 2β approximation for the optimal solution (possibly non-periodic).

Suppose S is a periodic schedule. We consider the schedule of one period, denoted by $\bar{S}$. $\bar{S}$ is not necessarily a permutation of $\mathcal{V}$. In fact, certain nodes in $\mathcal{V}$ may not show up in $\bar{S}$ at all, while some other nodes of $\mathcal{V}$ may show up multiple times in $\bar{S}$. We denote by $\tau(g)$ the number of times that a node $g \in \mathcal{V}$ appears in one full cycle. $\tau(g)$ is always an integer. $\tau(g) \geq 0$. Denote by T the length of one full cycle in the schedule S. Then we know $\sum_{g \in \mathcal{V}} \tau(g) = T \cdot k$.

Suppose we are able to guess the correct value of $\tau(g)$ and T, then we can simply reduce the problem to the special case studied in Theorem 5.2 and get an $O(\alpha)$ approximation. In the following we show how to obtain good values for $\tau(g)$ and T to achieve that.

If we define $f(g) = \frac{\tau(g)}{kT}$ as the frequency of occurrence of g in the schedule S, we have

$$\sum_{g \in V} f(g) = \sum_{g \in V} \frac{\tau(g)}{kT} = 1.$$

For each point $p \in \mathcal{D}$, define the number of times that p is covered in one cycle by $X(p) := \sum_{g \in G(p)} \tau(g)$. Thus, $T(p)$, the largest dark length in the schedule, is lower bounded by the average separation if all occurrences of nodes that cover p are uniformly spread out:

$$T(p) \geq \frac{T}{X(p)} = \frac{1}{\sum_{g \in G(p)} \tau(g)/T} = \frac{1}{k} \cdot \frac{1}{\sum_{g \in G(p)} f(g)}$$

The objective function for S is lower bounded in the following inequality:

$$A = \sum_{p \in \mathcal{D}} w(p) T(p) \geq \sum_{p \in \mathcal{D}} \frac{w(p)/k}{\sum_{g \in G(p)} f(g)} = B$$

Notice that in the above discussion we didn't make any assumption on what S is. Take the best periodic schedule S^* the inequality holds.

$$A^* = \sum_{p \in \mathcal{D}} w(p) T^*(p) \geq \sum_{p \in \mathcal{D}} \frac{w(p)/k}{\sum_{g \in G(p)} f^*(g)} = B(f^*(g), \forall g \in \mathcal{V})$$

The right hand side (i.e., B) is a lower bound of the optimal value, in which the only parameter is the frequency $f(g)$ of each node g. In the following we use two steps to get a good approximation.

1. We show that we can *minimize* the value B in polynomial time by choosing proper values of $f(g)$ for all $g \in \mathcal{V}$.

2. We show an algorithm that achieves a β approximation of the *minimum* possible B, which immediately gives us a β approximation to the best periodic schedule, or a 2β approximation to the optimal schedule.

Lemma 5.3. *The following optimization problem is convex and thus can be solved in polynomial time:*

$$\min \sum_{p \in D} \frac{w(p)}{\sum_{g \in G(p)} f(g)}$$

subject to

$$\sum_{g \in G} f(g) = 1, f(g) \geq 0$$

PROOF. In this minimization problem the only parameters are the values $f(g)$ for $g \in \mathcal{V}$. We only need to show that the objective function is convex. Since the sum of convex functions is still convex, we only need to verify that $1/\sum_{g \in G(p)} f(g)$ is convex. In other words,

$$\frac{1}{\sum_{g \in G(p)} [f(g) + f'(g)]} \leq \frac{1}{\sum_{g \in G(p)} f(g)} + \frac{1}{\sum_{g \in G(p)} f'(g)}$$

This can be verified easily. $\square$

Remark that in the optimization problem it is easiest to understand when the domain $\mathcal{D}$ has a discrete, finite number of target points. But the same can be applied when the domain $\mathcal{D}$ is continuous. In this case, we can find the arrangement defined by the sensor coverage ranges (i.e., the decomposition of the domain $\mathcal{D}$ into pieces $\mathcal{D}_i$ such that all points in each piece $\mathcal{D}_i$ are covered by the same subset of sensors). The number of pieces in the arrangement is bounded, typically in polynomial of n. For example, if line of sight constraints is the only constraint in sensor coverage, with n sensors inside a polygon of h vertices, the number of pieces in the arrangement is bounded by $O(n^2 h^2)$. Then the summation over all points of p boils down to a weighted sum of the finite number of pieces in the arrangement, which can be solved in the same manner.

Now we solve this optimization problem for the best frequency $f'(g)$ for all $g \in \mathcal{V}$, which are real numbers. Since $f'(g)$ minimizes the function B, we have

$$B(f'(g), \forall g \in \mathcal{V}) \leq B(f^*(g), \forall g \in \mathcal{V}) \leq A^*$$

We can approximate these real number $f'(g)$ by rational numbers $\hat{f}(g)$ while losing only a tiny factor $(1 + \varepsilon)$ for an $\varepsilon > 0$ that can be made arbitrarily small:

$$B(\hat{f}(g), \forall g \in \mathcal{V}) \leq B(f'(g), \forall g \in \mathcal{V})(1 + \varepsilon)$$

For rational numbers $\hat{f}(g)$ we can find integer values $\hat{\tau}(g)$ for all $g \in \mathcal{V}$ and integer value $\hat{T}$ such that $\hat{f}(g) = \hat{\tau}(g)/(k\hat{T})$. The scale of $\hat{T}$ depends on the resolution of $\hat{f}(g)$ which depends on the error ε. Thus for a fixed error ε, $\hat{T}$ is bounded.

Now we have a set of nodes $\mathcal{V}$ and for each node $g \in \mathcal{V}$ we repeat it $\hat{\tau}(g)$ times so that we have a new set of nodes $\hat{\mathcal{V}}$. By the discussion in the previous section, we simply choose a random permutation of $\hat{\mathcal{V}}$ and run round robin on it. This generates a schedule whose average dark length

value is at most α times the value $B(\hat{f}(g), \forall g \in \mathcal{V})$, according to Theorem 5.2. Therefore the random permutation using repetition $\hat{\tau}(g)$ for $g \in \mathcal{V}$ will generate average dark length value at most

$$\alpha \cdot B(\hat{f}(g), \forall g \in \mathcal{V}) \leq \alpha(1 + \varepsilon) \cdot A^*,$$

This finishes the proof for the following theorem.

Theorem 5.4. *In polynomial time we can generate a schedule which is round robin on a random permutation in which node g is repeated $\hat{\tau}(g)$ times, where $\hat{\tau}(g)$ is produced by the convex optimization procedure. This schedule achieves an $(2 + \varepsilon)\alpha$ approximation to the optimal schedule. Here $\alpha = O(1)$ if $k \geq \log n / \log \log n$, and $\alpha = O(\log n / \log \log n)$ otherwise.*

At the end, we would like to summarize the algorithm for the sake of clarity:

1. Solve the convex optimization problem for $f(g)$ for each node $g \in \mathcal{V}$:

$$\min \sum_{p \in \mathcal{D}} \frac{w(p)}{\sum_{g \in G(p)} f(g)}$$

subject to

$$\sum_{g \in \mathcal{V}} f(g) = 1, f(g) \geq 0$$

2. Suppose the optimal values for the above problem is $f'(g)$ for $g \in \mathcal{V}$, we take rational numbers $\hat{f}(g)$ to approximate $f'(g)$ such that the objective function is only losing at most a factor $(1 + \varepsilon)$ for $\varepsilon > 0$.

3. With $\hat{f}(g)$ being rational numbers, we find integers $\hat{\tau}(g)$ and $\hat{T}$ such that $\hat{f}(g) = \hat{\tau}(g)/(k\hat{T})$.

4. Given the nodes $\mathcal{V}$, we generate $\hat{\mathcal{V}}$ which is a random permutation in which node $g \in \mathcal{V}$ is repeated $\hat{\tau}(g)$ times. We run round robin on $\hat{\mathcal{V}}$.

We remark that although the analysis of the theoretical guarantee of the above algorithm is non-trivial, the algorithm itself is essential taking a random permutation which is sufficiently simple and practically appealing.

6. SIMULATION

We ran simulations to evaluate the performance of the scheduling algorithm. We consider the scenario of an indoor localization application, in which the topology of a room is represented by a polygon and the possible target points (the domain set $\mathcal{D}$) monitored are selected by a grid sampling of the interior of the polygon. Then we run a Greedy Set Cover to find the set of vertex guards (the node set $\mathcal{V}$) that can fully cover these target points since the general Art Gallery Problem is NP-hard. To maintain a full coverage of the room, all guards in the greedy set need to be turned on at all times. This is the baseline energy consumption.

Assume $k = 1$ and each time slot is 1 second in the localization system. Now we can only turn on one node per second instead of $|\mathcal{V}|$ nodes. So the targets are not covered in every time slot. We want to know the dark length of the target points to understand the detection delay in such a application.
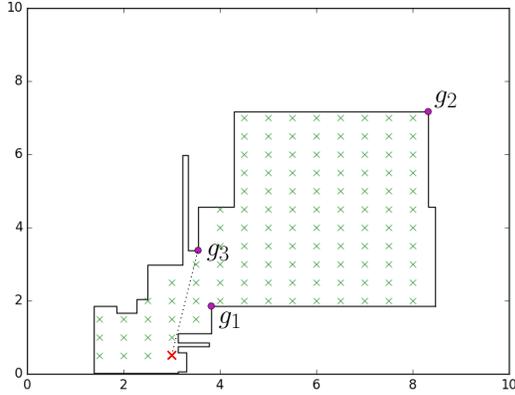
Figure 4. A topology with 113 targets in the domain and 3 nodes to cover the domain. The target point represented by a red cross is uniquely covered by node g_3.

Figure 4 illustrates an example with $|\mathcal{D}| = 113$ and $|\mathcal{V}| = 3$. There are three nodes where g_1 covers 109 targets, g_2 covers 108 targets and g_3 covers 101 targets. If we want to minimize the average dark length, we run the approximation algorithm in Section 5.2. After the optimization, the weight of g_1, g_2 and g_3 are 0.4173, 0.3883 and 0.1944 respectively. In a periodic schedule with the length of period = 100 time slots, g_1, g_2 and g_3 should appear 42, 39 and 19 times respectively. Random permutation according to the weights gives out an average dark length of 1.74 time slots for each target. The weight of g_3 is lowest because it covers the least number of target but is necessary as a target point (represented by red cross) is uniquely covered by g_3. The maximum dark length of this periodic schedule is very high as g_3 is not frequently scheduled so it is dominated by the dark length of this unique target point. This illustrates the difference between min average dark length and min max dark length.

We simulated with 10 different room topologies (A few figures are provided after the conclusion). If we want to minimize max dark length, we need to find the minimum cover. The Greedy Set Cover solution is a $O(\log n)$-approximation to the optimal cover. For min average dark length, we run the algorithm to get the periodic schedule. The result of the approximation scheduling algorithm is reported in Table 1. Note that this algorithm tries to minimize the average dark length but not the maximum dark length so there is a large difference between average and max dark length of the schedule.

For such a localization application and topology, if we can only turn on one sensor in each time slot, the mean detection delay is 3.588 time slots which could be less than 2 seconds in our experience of infrared localization system.

7. TESTBED EVALUATIONS

We had implemented the approximation scheduling algorithm on a camera network. The objective of this network is to detect the random occurrence of an event in a room. We had deployed networked camera nodes to monitor the testing environment. One camera node is shown in Figure 5. A full description of a similar network can be found in [39]. Sev-

Table 1. Approximation algorithm on 10 topologies.

Topology	$|\mathcal{D}|$	$|\mathcal{V}|$	Avg Dark Len.	Max Dark Len.
1	114	3	3.34	9
2	113	3	1.74	26
3	135	4	4.15	26
4	181	4	6.13	20
5	175	3	3.45	13
6	165	2	1.96	6
7	146	4	3.97	22
8	157	3	4.11	13
9	118	2	1.9	9
10	243	3	5.13	13

eral valuable points on the ground are marked (the weight $w(p)$ is non-zero). The coverage of those points with respect to each camera is recorded ($P(g)$ is known). The resource constraint on the system is designed as follows: only one camera node can be turned on in each 5 seconds time slot ($k = 1$) and the camera takes a picture in the beginning of the slot. With the information about the weight, coverage and resource constraint, we can run the scheduling algorithm to determine an efficient ordering for the nodes to be turned on in each time slot.

Figure 5. A camera node. Each node is able to communicate over WiFi, take pictures and store pictures in the USB drive.

An event is represented by a blinking LED light bulb controlled by a Raspberry-PI in Figure 6. This is analogous to a fire event or smoking behavior. Each blinking event lasts for 6 seconds. Therefore, if a camera that can cover the event point is turned on during the period, the event is detected and reported.

A case study of 4 target points and 3 camera nodes in a room is illustrated in Figure 7. A target point is represented by a red cross. Therefore, $P(g_1) = \{p_1, p_3\}, P(g_2) = \{p_1, p_2\}, P(g_3) = \{p_2, p_4\}$. If the targeta are equally important $w(p_1) = w(p_2) = w(p_3) = w(p_4)$, then $f(g_2)$ is set to be 0 in our approximation algorithm. Therefore, we will alternate between camera g_1 and g_3. Instead of turning on 3 cameras all the time, we will cut the energy consumption to $\frac{1}{3}$ while increasing the detection delay by one time slot (5 seconds). Our experiment confirms the delay but we cannot achieve exactly the theoretical energy saving because of system overhead. Keeping the cameras connected and ready
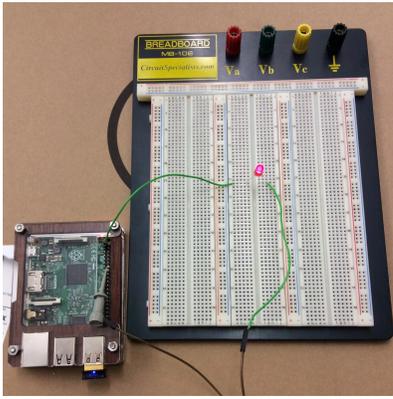
Figure 6. The LED bulb is connected and controlled by the program on Raspberry-PI. The timing of the random blinking event is recorded.

causes energy even when they are not taking photos. The energy saving is around 60% for the experiments.
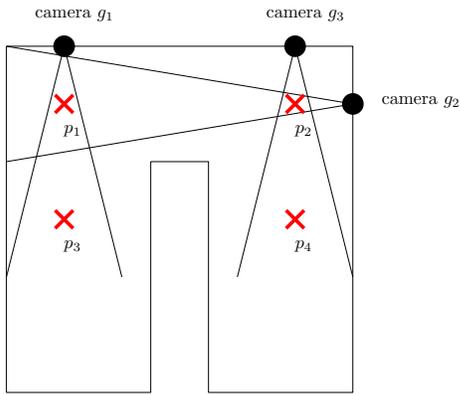
Figure 7. Physical layout of experiments on camera network in a room. 3 cameras (black dot) is deployed to cover 4 target points (red cross).

The weights of the target points also play an important role. Suppose that the event occurs more likely at p_1 and p_2. Then the weight of p_1 and p_2 is increased. For example, we can change the control program on Raspberry-PI so the probability of blinking event is higher at p_1 and p_2. Then the frequency of g_2 will be non-zero and higher in the schedule.

8. CONCLUSION

In this paper we proposed the first set of solutions for joint sensor duty cycle scheduling for indoor applications when the sensors collaboratively cover the domain with quality guarantee while significantly reducing energy consumption. We remark that the efficiency of the algorithm makes it appealing to be applied in practical applications. There are a couple of open questions that we wish to answer in future work. First we would like to know if it is NP-hard to find the optimal schedule in the second optimization metric (min max dark length). Second we wish to establish hard coverage frequency requirement for each target and develop efficient scheduling algorithms.
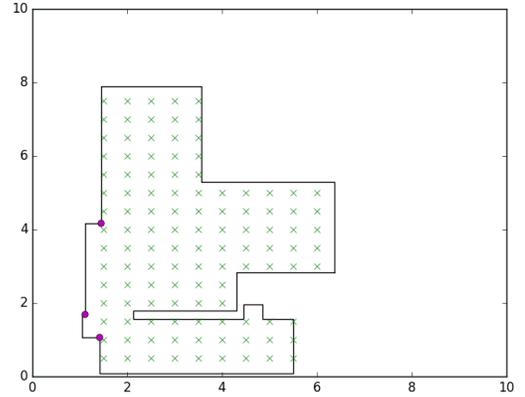
Figure 8. Topology 1
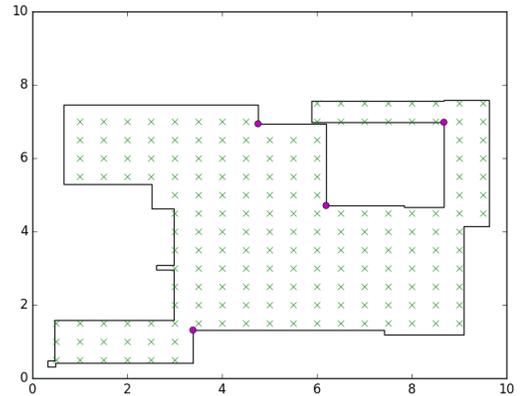
Figure 9. Topology 4
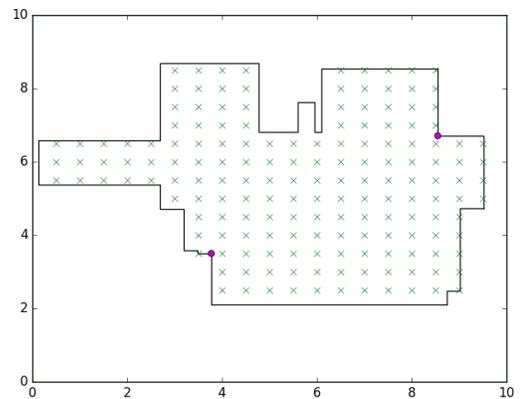
Figure 10. Topology 6
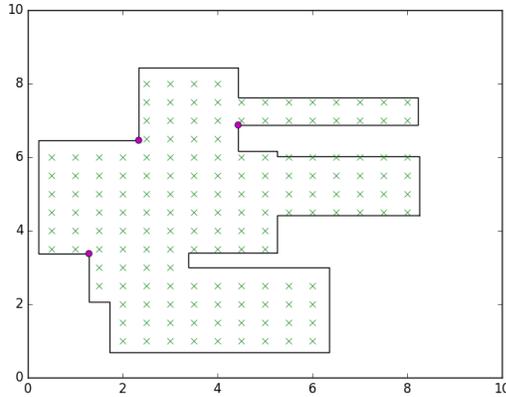
Figure 11. Topology 8

Acknowledgment

9. REFERENCES

[1] T. He, S. Krishnamurthy, L. Luo, T. Yan, L. Gu, R. Stoleru, G. Zhou, Q. Cao, P. Vicaire, J. A. Stankovic, T. F. Abdelzaher, J. Hui, and B. Krogh, "Vigilnet: An integrated sensor network system for energy-efficient surveillance," *ACM Trans. Sen. Netw.*, vol. 2, no. 1, pp. 1–38, 2006.

[2] M. Kushwaha and X. Koutsoukos, "3d target tracking in distributed smart camera networks with in-network aggregation," in *Proceedings of the Fourth ACM/IEEE International Conference on Distributed Smart Cameras*, ser. ICDSC '10. New York, NY, USA: ACM, 2010, pp. 25–32. [Online]. Available: http://doi.acm.org/10.1145/1865987.1865992

[3] J. Lu, D. Birru, and K. Whitehouse, "Using simple light sensors to achieve smart daylight harvesting," in *Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, ser. BuildSys '10, 2010.

[4] Y. Gu, A. Lo, and I. Niemegeers, "A survey of indoor positioning systems for wireless personal networks," *Commun. Surveys Tuts.*, 2009.

[5] D. Das, A. Abouzeid, and M. Codreanu, "Opportunistic scheduling and relaying in a cooperative cognitive network," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt), 2014 12th International Symposium on*, 2014.

[6] C. Kam, S. Kompella, A. Ephremides, and I. S. Moskowitz, "Sinr-based scheduling for minimum latency broadcast," in *2015 IEEE International Conference on Communications, ICC 2015, London, United Kingdom, June 8-12, 2015*.

[7] R. Li and A. Eryilmaz, "Scheduling for end-to-end deadline-constrained traffic with reliability requirements in multihop networks," *IEEE/ACM Trans. Netw.*

[8] J. Kwak, C.-H. Lee, and D. Y. Eun, "A high-order markov chain based scheduling algorithm for low delay in csma networks," in *INFOCOM, 2014 Proceedings IEEE*, 2014.

[9] M. G. Markakis, E. Modiano, and J. N. Tsitsiklis, "Max-weight scheduling in queueing networks with heavy-tailed traffic," *IEEE/ACM Trans. Netw.*, 2014.

[10] K. Jagannathan, L. Jiang, and E. Modiano, "On scheduling algorithms robust to heavy-tailed traffic," in *Information Theory and Applications Workshop (ITA), 2012*, 2012.

[11] R. Agrawal, V. Subramanian, and R. Berry, "Joint scheduling and resource allocation in cdma systems," in *Proc. of 2nd Workshop on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt '04)*, 2004.

[12] S. Shakkottai, "Effective capacity and qos for wireless scheduling," *Automatic Control, IEEE Transactions on*, 2008.

[13] S. Moharir and S. Shakkottai, "Maxweight vs. backpressure: Routing and scheduling in multi-channel relay networks," in *INFOCOM, 2013 Proceedings IEEE*, 2013.

[14] M. Gribaudo, D. Manini, A. Nordio, and C.-F. Chiasserini, "Analysis of ieee 802.15.4 sensor networks for event detection." in *GLOBECOM*, 2009.

[15] T. F. L. Porta, C. Petrioli, and D. Spenza, "Sensor-mission Assignment in Wireless Sensor Networks with Energy Harvesting," in *Proceedings of IEEE SECON 2011*, 2011.

[16] Y. Yang, I. Hou, J. C. Hou, M. Shankar, and N. S. V. Rao, "Sensor placement for detecting propagative sources in populated environments," in *INFOCOM 2009. 28th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 19-25 April 2009, Rio de Janeiro, Brazil*.

[17] J. He, L. Duan, F. Hou, P. Cheng, and J. Chen, "Multiperiod scheduling for wireless sensor networks: A distributed consensus approach." *IEEE Transactions on Signal Processing*, 2015.

[18] L. Huang, S. Zhang, M. Chen, and X. Liu, "When backpressure meets predictive scheduling," *Networking, IEEE/ACM Transactions on*, 2015.

[19] S. Chen, P. Sinha, N. B. Shroff, and C. Joo, "A simple asymptotically optimal joint energy allocation and routing scheme in rechargeable sensor networks," *IEEE/ACM Trans. Netw.*

[20] X. Xu, X.-Y. Li, P.-J. Wan, and S. Tang, "Efficient scheduling for periodic aggregation queries in multihop sensor networks," *IEEE/ACM Trans. Netw.*

[21] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the complexity of scheduling in wireless networks," in *Proceedings of the 12th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '06, 2006.

[22] K. Balachandran, J. Kang, and W. C. Lau, "Adaptive sleeping and awakening protocol (asap) for energy efficient adhoc sensor networks," in *Communications, 2005. ICC 2005. 2005 IEEE International Conference*

on, 2005.

[23] J. Ma, W. Lou, and X.-Y. Li, "Contiguous link scheduling for data aggregation in wireless sensor networks," *Parallel and Distributed Systems, IEEE Transactions on*, 2014.

[24] J. Kim, X. Lin, N. B. Shroff, and P. Sinha, "Minimizing delay and maximizing lifetime for wireless sensor networks with anycast," *IEEE/ACM Trans. Netw.*, 2010.

[25] H. Li, X. Cheng, K. Li, X. Xing, and T. Jing, "Utility-based cooperative spectrum sensing scheduling in cognitive radio networks," in *IEEE INFOCOM Mini-Conference*, 2013.

[26] Y. Yi, J. Zhang, and M. Chiang, "Delay and effective throughput of wireless scheduling in heavy traffic regimes: Vacation model for complexity," in *Proceedings of the Tenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '09.

[27] L. Yang, J. Zhang, and H. Poor, "Risk-aware day-ahead scheduling and real-time dispatch for electric vehicle charging," *IEEE Transactions on Smart Grid*, 2014.

[28] L. Huang and M. J. Neely, "Utility optimal scheduling in energy harvesting networks," in *Proceedings of the Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '11, 2011.

[29] L. Shi, M. Epstein, B. Sinopoli, and R. M. Murray, "Effective sensor scheduling schemes in a sensor network by employing feedback in the communication loop," Nov 2006.

[30] Q. Liu, X. Wang, and Y. Cui, "Scheduling of sequential periodic sensing for cognitive radios," in *INFOCOM, 2013 Proceedings IEEE*, 2013.

[31] L. Xu, J. Cao, S. Lin, H. Dai, X. Wu, and G. Chen, "Energy-efficient broadcast scheduling with minimum latency for low-duty-cycle wireless sensor networks," in *Mobile Ad-Hoc and Sensor Systems (MASS), 2013 IEEE 10th International Conference on*, 2013.

[32] T. He, S. Krishnamurthy, J. A. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh, "Energy-efficient surveillance system using wireless sensor networks," in *Proceedings of the 2Nd International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '04, 2004.

[33] J. Hsu, S. Zahedi, A. Kansal, M. Srivastava, and V. Raghunathan, "Adaptive duty cycling for energy harvesting systems," in *Low Power Electronics and Design, 2006. ISLPED'06. Proceedings of the 2006 International Symposium on*, 2006.

[34] H. Yoo, M. Shim, and D. Kim, "Dynamic duty-cycle scheduling schemes for energy-harvesting wireless sensor networks," *Communications Letters, IEEE*, 2012.

[35] H. Jaleel, A. Rahmani, and M. Egerstedt, "Duty cycle scheduling in dynamic sensor networks for controlling event detection probabilities," *American Control Conference (ACC)*, 2011.

[36] A. Makhoul and C. Pham, "Dynamic scheduling of cover-sets in randomly deployed wireless video sensor networks for surveillance applications," in *Wireless Days (WD), 2009 2nd IFIP*, 2009.

[37] T. Hobson and I. Clarkson, "Collaborative sensor scheduling for space situational awareness," in *Statistical Signal Processing Workshop (SSP), 2012 IEEE*, 2012.

[38] P. Kulkarni, P. Shenoy, and D. Ganesan, "Approximate initialization of camera sensor networks," in *Proceedings of the 4th European Conference on Wireless Sensor Networks*, ser. EWSN'07, 2007.

[39] H. Huang, C.-C. Ni, X. Ban, J. Gao, A. Schneider, and S. Lin, "Connected wireless camera network deployment with visibility coverage," in *INFOCOM, 2014 Proceedings IEEE*, 2014.

[40] V. P. Munishwar and N. B. Abu-Ghazaleh, "Coverage algorithms for visual sensor networks," *ACM Trans. Sen. Netw.*

[41] H.-C. Shih, H.-W. Lee, C.-L. Huang, and Y.-C. Liu, "Collaborative real-time scheduling for multiple objects tracking in ptz camera network," in *Intelligent Signal Processing and Communications Systems (ISPACS), 2012 International Symposium on*, 2012.

[42] C. Yu and G. Sharma, "Camera scheduling and energy allocation for lifetime maximization in user-centric visual sensor networks," *Image Processing, IEEE Transactions on*, 2010.

[43] B. Wang, "Coverage problems in sensor networks: A survey," *ACM Comput. Surv.*, vol. 43, no. 4, pp. 32:1–32:53, Oct. 2011. [Online]. Available: http://doi.acm.org/10.1145/1978802.1978811

[44] S. Ren, Q. Li, H. Wang, X. Chen, and X. Zhang, "Design and analysis of sensing scheduling algorithms under partial coverage for object detection in sensor networks," *Parallel and Distributed Systems, IEEE Transactions on*, 2007.

[45] V. Gupta, T. H. Chung, B. Hassibi, and R. M. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage," *Automatica*.

[46] Y. Li, C. Ai, Z. Cai, and R. Beyah, "Sensor scheduling for p-percent coverage in wireless sensor networks," *Cluster Computing*.

[47] G. H. Gonnet, "Expected length of the longest probe sequence in hash code searching," *J. ACM*, vol. 28, no. 2, pp. 289–304, Apr. 1981. [Online]. Available: http://doi.acm.org/10.1145/322248.322254

[48] M. Raab and A. Steger, ""balls into bins" - a simple and tight analysis," in *Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science*, ser. RANDOM '98. London, UK, UK: Springer-Verlag, 1998, pp. 159–170. [Online]. Available: http://dl.acm.org/citation.cfm?id=646975.711521