Joint 48th IEEE Conference on Decision and Control and
28th Chinese Control Conference
Shanghai, P.R. China, December 16-18, 2009

Resource Allocation for Signal Detection with Active Sensors

Jerome Le Ny, Michael M. Zavlanos and George J. Pappas

***Abstract*— We consider the problem of determining the existence of known constant signals over a set of sites, given noisy measurements obtained by a team of *active* sensors that can switch between different sites. Since the quality of detection depends on the time that the sensors allocate at every site, maximizing the total detection probability relies on selecting the sites and possibly the order in which these should be monitored. When the switching time between sites is negligible, as in steerable camera networks, we show that optimizing the global detection performance for a team of sensors with uncorrelated measurement noise is a convex problem. On the other hand, for significant switching times, which can be due to path planning for mobile robots in surveillance missions, the detection problem can be approximated by an integer program, known as the *orienteering problem*. Due to its hardness, even small instances of this problem are difficult to solve. Focusing on the single sensor problem, we propose a heuristic that employs the well studied *traveling salesman problem* to determine an optimal sequence of sites that maximizes the available time for detection. We finally show that when the switching penalties can be captured by a constraint on the number of sites to be observed, then submodularity of the unconstrained performance objective results in an effective greedy algorithm for selecting these sites.**

I. INTRODUCTION

The development of active sensor networks, steerable camera networks [1], and sophisticated mobile sensor platforms for Intelligence, Surveillance, and Reconnaissance missions (ISR) [2] is prompting researchers to revisit classical signal processing problems under the light of the newly offered capabilities. In these problems, studied under the names of sensor management, active learning, or optimal experiment design (see e.g. [3] and the references therein), deciding where and when the sensors should measure is strongly correlated with the sensing performance. A typical such example is the waveform selection problem [4]–[9], for target tracking using multifunction radars.

Introducing mobility in sensor networks couples sensor management with motion planning [10]. Recent work in this area includes for example controlling spacecraft formations for interferometric imaging [11], [12] and ocean sampling [13], and focuses mostly on estimation problems. Another well investigated problem concerns the optimization of the geometric configuration of a sensor network to improve the detection of instantaneous spatially distributed events [14]. Mobile sensor management for more dynamic detection problems forms essentially the topic of the search theory literature [15], [16], which often neglects the motion constraints and focus on single target single sensor problems (although there are exceptions, see [15]).

We consider a signal detection problem for multiple sites and multiple active sensors. The sensors must switch between the available sites and determine for each site whether a known signal exists or not, based on noisy observations. Since the probability of correct detection depends not only on the capabilities of the sensors, but also on the time allocated to each site, minimizing the time required to switch between sites is important. Applications of this problem range from camera networks in galleries or museums, to UAVs counting the number of events at certain locations of interest and reporting unusually high levels of activity. In this paper, we make the simplifying assumption of known constant signals distorted by white noise with known characteristics. The locations of the sites are also known and the overall task needs to be completed within given time limits.

In Section II we provide an overview of the results in signal detection theory that we rely on. In Section III we study the problem for negligible steering times between sites and show that finding the optimal time allocation per site is a convex problem. Note that after having completed the work described in this paper, we came across the paper [17], which is quite close to the model of section III. In Sections IV and V we consider the case of a single sensor, but with additional resource constraints. Namely, Section IV assumes that the sensor has enough energy to observe only a subset of given cardinality of the set of all sites. We show that the detection performance is a submodular set function, which allows us to characterize the performance of a greedy heuristic to select the subset of sites to observe. In Section V we assume that there are significant traveling times between the sites, which have to be subtracted from the total time available for detection. A discretized version of this problem reduces it to a combinatorial optimization problem known as the *orienteering problem*. Due to the computational complexity of this problem, a heuristic which consists in decoupling path planning and detection, and employing traveling salesman tours, provides solutions more efficiently.

II. HYPOTHESIS TESTING

A. A Single Sensor & A Single Site

Consider a single sensor monitoring a single site during a time interval $[0, T]$ and let the measurement be given by

$$x(t) = \theta s(t) + w(t), \quad 0 \le t \le T,$$

where $\theta \in \{0, 1\}$ is an unknown parameter, $s(t)$ is a *known* scalar signal, and $w(t)$ is a zero mean Gaussian white noise with known spectral height $N_0/2$. Our goal is to decide

This work was supported by the ONR-MURI award N00014-08-1-0696.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA jeromel, zavlanos, pappasg@seas.upenn.edu.

978-1-4244-3872-3/09/$25.00 ©2009 IEEE

between the hypothesis H_0 with $\theta = 0$, i.e., $x(t) = w(t)$, and H_1 with $\theta = 1$, i.e., $x(t) = s(t) + w(t)$. Equivalently, we want to infer from the measurements if the signal $s(t)$ is present at the site or not. An optimal detector [18] first passes the measured signal through a correlation filter $y = \int_0^T s(t)x(t)dt$. Then y is the realization of a Gaussian random variable Y with conditional mean

$$\mathbb{E}[Y|H_0] = 0,\ \mathbb{E}[Y|H_1] = \int_0^T |s(t)|^2 dt =: \|s\|_T^2 = \mu_1,$$

and variance under both hypothesis equal to $\frac{N_0}{2}\int_0^T |s(t)|^2 dt = \sigma_0^2$. Following the filter, the detector compares y to a threshold. More precisely, consider the likelihood ratio (LR)

$$L(y) = \frac{p_1(y)}{p_0(y)} = \frac{\frac{1}{\sqrt{2\pi\sigma_0^2}} e^{-(y-\mu_1)^2/2\sigma_0^2}}{\frac{1}{\sqrt{2\pi\sigma_0^2}} e^{-y^2/2\sigma_0^2}} = e^{y\mu_1/\sigma_0^2 - \mu_1^2/2\sigma_0^2}.$$

This LR is compared to a threshold τ that depends on the chosen performance criterion [18], [19]. Equivalently, we test the log-likelihood ratio

$$\frac{y\mu_1}{\sigma_0^2} - \frac{\mu_1^2}{2\sigma_0^2} \underset{H_0}{\overset{H_1}{\gtrless}} \log\tau \quad \Rightarrow \quad \frac{y}{\sigma_0} \underset{H_0}{\overset{H_1}{\gtrless}} \tilde{\tau} + \frac{\mu_1}{2\sigma_0},$$

with $\tilde{\tau} = (\sigma_0 \log\tau)/\mu_1$. Note that when $L(y) = \tau$ the optimal test needs in general to randomize the decision, but this case has no influence on the performance of the detector under the Gaussian assumption since it occurs with probability 0. For simplicity, we consider the Bayesian detector with uniform prior on H_0 and H_1 (or equivalently the minimax detector), in which case $\tau = 1$. Then, the probabilities of false alarm and missed detection are $P_F = \mathbb{P}\left(\frac{y}{\sigma_0} > \frac{\mu_1}{2\sigma_0}\ \middle|\ y \sim N(0, \sigma_0^2)\right)$ and $P_M = \mathbb{P}\left(\frac{y}{\sigma_0} < 0\ \middle|\ y \sim N(\mu_1, \sigma_0^2)\right)$, and both turn out to be equal to $1-\Phi(\mu_1/2\sigma_0)$, where Φ is the standard normal cumulative distribution function: $\Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^x e^{-u^2/2}\,du$. The total probability of error is

$$P_e = \frac{1}{2}P_M + \frac{1}{2}P_F = 1 - \Phi(\mu_1/2\sigma_0) = 1 - \Phi(\gamma),$$

where γ^2 denotes the signal-to-noise ratio (SNR)

$$\gamma^2 = \frac{\mu_1^2}{4\sigma_0^2} = \frac{\|s\|_T^4}{4\frac{N_0}{2}\|s\|_T^2} = \frac{\|s\|_T^2}{2N_0}.$$

Hence, the probability of error P_e depends on the observation interval $[0, T]$ through $\gamma = \|s\|_T/\sqrt{2N_0}$, with $\|s\|_T = \left(\int_0^T |s(t)|^2 dt\right)^{1/2}$.

B. Multiple Sensors & A Single Site

Consider now M sensors running their own correlation filters and all observing the same single site. Suppose sensor j observes the site during interval $I^j = [t^j, t'^j]$ (the sensors do not necessarily all start and stop observing at the same times t^j and t'^j). Under hypothesis H_1, sensor j's filter produces the random variable

$$\begin{aligned} Y^j &= \int_{t^j}^{t'^j} s(t)x(t)dt = \int_{I^j} s(t)x(t)dt \\ &= \int_{I^j} |s(t)|^2 dt + \int_{I^j} s(t)w^j(t)dt \\ &:= \mu^j(I^j) + \int_{I^j} s(t)w^j(t)dt. \end{aligned}$$

with conditional expectation

$$\mathbb{E}[Y^j|H_0] = 0,\ \mathbb{E}[Y^j|H_1] = \int_{I^j} |s(t)|^2 dt =: \mu^j(I^j)$$

and variance

$$\mathrm{Var}(Y^j|H_0) = \mathrm{Var}(Y^j|H_1) = \frac{N^j}{2}\int_{I^j} |s(t)|^2 dt = \frac{N^j}{2}\mu^j(I^j),$$

where $N^j/2$ is the spectral height of the Gaussian white measurement noise of sensor j. The cross-correlations between sensor signals now depend on the specific assumptions. In general,

$$\begin{aligned} \mathrm{Cov}(Y^{j_1}, Y^{j_2}|H_0) &= \mathrm{Cov}(Y^{j_1}, Y^{j_2}|H_1) \\ &= \mathbb{E}\left(\int_{I^{j_1}} s(t)w^{j_1}(t)dt \int_{I^{j_2}} s(t)w^{j_2}(t)dt\right). \end{aligned}$$

Letting $\mathbb{E}[w^{j_1}(t)w^{j_2}(t')] = \frac{N^{j_1 j_2}}{2}\delta(t-t')$ (with $N^{jj} = N^j$) we obtain the covariance matrix

$$\begin{aligned} \Sigma^{j_1 j_2}(I^{j_1}, I^{j_2}) &= \mathrm{Cov}(Y^{j_1}, Y^{j_2}|H_0) = \mathrm{Cov}(Y^{j_1}, Y^{j_2}|H_1) \\ &= \frac{N^{j_1 j_2}}{2}\int_{I^{j_1}\cap I^{j_2}} |s(t)|^2 dt. \end{aligned}$$

Assuming Σ is positive definite, then we can define $\Omega = \Sigma^{-1}$. Define further the quantities $Y = [Y^1 \dots Y^M]^T$ and $\mu = [\mu^1 \dots \mu^M]^T$. Then the log-likelihood ratio becomes

$$\begin{aligned} \log L(y) &= \log\frac{p_1(y)}{p_0(y)} = -\frac{1}{2}(y-\mu)^T\Omega(y-\mu) + \frac{1}{2}y^T\Omega y \\ &= \mu^T\Omega y - \frac{1}{2}\mu^T\Omega\mu. \end{aligned}$$

Under the assumption that all sensors share their information $y^1, \dots, y^M$, the log-likelihood ratio can be computed and results in the following test:

$$\mu^T\Omega y \underset{H_0}{\overset{H_1}{\gtrless}} \log\tau + \frac{1}{2}\mu^T\Omega\mu$$

or equivalently

$$\frac{\mu^T\Omega y}{\sqrt{\mu^T\Omega\mu}} \underset{H_0}{\overset{H_1}{\gtrless}} \frac{\log\tau}{\sqrt{\mu^T\Omega\mu}} + \frac{1}{2}\sqrt{\mu^T\Omega\mu}.$$

As before, consider the Bayesian detection problem with uniform prior (or minimax detection problem), for which $\tau = 1$ and so $\log\tau = 0$. Then $y \sim N(0, \Sigma)$ under H_0, and $\frac{\mu^T\Omega y}{\sqrt{\mu^T\Omega\mu}}$ is a normal random variable with mean 0 and variance

$$\frac{\mu^T\Omega E[yy^T]\Omega\mu}{\mu^T\Omega\mu} = \frac{\mu^T\Omega\Sigma\Omega\mu}{\mu^T\Omega\mu} = 1,$$

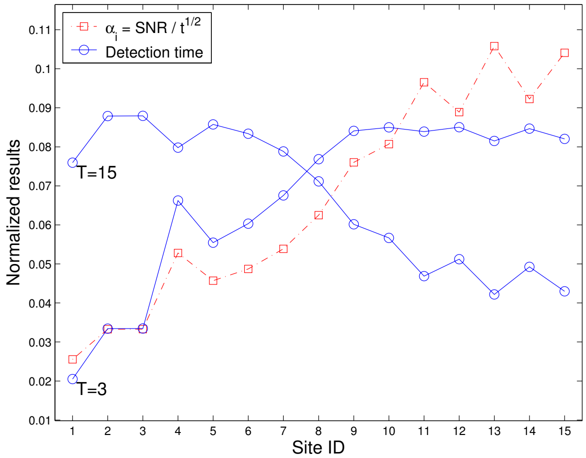

Fig. 1. Plot of the detection times per site and coefficients α_i for $n = 15$ sites and total detection times $T = 3, 15$. Observe that for small total time T, more time is spent at sites with large α_i, while the opposite is true for large T. This is due to the shape of the correct detection probability function Φ at any given site. Although for small time T sites with large coefficient α_i give a higher detection probability, this probability gets saturated at 1 much sooner than it does for sites with low α_i. Hence, when T is large, it is more beneficial to stay at sites with low α_i.

so the false alarm probability becomes $P_F = 1 - \Phi(\sqrt{\mu^T \Omega \mu}/2)$. By symmetry, this is also the probability of missed detection. The total probability of error is then

$$P_e = \frac{1}{2}P_F + \frac{1}{2}P_M = 1 - \Phi(\sqrt{\mu^T \Omega \mu}/2).$$

If all noises are independent, then $\Sigma^{j_1 j_2} = 0$ for $j_1 \neq j_2$ and so $\Omega^{jj} = 2/(\mu^j N^j)$. Hence,

$$\sqrt{SNR} = \frac{1}{2}\sqrt{\mu^T \Omega \mu} = \frac{1}{2}\sqrt{\sum_{j=1}^{M} \frac{2\,\mu^j (I^j)}{N^j}}$$

$$= \sqrt{\frac{1}{2}\sum_{j=1}^{M} \frac{\|s\|^2_{I^j}}{N^j}} = \sqrt{\sum_{j=1}^{M} SNR^j}.$$

In other words, with uncorrelated sensor noises, the total SNR can be obtained by adding the individual sensor SNRs, which greatly simplifies calculations.

III. Time Allocation without Switching Costs

In Section II we studied the signal detection problem for a single site and any number of sensors. When multiple sites are considered, an optimal allocation of the available time for detection at each site needs to be determined in order to minimize the total error probability. In this section, we address this problem for the case where the switching times between sites are negligible. This, for instance, is the case for networks of pan-tilt cameras [1].

A. A Single Sensor

Consider first the case of a single sensor observing N sites, one at a time, in order to determine whether *known* and *constant* signals $s_1(t) = s_1, s_2(t) = s_2, \ldots, s_N(t) = s_N$ are present or not. The observation noise when observing target i is Gaussian with spectral height $N_i/2$. Let t_i be the total time that the sensor spends observing target i, T the total time available to complete the detection task, and $\alpha_i = |s_i|/\sqrt{2N_i}$. Along the lines of Section II-A, our goal is to minimize a weighted sum of the error probabilities at all sites, or equivalently

$$\text{maximize} \sum_{i=1}^{N} w_i \Phi(\alpha_i \sqrt{t_i}) \tag{1}$$

$$\text{subject to} \sum_{i=1}^{N} t_i = T, \quad t_i \geq 0,\ \forall\, i = 1, \ldots, N.$$

where the weights satisfy $w_i > 0$ for $i = 1, \ldots, N$. Normalizing the times spent at every site, i.e., letting $p_i = t_i/T$ and $\beta_i = \alpha_i \sqrt{T}$, the optimization problem becomes (Fig. 1)

$$\text{maximize} \sum_{i=1}^{N} w_i \Phi(\beta_i \sqrt{p_i}) \tag{2}$$

$$\text{subject to } p = (p_1, \ldots, p_N) \in \Delta_{N-1},$$

where Δ_{N-1} is the $(N-1)$-dimensional simplex

$$\Delta_{N-1} = \{x \in \mathbb{R}^N | x_i \geq 0, \textstyle\sum_{i=1}^N x_i = 1\}.$$

In the following lemma we show that the function $p = (p_1, \ldots, p_N) \mapsto \Phi(\beta_i \sqrt{p_i})$ is concave. Since a sum of concave functions is concave, Lemma 1 implies that the optimization problem (2), or equivalently (1), is convex.

Lemma 1: For any $\beta_i \geq 0$ and any $i \in \{1, \ldots, N\}$, the function $p \mapsto \Phi(\beta_i \sqrt{p_i})$ is concave on Δ_{N-1}.

Proof: Observe first that the univariate function $p_i \mapsto \Phi(\beta_i \sqrt{p_i})$ is concave on $[0, 1]$. Indeed,

$$\frac{d}{dp_i}\Phi(\beta_i \sqrt{p_i}) = \frac{\beta_i}{2\sqrt{p_i}} \frac{1}{\sqrt{2\pi}} e^{-\frac{\beta_i^2 p_i}{2}}$$

and taking the second derivative we get

$$\frac{d^2}{dp_i^2}\Phi(\beta_i \sqrt{p_i}) = -1/8\, \frac{e^{-1/2\,\beta_i{}^2 p_i}\sqrt{2}\beta_i\,\left(\beta_i{}^2 p_i + 1\right)}{\sqrt{\pi} {p_i}^{3/2}},$$

which is negative in $(0, 1]$. To obtain the function of the lemma, we pre-compose with the projection on the i^{th} coordinate direction, which is linear, hence concavity is preserved [20, p.79]. ■

Remark 1: The objective function of problem (1) is separable by site and, since there is a single constraint tying the sites together, large-scale instances of this problem can be solved using the *dual decomposition* method, which decomposes at each subgradient step the large N-dimensional problem into N one-dimensional problems [21].

B. Multiple Independent Information-Sharing Sensors

The single sensor case described in Section III-A can be extended to the case of M sensors sharing all their information. As before, observations of site i by sensor j are subject to a Gaussian white noise with spectral height $N_i^j/2$, where noises of different sensors are assumed to be independent. Suppose that sensor j is available for a total time $T^j, j = 1, \ldots, M$ and let t_i^j be the time that sensor j spends observing site i. Using the results of section II-B, minimizing a weighted sum of the probability of detection error for all sites leads to the program

$$\text{maximize} \sum_{i=1}^{N} w_i \Phi\left(\frac{|s_i|}{\sqrt{2}}\sqrt{\sum_{j=1}^{M}\frac{t_i^j}{N_i^j}}\right) \quad (3)$$

$$\text{subject to} \sum_{i=1}^{N} t_i^j = T^j,\ j = 1, \ldots, M.$$

$$t_i^j \geq 0,\ i = 1, \ldots, N,\ j = 1, \ldots, M.$$

The decision variables are the times t_i^j and using Lemma 1, it is easy to see that (3) is a again a convex program. Moreover, it can also be decomposed and solved using dual decomposition methods.

IV. Time Allocation and Site Selection for a Single Sensor

While the detection problem takes a nice convex form when the switching costs between sites are negligible, this is usually not the case when additional resource constraints are introduced. In this section, we consider a single sensor and assume that due to energy constraints, it can only observe at most k of the N sites. As before, the the goal is to determine the existence of a constant signal $s_i(t) = s_i$ at site i and the measurements are subject to a noise with spectral height $N_i/2$ at that site. Then, the problem addressed in this section is that of selecting a subset of at most k sites to observe so that the overall weighted detection error probability is minimized.

For a uniform prior ($\mathbb{P}(\mathrm{H}_0) = \mathbb{P}(\mathrm{H}_1) = 1/2$) and the Bayesian detection problem under consideration, the error probability associated with not observing a site is $1/2$. Let $\mathcal{N} = \{1, \ldots, N\}$ denote the set of sites, and define the performance function $v(\cdot)$ for each subset $\mathcal{S} \subseteq \mathcal{N}$ as

$$v(\mathcal{S}) := \max_{\{t_i\}_{i\in\mathcal{S}}} \sum_{i\in\mathcal{S}} w_i(\Phi(\alpha_i\sqrt{t_i}) - 1) - \sum_{i\notin\mathcal{S}} \frac{w_i}{2} \quad (4)$$

$$\text{subject to} \sum_{i\in\mathcal{S}} t_i = T, \quad t_i \geq 0,\ i \in \mathcal{S}.$$

The overall weighted detection error probability associated with subset $\mathcal{S}$ is $-v(\mathcal{S})$. Note that for $t_i \geq 0$, we have $\Phi(\alpha_i\sqrt{t_i}) \geq \frac{1}{2}$. Rewriting $\tilde{\Phi}(\alpha_i\sqrt{t_i}) = \Phi(\alpha_i\sqrt{t_i}) - \frac{1}{2}$, we can equivalently rewrite the objective function as

$$v(\mathcal{S}) := \max_{\{t_i\}_{i\in\mathcal{S}}} \sum_{i\in\mathcal{S}} w_i\tilde{\Phi}(\alpha_i\sqrt{t_i}) - \sum_{i\in\mathcal{N}} \frac{w_i}{2}.$$

We obtain the optimization problem

$$\max_{\mathcal{S}\subseteq\mathcal{N}, |\mathcal{S}|\leq k} v(\mathcal{S}). \quad (5)$$

Clearly a solution $\mathcal{S}$ to (5) has $|\mathcal{S}| = k$. Still, comparing the value of $v(\mathcal{S})$ for all possible subsets of size k requires solving (4) $\binom{N}{k}$ times, which can be impractical as N grows. Note, for instance, that for $N = 2k$ we have $\binom{2k}{k} \approx \sqrt{\frac{2}{\pi}}\frac{4^k}{\sqrt{2k+1}}$, which for $k = 10$ gives approximately $1.8 \cdot 10^5$ possible subsets of size 10.

To address this problem we propose a greedy heuristic to find a subset $\mathcal{S}$ of cardinality k. We initialize the subset of sites as $\mathcal{S}_0 = \emptyset$ and for all future iterations $1 \leq t \leq k$ we obtain $\mathcal{S}_t$ by $\mathcal{S}_t = \mathcal{S}_{t-1} \cup \{j_t\}$ where

$$j_t \in \arg\max_{j\notin\mathcal{S}_{t-1}} \{v(\mathcal{S}_{t-1} \cup \{j\})\}.$$

This step can be executed by simply computing the value $v(\mathcal{S}_{t-1} \cup \{j\})$ of (4) for all $j \notin \mathcal{S}_{t-1}$. The final set is then $\mathcal{S}_k = \{j_1, \ldots, j_k\}$. Hence the greedy policy requires the computation of the program (4) $O(kN)$ times, and the computation at step t involves t variables. Characterizing the performance of this greedy heuristic relies on showing *submodularity* of the function $\mathcal{S} \mapsto v(\mathcal{S})$, i.e.,

$$v(\mathcal{S}_1) + v(\mathcal{S}_2) \geq v(\mathcal{S}_1\cap\mathcal{S}_2) + v(\mathcal{S}_1\cup\mathcal{S}_2),\ \forall\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{N}. \quad (6)$$

Proposition 1: The function $\mathcal{S} \mapsto v(\mathcal{S})$ is submodular.

Proof: The result follows from Theorem 3.5 in [22]: Let $N = |\mathcal{N}|$ be the number of sites, and $\delta \in \{0, 1\}^N$ be a vector with $0-1$ entries, where $\delta_i = 1$ if and only if site i is selected. Hence each vector δ corresponds to a subset $\mathcal{S} \subset \mathcal{N}$. Now let $f_i(p_i, \delta_i) = \delta_i w_i\tilde{\Phi}(\beta_i\sqrt{p_i})$. Then v is submodular if and only if the following function is submodular

$$\tilde{v}(\delta) = \max \sum_{i=1}^{N} f_i(p_i, \delta_i) \quad (7)$$

$$\text{subject to} \sum_{i=1}^{N} p_i = 1, \quad p_i \geq 0,\ i = 1, \ldots, N,$$

where we identify $\tilde{v}(\delta)$ and the corresponding $\tilde{v}(\mathcal{S})$. Note that summing over all indices $1 \leq i \leq N$ in the constraint is possible, since for any fixed parameter δ, if $\delta_i = 0$ for some i, then $f_i(p_i, 0) = 0$ implies $p_i = 0$ in the optimal solution. Hence the constraint $\sum_{i\in\mathcal{S}} p_i = 1$ can be replaced by the constraint $\sum_{i=1}^{N} p_i = 1$ without changing the value of $v(\delta)$.

Consider now two variables p_i and p_j of a vector p satisfying the single linear equality constraint $\mathbf{1}_N^T p = 1$ of (7). In the terminology of [22], we have trivially that p_i and p_j are $\mathcal{B}_0$-substitutes. Hence we deduce the following: First, for all $i, j \in \mathcal{N}$, the optimal value $p_i^*(\delta)$ is nonincreasing in δ_j. This means that for every δ^1, δ^2 with $\delta^1 \leq \delta^2$ and δ^1 and δ^2 differing only on the j^{th} component we have $p_i^*(\delta^2) \leq p_i^*(\delta^1)$. That is, adding one more site j to a set $\mathcal{S}$ will reduce the time spent at every site in $\mathcal{S}$ (in the optimal solution) before j was added. Second, $v(\delta)$ is submodular in δ_i and δ_j for all $i, j \in \mathcal{N}$. According to [22], this means that for every δ^1 and δ^2 with $\delta^1 \leq \delta^2$ and δ^1 and δ^2 differing only in the i^{th} and j^{th} components, we have

$$\tilde{v}(\delta^1) + \tilde{v}(\delta^2) \leq \tilde{v}(\ldots, \delta_i^1, \ldots, \delta_j^2, \ldots) + \tilde{v}(\ldots, \delta_i^2, \ldots, \delta_j^1, \ldots).$$

For δ^1 and δ^2 corresponding to $\mathcal{S}$ and $\mathcal{S}\cup\{i, j\}$, $i, j \in \mathcal{N}\setminus\mathcal{S}$, this inequality can be rewritten

$$\tilde{v}(S \cup \{i\}) + \tilde{v}(S \cup \{j\}) \geq \tilde{v}(S) + \tilde{v}(S \cup \{i, j\}),$$

which is a necessary and sufficient condition for $S \to \tilde{v}(S)$ to be submodular according to our definition (6) [23, p.767]. ∎

In addition to submodularity, it is also clear that the function $\mathcal{S} \mapsto v(\mathcal{S})$ is nondecreasing, i.e., $\mathcal{S} \subseteq \mathcal{T} \Rightarrow v(\mathcal{S}) \leq v(\mathcal{T})$, and that $v(\emptyset) = -\sum_{i \in \mathcal{N}} \frac{w_i}{2} =: -W$. Hence the following proposition is a consequence of Proposition 1 and a result of [24]. Here we denote by $\mathcal{S}_k$ the set of size k selected by the greedy heuristic, by $P_e(k, \mathcal{S}_k)$ its *weighted* detection error probability (i.e., $P_e(k, \mathcal{S}_k) = -v(\mathcal{S}_k)$), and by $P_e^*(k)$ the optimal weighted detection error probability achievable by some subset of size k.

Proposition 2: The greedy algorithm returns a set $\mathcal{S}_k$ such that

$$P_e(k, \mathcal{S}_k) - P_e^*(k) \leq \frac{1}{e}(W - P_e^*(k)).$$

Note here that W is the performance we would obtain if none of the sites were selected. Hence the proposition provides a performance guarantee for the greedy heuristic by bounding its difference with respect to the optimum performance achievable.

V. Joint Path Planning and Detection for a Single Mobile Sensor

In sensing applications involving mobile robots, simply abstracting the switching penalties by bounding the number of visited sites as in Section IV does not necessarily capture adequately the joint path planning and scheduling problem that arises. More realistically, there is in general a specified amount of time to complete the whole detection task, including traveling between sites, and minimizing the amount of time spent traveling is critical to maximize the overall probability of correct detection.

In this section, we are given the traveling times between the different sites, and we consider the joint signal detection and path planning problem for a single mobile sensor. We have developed two heuristics to solve this problem. The first, which we do not describe in details due to space constraints, relies on partitioning the total detection probability at every site in a finite number of equal intervals and requiring that the sensor stays at a site for the amount of time necessary to increase the detection probability by a number of such intervals, which is itself a decision variable of the problem. This discretization of the detection benefit and time at every site leads to a well studied integer programming formulation of the problem, known in the literature as the *orienteering problem* [25]–[29] Although the solution to this integer program tends to the optimal solution of the original problem as the discretization becomes finer, solving this integer program can be difficult in practice. To find solutions in a way that scales better with the number of sites we propose a second approximation to the problem, which essentially decouples path planning from detection. In particular, path planning relies on *traveling salesman* tours [30], for which scalable constant factor approximation algorithms and good solvers exist, while detection depends on the solution of problem (4). This heuristic works well in practice, however, characterizing its accuracy is still work in progress.

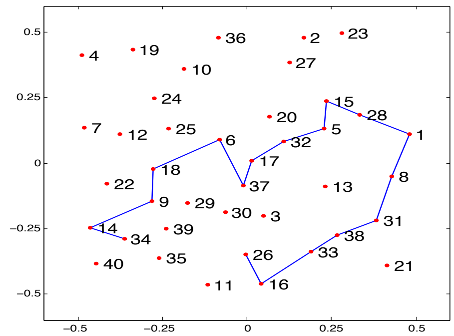

Fig. 2. Final TSP tour for a sensor observing 40 sites. Alg. 1 returns a tour containing 18 sites. The vehicle starts at site 34, and the edge closing the tour is not shown for clarity.

In more details, we compute for different subsets of sites a traveling salesman tour that minimizes the total travel time, while the remaining time is optimally allocated, by solving problem (4) for the sites in the tour[1] (Alg. 1). This approach relies on the availability of effective software tools, which are able to find optimal solutions for instances of the Traveling Salesman Problem (TSP) involving hundreds of sites. In particular we used LKH, Helsgaun's implementation of the Lin-Kernighan heuristic, available at [31]. To find a good subset of nodes to include in the tour, we sequentially remove sites one-by-one from the original set. Using this heuristic, results can be obtained for large numbers of sites.

We tested Alg. 1 for 40 sites randomly distributed in the unit square. The signals and noise levels at every site were randomly chosen from the intervals (1,5) and (1,2), respectively. The total available time for the combined motion and detection task was $T = 5$. Fig. 2 shows the sub-optimal TSP path that the agent follows. Fig. 3 shows the variation in the performance $\sum_{i \in \mathcal{S}} \tilde{\Phi}(\alpha_i \sqrt{t_i})$ for the successive tours computed by Alg. 1.

References

[1] M. Chu, J. Reich, and F. Zhao, "Distributed attention in large scale video sensor networks," in *Intelligent Distributed Surveilliance Systems*, 2004.

[2] UAV roadmap, "Unmanned aircraft systems roadmap 2005-2030," Office of the Secretary of Defense, Tech. Rep., 2005. [Online]. Available: http://www.acq.osd.mil/usd/Roadmap%20Final2.pdf

[3] J. Williams, "Information theoretic sensor management," Ph.D. dissertation, Massachusetts Institute of Technology, February 2007.

[4] M. Athans, "On the determination of optimal costly measurement strategies for linear stochastic systems," *Automatica*, vol. 8, pp. 397–412, 1972.

[5] D. Kershaw and R. Evans, "Optimal waveform selection for tracking systems," *IEEE Transactions on Information Theory*, vol. 40, no. 5, pp. 1536–1550, 1994.

[6] V. Krishnamurthy and R. Evans, "Hidden markov model multiarm bandits: a methodology for beam scheduling in multitarget tracking," *IEEE Transactions on Signal Processing*, vol. 49, no. 12, pp. 2893 – 2908, December 2001.

[1] In this paragraph, we assume that the sensor must start and end at the same site. A tour is a closed path. In Alg. 1 for simplicity, the initial site is not specified *a priori* but can be chosen to be any site on the final tour.

Algorithm 1 Decoupling planning and detection

Require: A set $\mathcal{S}$ of n available sites;

Compute the TSP tour $TSP(\mathcal{S})$ and the associated travel time $T_{TSP}(\mathcal{S})$;
if $T - T_{TSP}(\mathcal{S}) > 0$ **then**
  Solve problem (4) for the performance $v^\star_{|\mathcal{S}|} = v(\mathcal{S}, T - T_{TSP}(\mathcal{S}))$;
else
  Set $v^\star_{|\mathcal{S}|} = 0$;
end if
while $|\mathcal{S}| > 1$ **do**
  for $k \in \mathcal{S}$ **do**
    Compute the TSP tour $TSP(\mathcal{S}\backslash\{k\})$ and the associated travel time $T_{TSP}(\mathcal{S}\backslash\{k\})$;
    if $T - T_{TSP}(\mathcal{S}\backslash\{k\}) > 0$ **then**
      Solve problem (4) for the performance $v(\mathcal{S}\backslash\{k\}, T - T_{TSP}(\mathcal{S}\backslash\{k\}))$;
    else
      Set $v(\mathcal{S}\backslash\{k\}, T - T_{TSP}(\mathcal{S}\backslash\{k\})) = 0$;
    end if
  end for
  if $\max_{k\in\mathcal{S}} v(\mathcal{S}\backslash\{k\}, T - T_{TSP}(\mathcal{S}\backslash\{k\})) > 0$ **then**
    The site to be removed from the tour is $k^\star = \text{argmax}_{k\in\mathcal{S}} v(\mathcal{S}\backslash\{k\}, T - T_{TSP}(\mathcal{S}\backslash\{k\}))$;
  else
    Set $k^\star = \text{argmin}_{k\in\mathcal{S}} T_{TSP}(\mathcal{S}\backslash\{k\})$;
  end if
  Set $v^\star_{|\mathcal{S}|-1} = v(\mathcal{S}\backslash\{k^\star\}, T - T_{TSP}(\mathcal{S}\backslash\{k^\star\}))$;
  Remove site $k^\star$ from $\mathcal{S}$, i.e., $\mathcal{S} = \mathcal{S}\backslash\{k^\star\}$;
end while
The optimal tour and detection times are obtained for the j^{th} iteration of the algorithm satisfying $j = \text{argmax}_{i=1,\dots,n} v^\star_i$;

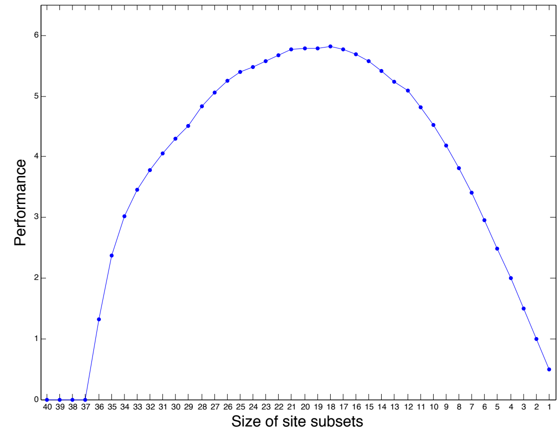

Fig. 3. Plot of the maximum performance (i.e., $\sum_{i\in\mathcal{S}} \tilde{\Phi}(\alpha_i\sqrt{t_i})$) associated with TSP tours of different size on the set of 40 sites. The best performance is obtained for the TSP tour on 18 sites, illustrated in Fig. 2.

[7] V. Gupta, T. Chung, B. Hassibi, and R. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage," *Automatica*, vol. 42, no. 2, pp. 251–260, 2006.

[8] J. Le Ny, M. Dahleh, and E. Feron, "Multi-UAV dynamic routing with partial observations using restless bandits allocation indices," LIDS, Massachusetts Institute of Technology, Tech. Rep., 2007. [Online]. Available: http://web.mit.edu/jleny/www/publications.html

[9] J. Le Ny, E. Feron, and M. Dahleh, "Scheduling Kalman filters in continuous time," in *Proceedings of the American Control Conference*, June 2009. [Online]. Available: http://arxiv.org/abs/0810.5148

[10] D. Uciński, *Optimal Measurement Methods for Distributed Parameter System Identification*. CRC Press, 2005.

[11] R. Beard, J. Lawton, and F. Hadaegh, "A coordination architecture for spacecraft formation control," *IEEE Transactions on Control Systems Technology*, vol. 9, no. 6, pp. 777–790, November 2001.

[12] I. I. Hussein and A. Bloch, "Dynamic coverage optimal control for interferometric imaging spacecraft formations," in *IEEE Conference on Decision and Control*, 2004.

[13] N. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. Fratantoni, and R. Davis, "Collective motion, sensor networks, and ocean sampling," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 48–74, January 2007.

[14] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, April 2004.

[15] S. Benkoski, M. Monticino, and J. Weisinger, "A survey of the search theory literature," *Naval Research Logistics*, vol. 38, pp. 469–494, 1991.

[16] L. Stone, *Theory of optimal search*, 2nd ed. Military application section, Operations Research Society of America, 1992.

[17] E. Duflos, M. de Vilmorin, and P. Vanheeghe, "Time allocation of a set of radars in a multitarget environment," in *International Conference on Information Fusion (FUSION 2007)*, Quebec, Canada, July 2007.

[18] H. Van Trees, *Detection, Estimation, and Modulation Theory: Part I*. New York: Wiley, 1968.

[19] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Springer, 1998.

[20] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2006.

[21] J. Bonnans, J. Gilbert, C. Lemarechal, and C. Sagastizábal, *Numerical Optimization*. Springer, 2006.

[22] J. Provan, "Substitutes and complements in constrained linear models," *SIAM Journal on Algebraic and Discrete Methods*, vol. 8, no. 4, pp. 585–603, October 1987.

[23] A. Schrijver, *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

[24] G. Nemhauser, L. Mosley, and M. Fisher, "An analysis of approximations for maximizing submodular set functions - I," *Mathematical Programming*, vol. 14, pp. 265–294, 1978.

[25] B. L. Golden, L. Levy, and R. Vohra, "The orienteering problem," *Naval Research Logistics*, vol. 34(3), pp. 307–318, 1987.

[26] C. P. Keller, "Algorithms to solve the orienteering problem: A comparison," *European Journal of Operational Research*, vol. 41, no. 2, pp. 224–231, July 1989.

[27] R. Ramesh, Y.-S. Yoon, and M. H. Karwan, "An optimal algorithm for the orienteering tour problem," *ORSA Journal on Computing*, vol. 4(2), pp. 155–165, 1992.

[28] M. Fischetti and P. Toth, "Solving the orienteering problem through branch-and-cut," *INFORMS Journal on Computing*, vol. 10, pp. 133–148, 1998.

[29] K. Chen and S. Har-Peled, "The Euclidean orienteering problem revisited," *SIAM J. Comput.*, vol. 38, no. 1, pp. 385–397, 2008.

[30] G. Gutin and A. P. Punnen, *The Traveling Salesman Problem and Its Variations*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2007.

[31] K. Helsgaun, "An effective implementation of the Lin-Kernighan traveling salesman heuristic," *European Journal of Operational Research*, vol. 126, no. 1, pp. 106–130, 2000. [Online]. Available: http://www.akira.ruc.dk/ keld/research/LKH/