Quality of Monitoring of Stochastic Events by Periodic & Proportional-Share Scheduling of Sensor Coverage *

David K. Y. Yau, Nung Kwan Yip, Chris Y. T. Ma
Purdue University, West Lafayette, IN

Nageswara S. Rao, Mallikarjun Shankar
Oak Ridge National Lab, TN

ABSTRACT

We analyze the quality of monitoring (QoM) of stochastic events by a periodic sensor which monitors a point of interest (PoI) for q time every p time. We show how the amount of information captured at a PoI is affected by the proportion q/p, the time interval p over which the proportion is achieved, the event type, and the stochastic event arrival dynamics and staying times. The periodic PoI sensor schedule happens in two broad contexts. In the case of static sensors, a sensor monitoring a PoI may be periodically turned off to conserve energy, thereby extending the lifetime of the monitoring until the sensor can be recharged or replaced. In the case of mobile sensors, a sensor may move between the PoIs in a repeating visit schedule. In this case, the PoIs may vary in importance, and the scheduling objective is to distribute the sensor's coverage time in proportion to the importance levels of the PoIs. Based on our QoM analysis, we optimize a class of periodic mobile coverage schedules that can achieve such proportional sharing while maximizing the QoM of the total system.

1. INTRODUCTION

There is considerable interest in using sensors to protect populated areas against physical hazards, such as chemical, biological, nuclear, radiational, and explosive (CBNRE) leaks/attacks. Real-world sensors have limited ranges of tens to hundreds of feet. If the area to be protected is large, it may be difficult to deploy a sufficient number of sensors to cover the entire area. This leads to strong interest in the use of mobile sensors to expand the area of coverage, so that one sensor can cover multiple *points of interest* (PoIs) where interesting events may dynamically appear and disappear.

Note that there are many real-world examples of mobility in monitoring tasks. In traditional public safety work, policemen are on patrol schedules around town to detect crimes. In national security, reconnaissance planes routinely fly over mission areas to collect surveillance images, since the installation of (static) video cameras in the mission areas may be out of the question (e.g., they are foreign or enemy territories). In the case of sensor networks, certain sensors are expensive and complete area coverage by static sensors would have prohibitive costs. For example, in the Memphis Port deployment against water pollution/poisoning [10], the engineers emphasize in the project report that with the high procurement, installation, and management costs of the Smith APD2000 chemical sensors, it was not possible to have complete area coverage. They then made the difficult decision to (statically) place the sensors where the impact on people protection is highest.

There are also situations in which, independent of costs, mobility is required for the sensor network. For example, when (static) sensors are placed at PoIs where long-range data communication is difficult (e.g., underwater [1], where the high wireless signal attenuation makes it infeasible to transmit sensor data over long distances, or in an underground system of ducts where complex pathways connect the PoIs so that the placement of communication nodes to reliably get data out from underground is extremely hard), a mobile node is necessary to move between the PoIs to collect the sensor data and carry them to a data center for analysis. In this case, data may be buffered at a sensor before they are read, but the buffer capacity is limited so that unread data may be replaced by newer data and lost. Hence, the data available for reading are

*Research supported in part by SensorNet program managed by UT-Battelle, LLC for U.S. Department of Energy under contract number DE-AC05-00OR22725, in part by the U.S. National Science Foundation under grant numbers CNS-0305496 and DMS-0707926, and in part by a summer Purdue Research Foundation Fellowship awarded to C. Ma.

ACM CoNEXT 2008, December 9-12, 2008, Madrid, SPAIN

similar to stochastic events that stay for a time duration after which they will disappear.[1] The use of mobile node for data collection also has the advantage of reducing the sending energy of the sensors [3].

At the same time, it is recognized that the surveillance region may not be homogeneous, but different PoIs may vary in importance. For example, a realized but undetected threat in some parts may impact a more densely populated region than in others, as in the Memphis Port deployment. In this case, simple area of coverage is no longer sufficient. An arguably more suitable goal is to allocate sensing resources to the different parts in proportion to their importance levels. Proportional sharing of resources is not a new concept. The notion has been employed in the scheduling of CPU time, network bandwidth, etc [6, 8], where the performance impact on the rates of computation and communication has been well studied. In the case of sensor coverage, however, proportional sharing must be evaluated in terms of its impact on the *quality of monitoring* (QoM), which can be expressed as the number of interesting events captured, or the total amount of information captured about these events. The problem is not well understood.

In this paper, we target the problem of information capture about stochastic events at a set of PoIs that may have different importance levels. The events are detected by a mobile sensor which allocates its coverage time among the PoIs in proportion to the importance levels. **Our contributions** are two fold, summarized as follow.

First, we provide extensive analysis to answer the following questions as a function of the event dynamics and type of event: (1) What is the QoM of a sensor that covers a PoI for q time every p time? Does a higher proportional share of q/p imply a linearly higher QoM? (2) For the same q/p, what is the impact of the period p that controls the fairness granularity of the proportional sharing? Under what situations is finer/coarser time-scale sharing preferred over the other? (3) What is the scaling law of mobile coverage, i.e., when a mobile sensor is allocated among k out of n PoIs, how is the average QoM over all the PoIs affected as k increases? Can mobility fundamentally improve the sensing by increasing the achievable QoM?

Second, based on the QoM analysis, we will analyze the performance of a class of periodic coverage algorithms considering the travel time overhead between PoIs. We first optimize a *linear periodic* sensor schedule for maximum total QoM that achieves given proportional shares of the coverage time. We then discuss the optimization of general periodic schedules. We present a *simulated annealing* algorithm, which can within a practical time budget, find a solution close to the global optimal.

We also mention in passing that independent of mobility, the analysis of periodic sensor schedules has obvious applications in energy-efficient sensing. In this case, an energy-constrained *static* sensor may be deployed at each PoI, and there is a need to periodically turn off the sensor to conserve energy, so that the sensor will last long enough until it can be recharged or replaced. Our analysis in Section 4 gives directly the QoM of such a periodic sensor. In particular, our results show that for events that stay, the QoM of a sensor working for q/p of the time may capture a fraction of information much higher than q/p. Hence, such periodic scheduling of the sensors can be quite productive. Our results also show where it is useful (or not) to have finer granularity of the periodic scheduling, in terms of a smaller p, to achieve a higher QoM. In this case, however, it is clear that the benefits of extremely fine grained periodic scheduling may be limited/offset by the latency and energy costs of turning on/off a sensor. The development of the full details such as the energy models is out of the scope of the present paper. Also, for simplicity of exposition, we will develop our problem from the point of view of mobile coverage only in the rest of the paper.

2. PROBLEM STATEMENT

We assume that events appear and disappear at given points of interest (PoIs) and are to be monitored by a sensor of sensing range R. The PoIs are located on a 2D plane. A pair of PoIs, say i and j, are connected by a road, given by E_{ij}, of distance d_{ij}. If there is no road that directly connects i and j, $d_{ij} = \infty$. Otherwise, the sensor traveling at speed v from i to j takes time d_{ij}/v to complete the trip.

The next set of assumptions concerns the event dynamics. The events appear at PoI i one after another. After appearing, each event stays for a duration of time, which we call the *event staying time*, and then disappears. The next event appears after another duration of time, which we call the *event absent time*. We denote the sequential staying and absent times by $\{X_k^i\}_{k\geq 1}$ and $\{Y_k^i\}_{k\geq 1}$. The *event inter-arrival time* is then denoted by $Z_k^i = X_k^i + Y_k^i$. We assume that (for each i) the $\{(X_k^i, Y_k^i)\}_{k\geq 1}$ are i.i.d. random variables drawn from a common distribution (X^i, Y^i), even though the X_k^i and Y_k^i may be dependent. However, the event dynamics at different PoIs are assumed to be independent. Lastly the commonly known *event arrival times* can be re-

[1] The sensing range in Section 2 will then correspond to the communication range between the sensor and the mobile node. The event utility function there might account for the time needed for different useful fractions of the total data (about a physical world event) to be uploaded from the sensor to the mobile node.

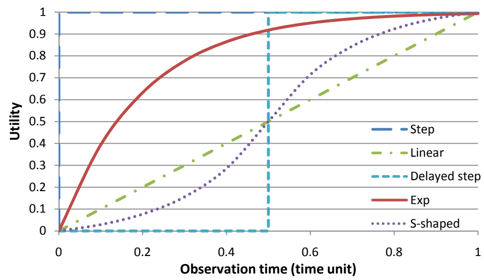

Figure 1: The utility functions.

covered by the formula: $T_0^i = 0$, $T_k^i = T_{k-1}^i + Z_k^i$ for $k, i \geq 1$ even though the T_k^i's will not be used in the analysis.

We further classify the events as follows. When the staying time drawn from X^i is always an infinitesimally small ϵ amount of time, the corresponding events are like "blips", i.e. they do not stay but disappear instantaneously after arrival. Another type of events are those which stay, i.e. there is an $0 < \epsilon \ll 1$ such that $P(X \geq \epsilon) = 1$. An event at a PoI is captured by the sensor provided that the PoI is within range of the sensor during the event's lifetime. We assume that events are *identifiable*, i.e. when the sensor that sees an event at a PoI, leaves the PoI, but comes back later to see the same event, it will know that it is the same event. We assume that as the sensor observes an event, the information it accumulates about the event increases as the observation time increases. We quantify the sensing quality as a utility function that increases monotonically from zero to one as a function of the total observation time. Fig. 1 illustrates the following five examples of the utility function:

(a) *Step function*: $U_I(x) = 1$ for $x \geq 0$. Full information about an event is obtained instantaneously on detection. (b) *Exponential function*: $U_E(x) = 1 - e^{-Ax}$. Much of the information about an event is obtained at the beginning but the marginal gain decreases as the observation time gets longer. (c) *Linear function*: $U_L(x) = Mx$ for $0 \leq x \leq \frac{1}{M}$ and $U_L(x) = 1$ for $x \geq \frac{1}{M}$. Information obtained increases linearly with the observation time until the full information is achieved. (d) *S-shaped function* $U_S(x)$. The initial observation gains little information until a critical observation time is reached, at which point there is a large marginal gain of information in a short time, and afterwards the marginal gain drops sharply as the full information is approached. (e) *Delayed step function* $U_D(x) = U_I(x - D)$. No information is gained until the total time of observation exceeds a threshold value D, after which the full information is captured instantaneously. We view (a) and (e) as extreme cases. All of the above, excepting (d), are quite amenable to analytical formulations.

When PoI i falls within the range of the sensor, we say that the sensor is *present* at i. Otherwise, the sensor is *absent* from i. Since we are interested in the resource competition between different PoIs, we make the following assumption.

Assumption 1. *The PoIs and the roads between them are separated such that (1) no two PoIs fall within the range of the sensor at the same time; (2) for the sensor traveling from PoI i to PoI j on E_{ij} at speed v, i will be within range of the sensor for R/v time before the sensor leaves i, and j will be within range of the sensor for R/v time until the sensor reaches j, and (3) no PoI other than i and j falls within the range of the sensor during the trip on E_{ij}. In general, however, the sensor can vary its speed while traveling on a road.*

2.1 Definition of QoM

We now define the quantitative measurement of the QoM at a PoI or for the whole protected area. In the course of a deployment, denote by $e_1^i, \ldots, e_{m_i}^i$ the sequence of events appearing at PoI i over the duration $[0, T]$ of the deployment. For the event e_j^i, assume that it is within range of the sensor for a total (but not necessarily contiguous) amount of time t_j^i, where $t_j^i \geq 0$. The sensor will then gain a certain amount of information, $U_j^i(t_j^i)$, about e_j^i, where $U_j^i(\cdot)$ is the utility function of e_j^i. The total information gained by the sensor at i is defined by $E_i(T) = \sum_{1 \leq j \leq m_i} U_j^i(t_j^i)$, and the average information gained per event at i during the whole deployment period is then $\bar{E}_i(T) = E_i(T)/m_i$. Similarly, the total information gained by the sensor in the whole deployment is $E_*(T) = \sum_{1 \leq i \leq n} E_i(T)$, where n is the number of PoIs in the protected area. The average information gained per event in the whole deployment is then

$$\bar{E}_*(T) = \left(\sum_{1 \leq i \leq n} m_i \bar{E}_i(T) \right) / (\sum_{1 \leq i \leq n} m_i).$$

By means of the strong law of large numbers and renewal theory, $\bar{E}_i(T)$ and $\bar{E}_*(T)$ will converge to a deterministic number as $T \longrightarrow \infty$. Hence we define the QoM of PoI i and the whole covered area as:

$$Q_i = \lim_{T \to \infty} \bar{E}_i(T), \text{ and } Q_* = \lim_{T \to \infty} \bar{E}_*(T). \tag{1}$$

Furthermore, they are related by:

$$Q_* = \frac{1}{\mu_*} \sum_{1 \leq i \leq n} \mu_i Q_i, \tag{2}$$

where $\mu_i = \frac{1}{E(Z)}$ is the mean event arrival rate at PoI i and $\mu_* = \sum_{1 \leq i \leq n} \mu_i$.

Note that in defining the QoM, we should in principle divide not by the number of events m, but by the maximum possible utility achievable for an event: $\int_0^\infty U(x) f(x)\, dx$, where $f(x)$ is the pdf of the event staying time distribution. The latter may be less than 1 if the events do not stay infinitely long. However, the difference is by a proportionality constant only, and will not affect our comparison results. Unless otherwise stated, we will further assume that all the events at i have the same utility function, and denote this function by $U^i(\cdot)$.

3. RELATED WORK

Area coverage in a sensor network has been well studied [4, 12]. Protocols have been proposed to task subsets of sensors in a dense network to provide maximum lifetime area coverage [15]. Simple area coverage does not consider the varying importance of different sub-regions. Our work addresses the heterogeneity of sub-regions by proportional-share coverage. Proportional-share resource allocation has been proposed for CPU/network scheduling [6, 8]. Mobile coverage has the additional challenge that sensor schedules can be severely constrained by the adjacencies and distances between the PoIs.

The importance of the sensing time in accurately assessing various physical phenomena has been well documented [9]. The need for non-negligible sensing durations to obtain useful information is due to noises in the measurement process and the probabilistic nature of the phenomena under observation. The impact of the sensing time is captured by the event utility functions in our problem statement.

Mobility has been discussed in delay-tolerant, vehicular, and sensor networks. Passive mobility has been analyzed for its effects on providing communication opportunities [5, 16]. Mobility control has been used to deploy ferries and data mules among a number of data sources, to optimize communication of the source data to the data sink [13, 17].

The dynamics of real-world events are frequently modeled as stochastic processes. Poisson arrivals are generally accurate characterizations of a large number of independent event occurrences, whose event inter-arrival times are Exponentially distributed. Real-world network/computing workloads have properties that are found to be long-range dependent [11], which follow the Pareto distribution. In a sensor network, the target events may have similar dynamic behaviors. For example, radioactive particles arriving at a Geiger-Müller counter follow a Poisson process [9]; a chemical leak at a facility may occur with a probability, and the leak may persist for a random duration until the chemical has been dispersed.

The sensing of stochastic events by a mobile sensor has been studied in [1]. Our problem in this paper is quite different. First, we consider differential coverage of PoIs by proportional sharing whereas they do not. In particular, we analyze the QoM of periodic sensor schedules, as a function of the proportional share q/p and the period p. Such analysis has applications besides mobile coverage, e.g., energy-efficient sensing by periodically turning off a sensor. Second, we consider sensing tasks with the temporal dimension as defined by the event utility function, whereas they focus on the number of captured events only, where an event is captured whenever it falls within the sensing range of a sensor, no matter how brief the sensing time. Third, we define the concepts of linear and general periodic schedules among the PoIs, and design optimal algorithms for both kinds of schedule.

4. SINGLE-POI ANALYSIS OF QOM

We explain the impact on the QoM by the coverage schedule of a sensor at a given PoI. The schedule specifies the time intervals over which the sensor is present at or absent from the PoI. A given schedule is achieved by how the sensor moves between the PoIs according to some movement algorithm. The problem of the algorithm design and the feasibility of a set of PoI schedules are the subject of Section 5.

We can already illustrate some interesting QoM properties of proportional-share mobile coverage by considering only *periodic* schedules at individual PoIs. Specifically, we assume that the sensor is alternately present and absent at a PoI, say i, for q_i and $p_i - q_i$ time units, respectively. For example, let S_1 be the following the coverage schedule of i:

$$S_1 = \{PAAAPAAA\ldots\}$$

for $q_i = 1$ and $p_i = 4$. In the schedule, P denotes one time unit of the sensor's presence and A denotes one time unit of the sensor's absence. Thus the proportional share equals $q_i/p_i = 25\%$ of the sensor's total coverage time.

Clearly, a given proportional share for i can be achieved in many different ways. For example, $q_i = 2$ and $p_i = 8$ give the following schedule S_2 with the same 25% share for i:

$$S_2 = \{PPAAAAAAPPAAAAAA\ldots\}.$$

While S_1 and S_2 are equivalent from the proportional-share point of view, they differ in terms of the time scale over which the proportional share is achieved. Specifically, S_1 achieves the 25% share over a time period of 4 time units, whereas S_2 achieves the same share over a period of 8 time units. We say that S_1 has a finer *fairness granularity* than S_2, and will use p_i to quantify this fairness granularity. Notice that for a fixed proportional share, a smaller p_i implies a proportionately smaller q_i.

The main purpose of this section is to analyze the dependence of the QoM on the utility function and the fairness granularity. In this section, as we will focus on a single PoI, the subscript i will be omitted where there is no confusion. We will frequently denote the proportional share $\frac{q}{p}$ by γ. For simplicity, we use $P_j = [(j-1)p, (j-1)p + q]$ and $A_j = [(j-1)p + q, jp]$ to denote the j-th sensor present and absent periods, respectively. For many of the proofs, it is sufficient to consider just the case $j = 1$, i.e. $P_1 = [0, q]$ and $A_1 = [q, p]$.

The problem as formulated in Section 2 fits perfectly well in the realm of renewal theory. One of its main conclusions is that in the long run, the expected number of arrivals in an interval dt equals $\mu\, dt$ where $\mu = 1/E(Z)$.

The following two types of event staying time distribution will be considered in this paper, where $f(x)$ is the pdf of X:

- Exponential Distribution ($\lambda > 0$):
$$f(x) = \lambda e^{-\lambda x}, \quad x > 0, \quad \text{mean} = \frac{1}{\lambda}.$$
- Pareto Distribution ($\alpha, \beta > 0$):
$$f(x) = \frac{\alpha\beta^{\alpha}}{x^{\alpha+1}}, \quad x > \beta, \quad \text{mean} = \frac{\alpha\beta}{\alpha - 1}.$$

Now we proceed to present our results. All of the proofs will only be outlined due to space constraints, but can be made fully rigorous.

4.1 Step utility function

We begin our discussion with events that have the step utility function (see Fig. 1). In this case, since the utility reaches one instantaneously, the QoM is equivalent to the fraction of events captured. The next result illustrates the effect on the QoM by a periodic sensor schedule with parameters p and q at a fixed PoI.

THEOREM 1. *For independent arrivals of events that have the step utility function and do not stay, i.e. ,"blip events", the QoM at any PoI is directly proportional to its share of coverage time q/p. In particular, the achieved QoM does not depend on the fairness granularity.*

PROOF. The statement is a consequence of the fact that an event is completely captured if and only if it arrives when the sensor is present. Hence the QoM is simply the ratio between the expected number (per unit time) of arrivals during the sensor present period and the total period. □

The above scenario shows that only the proportional sharing information determines the QoM. On the other hand, for events that do stay, the QoM depends on the relationship between the event staying time distribution and the parameters p and q. Specifically, we have the following result.

THEOREM 2. *For independent arrivals of events that stay and have the step utility function, the QoM at a PoI is given by*

$$Q = \frac{q}{p} + \frac{1}{p}\int_0^{p-q} \Pr(X \geq t)\, dt. \qquad (3)$$

PROOF. As the utility function is a step function, the overall utility is given by the total number of events captured when the sensor is present. Note that an event will be captured if (a) it arrives during the sensor present period $[0, q]$; (b) it arrives during the sensor absent period $[q, p]$, but stays long enough to be captured during the *next* sensor present period $[p, p + q]$. The contribution of (a) to the QoM is given by $\frac{q}{p}$, while that of (b) is given by $\frac{1}{p}\int_q^p \Pr(X + t \geq p)\, dt$, which is the second term of Equation 3 after a simple change of variable. □

Theorem 2 implies that the sensor that stays at a PoI for q/p of the time may be able to capture a significantly *larger* fraction of events than q/p. The following two corollaries give further statements due to this extra fraction of events.

COROLLARY 1. *Under the setting of Theorem 2, with the fairness granularity p kept constant, we have:*

$$\lim_{\gamma \to 0} Q = \frac{1}{p}\int_0^p \Pr(X \geq t)\, dt.$$

PROOF. The proof is a direct consequence of Equation (3), upon taking the limit $\gamma \longrightarrow 0$. □

This result clearly indicates that no matter how small the proportional share is, there is always some definite, positive gain of information. This is due to the fact that the events stay.

COROLLARY 2. *Under the setting of Theorem 2, the QoM of a given fixed proportional share is a monotonically decreasing function of the fairness granularity, i.e., Q decreases as p increases. Furthermore,*

$$\lim_{p \to 0} Q(p) = 1, \quad \textit{and} \quad \lim_{p \to \infty} Q(p) = \frac{q}{p}.$$

PROOF. The statement again is a simple consequence of Equation (3) which is re-written in the following form:

$$Q = \gamma + (1 - \gamma)\frac{1}{(1 - \gamma)p}\int_0^{(1-\gamma)p} \Pr(X \geq t)\, dt.$$

where $\frac{q}{p} = \gamma$. The results follow by taking the corresponding limits. □

In contrast to Theorem 1 for blip events, Corollary 2 implies that finer-grained fairness *does* generally improve the QoM for staying events having Step utility. In particular, no matter how small the proportional share is, an arbitrarily high QoM can be achieved by an extremely fine fairness granularity.

We now consider a scaling result for mobile sensor coverage among k out of n PoIs, whose event arrival and departure processes are i.i.d., as k increases. Assume that initially, the sensor achieves periodic schedules among k of the n PoIs such that $q_i = \delta$ and $p_i = k\delta$, for $1 \leq i \leq k$, where δ is a unit of time. The following theorem holds.

THEOREM 3. *The expected fraction of events captured is an increasing function of k, the number of PoIs covered.*

Proof. The expected fraction of the events captured in the schedule is

$$Q_* = \frac{1}{n}\sum_{1\le j\le k}\left[\frac{1}{k}+\frac{1}{k\delta}\int_0^{(k-1)\delta} P(X\ge t)\,dt\right] = \frac{1}{n}\left[1+\frac{1}{\delta}\int_0^{(k-1)\delta} P(X\ge t)\,dt\right]$$

which is clearly an increasing function of k. □

Theorem 3 provides a formal justification for mobile coverage, namely that the amount of information captured increases as the sensor moves among more PoIs to search for interesting information.

4.2 General utility function

We now turn our attention to events that have a general utility function $U(\cdot)$. In this case, we have the following QoM result.

Theorem 4. *For independent arrivals of events at a PoI that have the utility function $U(\cdot)$ and whose event staying time pdf is given by $f(x)$, the achieved QoM equals ($\xi_i = iq - t$, $\eta_i = x + ip - t$):*

$$\int_0^q \left[\int_0^{q-t} U(x)f(x)\,dx + \sum_{i=1}^{\infty}\int_0^q U(\xi_i + x)f(\eta_i)\,dx + \sum_{i=1}^{\infty} U(\xi_i)\int_{-(p-q)}^0 f(\eta_i)\,dx\right] dt \quad (4)$$

$$+ \int_q^p \left[\sum_{i=1}^{\infty}\int_0^q U(\xi_i - q)f(\eta_i)\,dx + \sum_{i=1}^{\infty} U(\xi_i + t)\int_q^p f(\eta_i)\,dx\right] dt. \quad (5)$$

Proof. The above formula follows from the fact that the overall utility available for any particular event depends on the *total* length of the intersecting region (which might be discontinuous) during which both the event and sensor are present. The various summands in integral (4) and (5) correspond to the cases that the event arrives when the sensor is present or absent. □

The formula above can have a complicated analytical form in general, but it is certainly amenable to numerical computation. Nevertheless, we first present two exact analytical results. (Recall $\gamma = \frac{q}{p}$.)

(1) Exponential utility function U_E and Exponential staying time: $f(x) = \lambda e^{-\lambda x}$.

$$\begin{aligned} Q &= \frac{A\gamma}{A+\lambda} - \frac{1-e^{-\lambda q}}{\lambda p} + \frac{\lambda(1-e^{-(A+\lambda)q})}{(A+\lambda)^2 p} \\ &+ \frac{(e^{\lambda q}-1)^2}{\lambda p e^{\lambda q}(e^{\lambda p}-1)} - \frac{\lambda(e^{(A+\lambda)q}-1)^2 e^{-(A+\lambda)q}}{(A+\lambda)^2 p(e^{(Aq+\lambda p)}-1)} \\ &+ \frac{2(e^{\lambda(p-q)}-1)}{p} \\ &\times \left[\frac{e^{\lambda q}-1}{\lambda(e^{\lambda p}-1)} - \frac{e^{(A+\lambda)q}-1}{(A+\lambda)(e^{(Aq+\lambda p)}-1)}\right] \\ &+ \frac{(e^{Aq}-1)e^{\lambda q}(e^{\lambda(p-q)}-1)^2}{\lambda p(e^{\lambda p}-1)(e^{(Aq+\lambda p}-1)}. \end{aligned} \quad (6)$$

Note that the above leads to

$$\lim_{p\to 0} Q = \frac{A\gamma}{A\gamma+\lambda}, \quad \lim_{p\to\infty} Q = \frac{A\gamma}{A+\lambda}. \quad (7)$$

(2) Delayed utility function U_D and Exponential staying time: $f(x) = \lambda e^{-\lambda x}$.

When p is very small such that D is an integral multiple of q, i.e. $D = kq$ for $k = 1, 2, \ldots$, we have:

$$Q = e^{-\frac{\lambda D}{\gamma}}\left[\gamma + \frac{e^{\lambda(1-\gamma)p}-1}{\lambda p}\right]. \quad (8)$$

On the other hand, when p is very large, specifically, when $q > D$, then

$$Q = e^{-\lambda D}\left[\gamma + \left(\frac{1}{\lambda} - D\right)\frac{1-e^{-\lambda(p-q)}}{p}\right]. \quad (9)$$

Combining Equations (8) and (9), we have:

$$\lim_{p\to 0} Q = e^{-\lambda\frac{D}{\gamma}}, \quad \lim_{p\to\infty} Q = \gamma e^{-\lambda D}. \quad (10)$$

The above analytical results can be intuitively understood in many ways, which are instructive to discuss.

4.3 Implications of theoretical results

The first three discussion points concern various limiting cases.

(i) Let the fairness granularity p and the proportional share γ be fixed. Then as the event staying time goes to infinity, every event will always be captured and the maximum value 1 for the utility can be achieved. Furthermore, the QoM is an increasing function of the mean event staying time. Note that this scenario corresponds to $\lambda \longrightarrow 0$ for the exponential staying time distribution, and $\beta \longrightarrow \infty$ for the Pareto distribution.

(ii) In the limit of $p \longrightarrow 0$, every event which stays will always be captured. However, the total observation time is only γ fraction of the event's duration. Hence the average utility achieved is:

$$Q_0 = \int_0^\infty U(\gamma x)f(x)\,dx. \quad (11)$$

This result is consistent with the explicit results (7) and (10).

(iii) In the limit of $p \longrightarrow \infty$, each event, if captured, will essentially be observed for its whole duration. On the other hand, only γ fraction of the events will be captured. Hence the QoM is given by:

$$Q_\infty = \gamma\int_0^\infty U(x)f(x)\,dx, \quad (12)$$

which is also consistent with the explicit results (7) and (10).

The next two discussion points concern the two most important qualitative descriptions of the QoM function.

(iv) For the step and exponential utility functions, the QoMs are monotonically decreasing functions of p. This is because both utility functions are *concave* functions of the observation time. Hence it is advantageous to capture as many *new* events as possible rather than to gain information for the same event. A finer fairness granularity exactly achieves this. (This is consistent with Theorem 2 and the analytical formula (6).)

(v) However, the key feature is that for certain utility functions, the *maximum* QoM is only achieved at some *optimal* fairness granularity. We spend a moment to explain this important phenomenon.

This observation is easiest to explain for the delayed step utility U_D. In the limit of $p \longrightarrow 0$, any event can always be captured. This is essentially the statement of Corollary 2. However, in order to gain enough information about the event, it is necessary that the event staying time be at least $\frac{D}{\gamma}$ long. This probability is given by $\Pr(X \geq \frac{D}{\gamma})$. However, when p is positive (no matter how small it is), this is not absolutely necessary. In fact, if the event arrives right at the beginning of a sensor present period, then the event staying time just needs to be at least $\frac{D}{\gamma} - (1-\gamma)p$ long. It is this saving that increases the QoM. Hence initially, the QoM is an *increasing* function of p for *small* p. (This can also be seen analytically from Equation (8).)

The behavior of QoM when p is *large* is also interesting and quite intricate. From Equation (9), observe that the QoM is a *decreasing*, *constant*, or *increasing* function of p for λ *less than*, *equal to*, or *greater than* $\frac{1}{D}$, respectively. This is due to the competitive effect (for p large) of the *loss* of utility for events arriving near the end of a sensor present period and the *gain* of utility for events arriving before the sensor present period. Hence for $\lambda < \frac{1}{D}$, the QoM initially increases and then decreases as a function of p. Thus it is optimal at some *intermediate* p value.

All of the above implications are supported by the simulation results in Section 6.

5. COVERAGE ALGORITHMS

The previous section discussed the QoM of periodic schedules at a specific single PoI. We now address the problem of covering n PoIs by the sensor. This is achieved by a visit schedule of the sensor to all the PoIs under a coverage algorithm to be designed.

We will analyze the QoM of *periodic* coverage of n PoIs. By this we mean that the schedule is realized by a periodic visit schedule of the sensor to the PoIs, in which the visit schedule in the smallest period is denoted by

$$S = \{(L_1, C_1), \ldots, (L_m, C_m)\}, \qquad (13)$$

where L_j denotes the jth PoI visited for a time of C_j in the sensor schedule, $L_j \neq L_{(j \bmod m)+1}$, and each of the n distinct PoIs appears at least once in S. Recall from Assumption 1 on Page that the sensor cannot be present at more than one PoI at a time. If $m = n$, each PoI appears in S exactly once, then we call S a *linear periodic schedule*. However, it is clear that not all periodic schedules are linear. For example, $S = \{(1,\delta), (2, 3\delta), (1,\delta), (3, 2\delta)\}$, where δ is a unit of time, is not. In the definition (13), if $m > n$, we call the periodic schedule *nonlinear*. We restrict our attention to periodic schedules for now.

Given a sensor schedule S, we define its *maximum feasible utilization* as

$$U_*(S) = \sup \sum_{1 \leq i \leq n} \frac{q_i}{p_i},$$

where the sup is taken over all possible sensor movements that realize S. The utilization is affected by the travel time overhead between two adjacent PoIs in S during which the sensor is not present at any PoI. Using $d(i,j)$ as an equivalent notation to d_{ij} for the distance between i and j, we define for $j = 1, \ldots, m$:

$$a_j = \frac{1}{v_{\max}} \Big[d(L_j, L_{(j \bmod m)+1}) - 2R \Big]$$

as the minimum travel time overhead from L_j to $L_{(j \bmod m)+1}$ for the sensor moving at maximum speed $v_{\max}$. Then the following statement holds.

THEOREM 5. *(For linear periodic schedule), the maximum feasible utilization of S is*

$$U_*(S) = \sup \left[1 - \frac{\sum_{1 \leq j \leq m} a_j}{\sum_{1 \leq j \leq m} (C_j + a_j)} \right],$$

where the sup *is taken over all possible sensor movements realizing S.*

PROOF. Completing one period of the sensor schedule requires $P_* = \sum_{1 \leq j \leq m}(C_j + a_j)$ time units. Hence the proportional share for PoI i is given by $\frac{C_j}{P_*}$. The result thus follows from: $\sum_j \frac{q_j}{p_j} = \sum_j \frac{C_j}{P_*} = 1 - \frac{1}{P_*} \sum_j a_j$. $\square$

Theorem 5 shows that 100% sensor utilization is feasible if and only if each adjacent pair of PoIs in S are exactly $2R$ apart. In actual application, we would like to maximize $U_*(S)$. As its form is a decreasing function of the sum $\sum_{i \leq j \leq m} a_j$, we would indeed want to minimize the travel overhead.

5.1 Linear periodic schedule optimization

Here we discuss the optimization of the QoM Q_* (defined in Section 2) for the overall system in the realm of linear periodic schedules. The solution must satisfy a given proportional fairness objective, i.e., for each pair of PoIs, say i and j, we must achieve a given ratio, γ_{ij}, of their shares of coverage time. I.e., for the periodic schedules induced by S at i and j, we have $\frac{q_i/p_i}{q_j/p_j} = \gamma_{ij}$.

A linear periodic schedule exists if there is a Hamiltonian circuit of the PoIs. An optimization approach for linear periodic schedules works as follows. We first determine the visit order of the PoIs in S that will minimize $\sum_{1\le j\le m} a_j$. The problem is the Traveling Salesman Problem and is NP hard, but practical approaches exist that give solutions within a few percent of the optimal for problem sizes of up to 100,000 [7]. Once the visit order is determined, a_j, $j = 1,\ldots,m$, is known, and it remains to determine the $C_j, j = 1,\ldots,m$. Notice that in a linear periodic schedule, $m = n$, $C_j = q_j$, and $p_1 = \ldots = p_n = \sum_j (C_j + a_j) = P_*$. We first select each C_j to satisfy $C_j = \gamma_{j1} C_1$ so that all the coverage times can be expressed in terms of C_1 only. This greatly simplifies the problem as it becomes a purely one-dimensional optimization problem. The choice of C_1 that optimizes Q_* depends on the event utility function U.

We illustrate the above approach by a simple example. Consider first blip events and the step utility function U_I, If $\sum_j a_j = 0$, then any choice of C_1 is optimal as the QoM is simply the fraction of events captured at the PoIs. More precisely,

$$Q_* = \frac{1}{nP_*}\sum_j C_j = \frac{1}{n}.$$

On the other hand, if $\sum_j a_j > 0$, then the optimal Q_* cannot be attained but it can be approached as close as possible by using a finite but sufficiently large value of C_1.

For general event utility functions, we need to compute the corresponding QoM Q_i for each i using Theorem 4. Recall that $C_i = \gamma_{i1} C_1$, and Q_* is expressible as a weighted sum of the individual Q_i's (from Equation 2):

$$Q_* = \frac{1}{\mu_*}\sum_j \mu_j Q_i\left(\frac{\gamma_{j1} C_1}{P_*}\right).$$

Therefore Q_* is a function of C_1 only. The value of C_1 that optimizes QoM Q_* can be computed by solving

$$\frac{dQ_*}{dC_1} = 0, \quad \text{and} \quad \frac{d^2 Q_*}{dC_1^2} < 0.$$

Note that Q_* can possibly have multiple local maxima as each Q_i has its own optimal C_i's. But the issue can be easily resolved by a numerical search since the problem is one-dimensional.

5.2 General periodic coverage

The previous section discussed the optimization of linear periodic sensor schedules. However, a linear periodic schedule does not exist if there is no Hamiltonian circuit of the PoIs. Even if it exists, a linear schedule is in general sub-optimal as the QoM depends on the fairness granularity (Corollary 2). This is illustrated by the following example. Consider three PoIs, located such that $d_{12} = d_{13} = d_{23} = 2R$, and the proportional fairness objective of $\gamma_{12} = n/(n-1)$ and $\gamma_{13} = n$. For events that stay and have the step utility function, the optimal linear periodic sensor schedule is $\{(1, n\delta), (2, (n-1)\delta), (3, \delta)\}$, where $\delta = 2R/v_{\max}$ is the minimum presence time of the sensor arriving at and then leaving a PoI — recall Assumption 1. From Theorem 2, however, we know that the QoM at i increases as the fairness granularity decreases. Hence, the optimal non-linear periodic schedule $\{\underbrace{(1,\delta),(2,\delta),\ldots,(1,\delta),(2,\delta),}_{n-1 \text{ times}}(1,\delta),(3,\delta)\}$ increases the QoM at 1 and 2 without affecting either the travel overhead or the QoM at 3. When n is large, the performance loss of the optimal linear schedule can be significant for certain distributions of the event staying time, e.g., when the mean event staying time is on the order of δ.

The above argues for the need to search for general periodic schedules with better performance. A beginning observation is that a new and potentially better periodic schedule can be obtained by rearranging the PoI order in an original schedule. Changing the PoI order affects the fairness granularity as discussed above, but it also affects the travel overhead between the adjacent PoIs visited. Since the travel time overhead is known given a PoI visit order, the achieved Q_* measure of the new schedule can be computed by applying Theorem 4 with a modification for non-linear periodic schedules.

For the case of the step utility function U_I, the QoM is in fact simply a weighted sum of the QoMs for the linear periodic sub-schedules which constitute the whole schedule (see the next Theorem). For simplicity, we ignore the travel overhead (which can be easily incorporated). To set up the notation, for a general periodic schedule, let $\{p_k^i - q_k^i, q_k^i\}_{1\le k\le K_i}$ be the consecutive sensor absent and present times for PoI i. Note that $p_* = \sum_{1\le k\le K_i} p_k^i$ is the total period of the schedule (which is the same for all i's). Then we have the following result.

THEOREM 6 (STEP UTILITY FUNCTION). *The QoM of PoI i is given by*

$$Q_i = \sum_{k=1}^{K_i} \frac{p_k^i}{p_*}\left[\frac{q_k^i}{p_k^i} + \frac{1}{p_k^i}\int_0^{p_k^i - q_k^i} \Pr(X \ge t)\, dt\right].$$

In particular, the QoM is a linear combination of the QoM of each individual sub-linear periodic schedules which constitute the overall nonlinear periodic schedule.

PROOF. The proof follows the same line as Theorem 2. The key observation that makes the proof go through is that if an event arriving during the absent period $p_k^i - q_k^i$ is ever captured, then it must be first captured in the next present period q_k^i. □

5.3 General periodic schedule optimization

Simulated Annealing Algorithm

```
best = s = initial periodic schedule
Qbest = Qs = QoM(best)
for (i = 0; i < computation_budget; i++)
    p1, p2 = random positions in s
        subjected to selection criteria
    new = s with p1, p2 swapped
    if (new is physically infeasible)
        continue
    Qnew = QoM(new)
    if (Qnew >= Qs)
        s = new, Qs = Qnew
        if (Qnew > Qbest)
            best = new, Qbest = Qnew
    else    // simulated annealing
        if (random < exp((Qnew - Qs) * i))
            s = new, Qs = Qnew
return best
```

Figure 2: Simulated annealing algorithm for optimal periodic schedule.

We now illustrate how the above Theorem is used to optimize a general periodic schedule for Step utility. Starting with any initial periodic schedule of length n, there are $n!$ straightforward permutations of the schedule to obtain a general periodic schedule. An exhaustive search for an optimal schedule is computationally infeasible for large n. To overcome the challenge, we use a simulated annealing algorithm to search for a general nonlinear periodic schedule with its Q_* value as close to the optimal as possible. (The use of simulated annealing for optimization problems was presented in [7]. See also [2] for an application to the traveling salesmen problem.) To apply simulated annealing in our problem domain, care must be taken to consider the physical constraints of mobility including the finite speed of the sensor and the adjacencies of the PoIs.

The optimization algorithm is specified in Fig. 2. We initialize the current search candidate s to some initial periodic schedule, and keep track of the current best schedule *best* seen so far. We then randomly select two elements in s, say (L_i, C_i) and (L_j, C_j), and swap $k_i\delta$ cover time from C_i with $k_j\delta$ time from C_j, to obtain a new schedule denoted by *new*, where $\delta = 2R/v_{\max}$, k_i and k_j are randomly selected positive integers such that $k_i\delta \le C_i$ and $k_j\delta \le C_j$. To avoid a cover time of less than δ for any element, we have the additional rule that any fractional δ time left by itself after a swap will be moved together with the associated whole number multiple of δ time moved. If two adjacent PoIs in *new* have distance ∞ between them, *new* is rejected as physically infeasible.

For general utility functions, the closed analytical form of the QoM for a general (non-linear) periodic schedule can be complicated. In particular, it will not be a weighted sum of the QoMs of the linear periodic sub-schedules. Nevertheless, in order to apply the simulated annealing, one can still write down an analytical formula for the QoM (Theorem 7) and resort to numerical integration to compute its value.

THEOREM 7 (GENERAL UTILITY FUNCTION). *Let U^i be the utility function of the events at PoI i and $f(x)$ be the pdf of the event staying time. Then*

$$Q_i = \frac{1}{p_*}\int_0^{p_*}\int_0^{\infty}[U]^i(t,x)f(x)\,dx\,dt,$$

where $[U]^i(t,x) = U^i\left(\int_t^{t+x} p^i(s)\,ds\right)$ and $p^i(s)$ is a function which takes the value 1 when the sensor is present at PoI i at time s, and 0 otherwise.

PROOF. The proof is the same as Theorem 4 with the following understanding. The variable t refers to the event arrival time, x refers to the event staying time, and $\int_x^{x+t} p^i(s)\,ds$ is the total time the event is observed by the sensor. $\square$

Note that by increasing the duration of the optimization period, the algorithm will optimize over an increasingly larger set of the candidate schedules, which can be quite general when the period is sufficiently long.

6. SIMULATION RESULTS

6.1 Single-PoI QoM

We present simulation results to illustrate the analytical results in Section 4. Recall the use of X and Y to denote the event staying and absent time variables, respectively. We measure the QoM Q_i achieved over 1,000,000 time units in a simulation run, and report the average Q_i of 10 different runs. The different runs produce results that have extremely small differences. Hence, we omit the error bars in the reported results. Note that not all the events in a simulation stay long enough to be captured at the full utility. The maximum information available for capture is given by $\int_0^\infty U(x)f(x)\,dx$ as explained in Section 2.1. Each reported experiment uses the same distribution for both the event staying and absent times, which is either Exponential with varying λ, or Pareto with varying β (and α is kept to be 2).

6.1.1 Step utility

We first present results for the Step utility function U_I and Exponential event dynamics.[2] Figure 3(a) shows the QoM as a function of the proportional share q/p. The results agree with Theorem 2 and its instantiations for the distribution. Note that the fraction of events captured can be significantly higher than the proportional share, e.g., a QoM of close to 0.4 is achieved for $\text{Exp}(\lambda = 0.25)$ even when the share is only slightly positive (see Corollary 1). The observation time of the events increases as the events stay longer, and so the QoM is higher when

[2]Results for Pareto event dynamics support similar conclusions and can be found in [14].

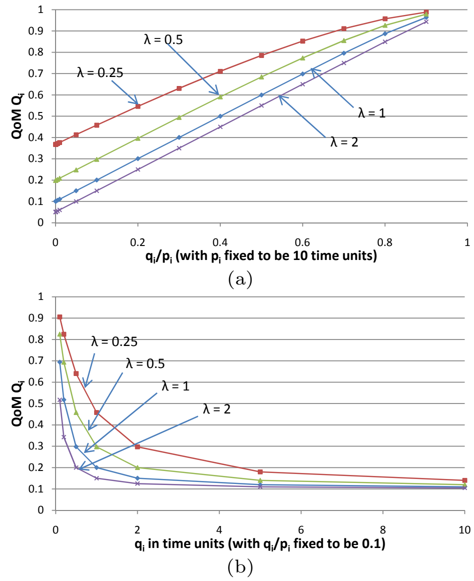

(a)

(b)

Figure 3: Achieved QoM for events that stay and have the Step utility function U_I; $X \sim Y \sim$ Exp(λ).

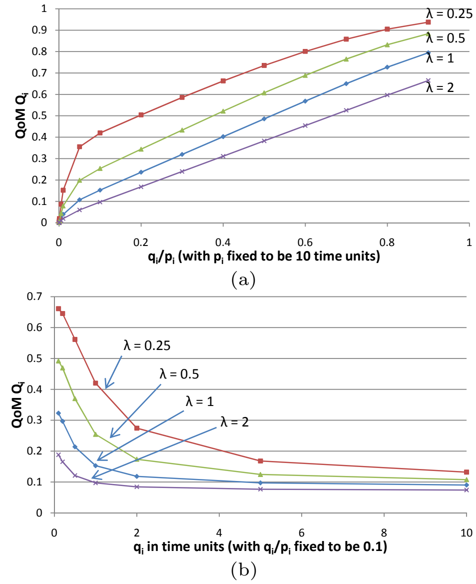

(a)

(b)

Figure 4: Achieved QoM for events that stay and have the Exponential utility function U_E, $A = 5$; $X \sim Y \sim$ Exp(λ).

λ is smaller for the Exponential event dynamics (see Sec. 4.3(i)). In general, the QoM is not linear in the proportional share.

Figure 3(b) shows the QoM as a function of the fairness granularity. As predicted by Corollary 2, the QoM is a monotonically decreasing function of p, meaning that finer grained fairness will improve performance. As explained before, the QoM increases as λ decreases. Furthermore, the QoM converges to the maximum value one and the proportional share $\gamma = q/p$ as p converges to 0 and ∞ (see Corollary 2).

6.1.2 Exponential utility

We now present results for the Exponential utility function U_E (with $A = 5$). Figure 4(a) shows the achieved QoM as a function of the proportional share for Exponential event dynamics. Unlike Step utility, the achieved QoM is close to zero when the share is only slightly positive. This is due to the need to accumulate information for Exponential utility. As the share increases initially, however, there is a sharp gain in the QoM. This is because most information is gained during the initial observation of an event for Exponential utility. Moreover, the initial gain is higher when the events stay longer (i.e., smaller λ), because longer staying events are more likely to be captured even if they arrive when the sensor is not present. As the share further increases, the marginal gain in the QoM becomes smaller, again mimicking the decreasing marginal gain of information with longer observation time for the type of event. Note that for the larger λ values (e.g., $\lambda = 2$), the QoM is significantly smaller than one even for a large share. This is in part because at those parameter values, the events do not stay long enough to be captured at their full utility.

Figure 4(b) shows the achieved QoM as a function of the fairness granularity. For the Exponential utility, the results agree with Equations 6 and 7. (See also Section 4.3(iv).) In particular, it shows that the QoM is monotonically decreasing in p and gives the correct QoM limits in Eq. (7) as $p \to 0$ and $p \to \infty$. In addition, the QoM increases when λ decreases.

6.1.3 Delayed Step utility

We now present simulation results for the Delayed Step utility U_D ($D = 0.5$ time units) with Exponential and Pareto event staying time distributions. Similar results hold for the S-shaped utility but they are not shown due to space.

Figures 5(a) and 5(b) show the achieved QoM as a function of the proportional share for the Exponential and Pareto event dynamics, respectively. They show that the QoM is monotonically increasing in the proportional share, and the QoM is higher when the events stay longer (i.e., smaller λ or larger β).

Figures 5(c) and 5(d) show the achieved QoM as a function of the fairness granularity. Note that in this case, the QoM is no longer monotonically decreasing in p, but the optimal fairness occurs at an intermediate value. Note also that for $\lambda = 2 = \frac{1}{D}$, the QoM is a constant function of p for large p. These properties are all discussed in Section 4.3(v).

6.2 General periodic coverage optimization

We present simulation results to illustrate the performance of the optimization algorithm in Section 5 for periodic schedules. We use 3 PoIs, de-

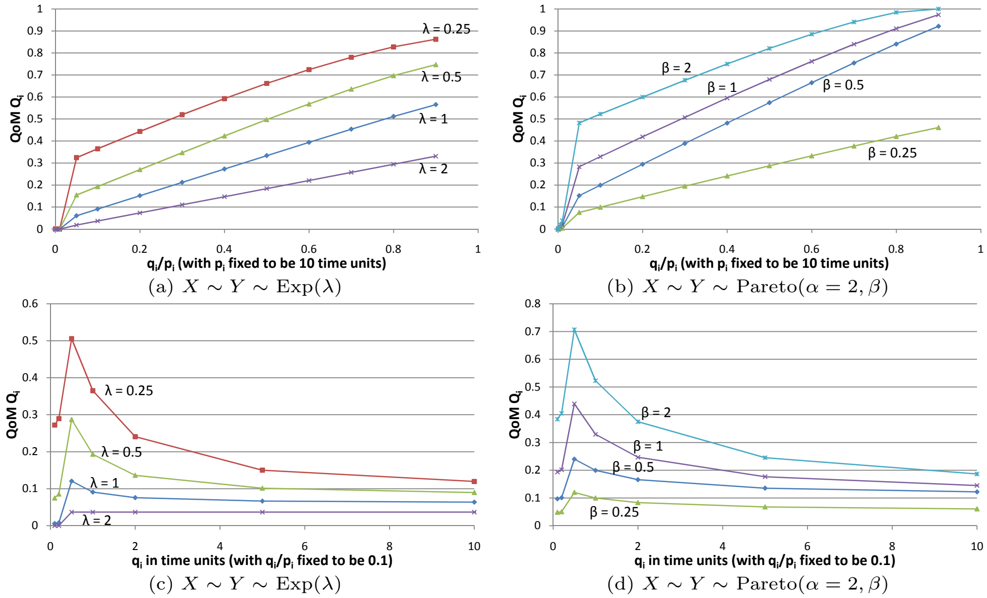

(a) $X \sim Y \sim \text{Exp}(\lambda)$ (b) $X \sim Y \sim \text{Pareto}(\alpha = 2, \beta)$

(c) $X \sim Y \sim \text{Exp}(\lambda)$ (d) $X \sim Y \sim \text{Pareto}(\alpha = 2, \beta)$

Figure 5: Achieved QoM for Delayed Step utility with delay of 0.5 time unit.

noted as 1, 2, and 3, such that $d_{12} = d_{13} = d_{23} = 2R$, where R is the sensing range. The maximum speed of the sensor is such that it will take one time unit to cover a distance of $2R$. In a coverage schedule, therefore, the minimum staying time of the sensor at any PoI is $\delta = 1$ time unit. For each experiment, we report the average of 20 runs of the algorithm. The differences in the measurements are small. We will thus omit the error bars, although in the case of the deployment QoM, we will also report the maximum Q_* achieved in the 20 runs.

6.2.1 *Revisit of example (Section 5.2)*

This example motivates the use of optimized general periodic schedules. The proportional shares of 1, 2, and 3 are in ratios of 50:49:1. We show the optimizations over schedules of period m, where $m = 100$ time units. The algorithm in Fig. 2 is run with the initial schedule set to be the optimal *linear* periodic schedule of the given length. Figure 6(a) plots the maximum and average deployment QoM Q_* achieved by the simulated annealing algorithm for small computation budgets of up to 1000 iterations. The optimal deployment QoM is also shown as the horizontal green line in the figure. Figure 6(b) plots the corresponding results for larger computation budgets of up to 100000 iterations.

From the smaller computation budget results, note that the optimal linear periodic schedule is suboptimal in general but the simulated annealing produce schedules that can very quickly approach the optimal. From the larger computation budget results, the algorithm can find a solution extremely close to the optimal (within 2%). When $m = 400$ time units, the results (not shown due to space) are similar and a solution which is close to the optimal is found within 100000 iterations. Initially, however, Q_* increases more slowly with the number of iterations than $m = 100$ time units. This is because in this particular experiment, the globally optimal schedule can be found with a period length of 100 time units. Increasing the optimization period to 400 time units will not increase the potential to find a better solution, but will increase the search space for the optimal solution.

We have measured the run time of the simulated annealing, written in C#, on a Pentium-4 3.4 GHz PC with L1/L2 cache sizes of 8 KB/512 KB and 2 GB of RAM. The results (not shown due to space) indicate that the run time is linear in the number of iterations, and is about 3.5 s and 7.7 s for 100000 iterations and an optimization period of 100 and 400 time units, respectively.

6.2.2 *Other proportional shares*

We now use proportional share ratios of 53:29:17 for the 3 PoIs. The results are shown in Figure 7 for up to 5000 iterations when the optimization period is 99 time units. The search can approach the optimal Q_* value very quickly, within a few thousand iterations. The algorithm takes about 0.1s to complete 5000 iterations for an optimization period of 99 time units.

7. CONCLUSIONS

We have presented extensive analysis to understand the QoM properties of proportional-share mobile sensor coverage. We show that (1) A higher share of the coverage time generally increases the QoM, but the relationship is not linear except for blip events. (2) For staying events, the QoM can be much higher than the proportional share, due to the

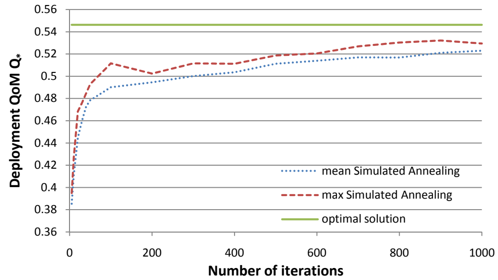

(a) computation budgets up to 1,000 iterations

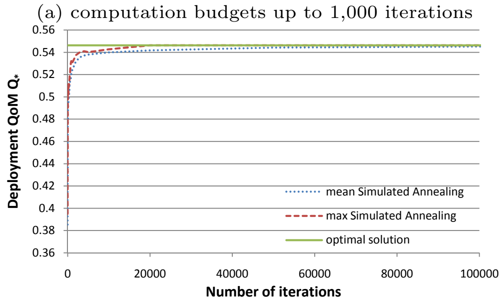

(b) computation budgets up to 100,000 iterations

Figure 6: Achieved deployment QoM Q_* for staying events with Step utility and proportional share ratios of 50:49:1. $X \sim Y \sim$ Exp($\lambda = 1$), period = 100 time units.

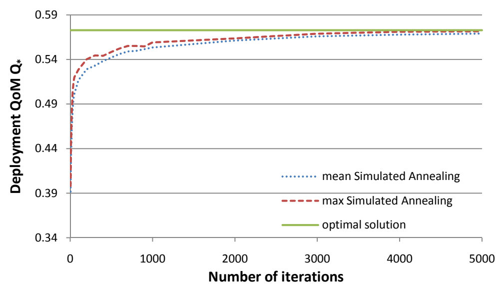

Figure 7: Achieved deployment QoM for proportional share ratios of 53:29:17. $X \sim Y \sim$ Exp($\lambda = 1$). Step utility function, period = 99 time units.

observation of "extra" events that arrive when the sensor is not present. This justifies mobile coverage from an information-capture point of view, i.e., the sensor gains by moving between places to search for new information. (3) The event utility function is important in determining the optimal fairness granularity p. For concave utility functions such as Step, Exponential, and Linear utilities, the QoM monotonically decreases with p (though for Linear, whose results were not shown due to space, it is initially flat for some range of p), whereas for Delayed Step and S-Shaped utilities, the QoM generally peaks at an intermediate p. Our analysis for Exponential/Pareto event dynamics and different forms of the utility function is all supported by the simulation results. We presented optimization algorithms for both linear and general proportional-share periodic coverage. Implementation results show that the simulated annealing algorithm can efficiently compute a periodic schedule that practically maximizes the total QoM, even for huge search spaces implied by long scheduling periods.

8. REFERENCES

[1] N. Bisnik, A. Abouzeid, and V. Isler. Stochastic event capture using mobile sensors subject to a quality metric. In *ACM MOBICOM*, 2006.

[2] V. Cerny. Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications*, 45(1), 1985.

[3] S. Chakraborty, Y. Dong, D. K. Y. Yau, and J. C. S. Lui. On the effectiveness of movement prediction to reduce energy consumption in wireless communication. *IEEE Transactions on Mobile Computing*, 5(2), February 2006.

[4] H. Gupta, S. Das, and Q. Gu. Connected sensor cover: self organization of wireless sensor networks. In *Mobihoc*, 2003.

[5] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. K. Miu, E. Shih, H. Balakrishnan, and S. Madden. Cartel: A distributed mobile sensor computing system. In *Proc. ACM Sensys*, 2006.

[6] K. Jeffay, F. D. Smith, A. Moorthy, and J. Anderson. Proportional share scheduling of operating system services for real-time applications. In *IEEE RTSS*, 1998.

[7] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science, New Series*, 220(4598):671–680, May 1983.

[8] A. K.Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Trans Networking*, 1(3), 1993.

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

[10] R. W. Lee and J. J. Kulesz. A risk-based sensor placement methodology. Technical report, Computational Sciences and Engineering Division, ORNL, 2006.

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

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

[13] W. Wang, V. Srinivasan, and K. C. Chua. Using mobile relays to prolong the lifetime of wireless sensor networks. In *ACM Mobicom*, 2005.

[14] D. K. Y. Yau, N. K. Yip, C. Y. T. Ma, N. S. Rao, and M. Shankar. Quality of monitoring of stochastic events by proportional-share mobile sensor coverage `www.cs.purdue.edu/homes/lans/publications/tr-08-011.pdf`. Technical Report 08-011, Purdue University, April 2008.

[15] H. Zhang and J. C. Hou. Maintaining sensing coverage and connectivity in large sensor networks. *Wireless Ad-hoc and Sensor Networks*, 1(1–2), January 2005.

[16] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a bus-based disruption-tolerant network. In *ACM Mobicom*, 2007.

[17] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *ACM Mobihoc*, 1994.