Energy Efficient Policies for Distributed Target Tracking in Multihop Sensor Networks

Shuchin Aeron, Venkatesh Saligrama, David A. Castañón

Abstract— **We consider the problem of distributed target tracking in a sensor network under communication constraints between the sensor nodes, a problem that has recently received significant attention. The problem requires the dynamic selection of which sensor nodes will communicate their information and the selection of a corresponding fusion center which will process the collected information. Ideally, selection of which sensors will communicate and where will fusion take place will be a dynamic process, adapting to new information, to trade off tracking accuracy versus communications usage. The resulting coupled problem is generally intractable and significant effort has been devoted towards proposing simple strategies under various performance criteria. In this paper, we propose an adaptive dynamic strategy for sensor selection and fusion location using a certainty equivalence approach that seeks to optimize a tradeoff between tracking error and communications cost. We define a certainty equivalent optimization problem for dynamic relocation of the fusion center that uses measures of average multi hop communications cost and average tracking errors, and solve the resulting optimal control problem for classes of tracking problems. The optimal strategy is a hybrid switching strategy, where the fusion center location and reporting sensors are held stationary unless the target estimates move outside of a threshold radius around the sensors. We illustrate the performance of our algorithms on sample tracking experiments with sensor networks.**

I. Introduction

Consider the problem of tracking a target moving in region populated by sensor nodes that have limited wireless communication capabilities. The sensors have limited sensing range and the quality of measurements degrades with distance from the target. In such networks, the ideal tracking system would process measurements from every sensor at a single location to estimate the target state. However, this centralization is not possible due to communication constraints on individual sensors.

In the presence of communication constraints, one must use a distributed mechanism for selecting which information should be communicated, and where should the information be processed. In this paper, we explore a particular distributed protocol for tracking in sensor networks with communication constraints. We assume that there is no central authority governing network operation. Since measurement quality degrades with distance, only sensors sufficiently close to the target should share their data. Furthermore, since there is no designated central authority an arbitrary sensor is designated as the fusion center (or *leader node*) for a given time period and the data is processed at this sensor. Therefore, the entire mechanism can be broken down into three tasks:

1) Fusion center estimates the current target position based on sensed information from sensor set.
2) Fusion center computes the new sensor set as well as the new fusion center for the next time period based on current estimated target state.
3) Fusion center communicates state information to the new fusion center.

Communication costs arise while transmitting data from sensors to the fusion center and while transmitting current state information to a subsequent fusion center.

An early approach for finding the best informative sensor along with communication costs was taken in [1],[2]. In [3] non-myopic strategies for sensor selection was proposed with an information utility measure. However the cost of communication between the successive sensors and cost of aggregation was not taken into account. In [4] the problem of tracking with leader selection and communication costs was formulated as a dynamic program with information theoretic utility measures as cost per stage, and constraints on the total communications. Due to analytical intractability they used approximate dynamic programming methods to solve finite horizon DP.

In this paper, we propose and analyze an alternative formulation that captures the inherent tradeoffs between communication costs and estimation performance, based on a continuum approximation to the number of sensors. We develop and solve a fully observed Markov decision problem for sensor and fusion center selection. We show that, for classes of target dynamic models, an optimal strategy is a switching strategy characterized by dead zones. The reporting sensor locations are stationary until the target exits the dead zone and then locations are switched to best sensor locations around the target.

In order to apply our results to the control of sensor networks with partial information about target states, we adopt a certainty equivalent perspective, where target state estimates are treated as corresponding to actual current states. We implement our sensor network control algorithms in several tracking scenarios, and show the resulting reductions in communication that arise from effective network control in contrast to alternative algorithms.

The paper is organized as follows. In section II we will present the target tracking problem of interest, and discuss the difficulties with existing approaches. In section III we

The authors are with the Dept. of Electrical and Computer Engineering at Boston University, Boston, MA-02215.(shuchin,srv,dac)@bu.edu. This research was supported by the ONR Young Investigator Program and Presidential Early Career Award (PECASE) N00014-02-100362, NSF CAREER award ECS 0449194, and NSF Grants DMI 0330171, CCF 0430983 and CNS-0435353

will formulate and solve an approximation for the tracking problem with two different models for target dynamics. In section IV, we illustrate the performance of our approximate approach compared to alternative approaches in two examples. Section V summarizes the results.

II. General Problem

Consider a collection, $\mathcal{S}$ of n sensors at positions, $s_1, s_2, \ldots, s_n$. At any time k the observation at the sensor s_i is given by,

$$Z_k^i = F(X_k, s_i) + V_k^i$$

where X_k is the state at time k of the target, typically consisting of position and velocity, and V_k^i is additive additive, white Gaussian noise of variance Σ_v and is independent of $X_{k'}$ for all k'. The target state evolves as an autonomous system as follows:

$$X_{k+1} = f(X_k, k) + w_k \tag{1}$$

where w_k is zero mean white Gaussian noise of variance Σ_w which is assumed independent of $V_{k'}^i$ for all i, k' and independent of the initial condition X_0.

The problem of interest is to select the fusion center at time $k+1$ and the set of active sensors to report at time $k+1$, based on the information available at time k, with a goal of minimizing a tradeoff between communications cost and tracking error. A similar problem was studied in [4], where they sought to maximize the information collected over time subject to a communications constraint. The resulting formulation was a Partially Observed Markov Decision problem with a combinatorial decision space whose exact solution was intractable.

In this paper, we focus on the problem of locating the fusion center, assuming the reporting sensors will be in a prespecified neighborhood of the fusion center. This reduces the combinatorial complexity of the decision problem while focusing on the fundamental issue of trading tracking accuracy with communications cost.

Let ℓ_k denote the location of the fusion center at stage k. Let $\mathcal{I}_k$ denote the information available for decision at time k, which includes all of the past observations collected and the past locations of the fusion center, including the observations at time k. The control problem at time k is to select a strategy $\mu_k : \mathcal{I}_k \to \mathcal{S}$ that determines where the fusion center will be located at time $k-1$.

The performance objective for selecting a control strategy is a tradeoff between communications costs and tracking error. We express this cost over a finite horizon N

$$E\{\sum_{k=0}^{N-1}[Kc_1(\ell_{k+1} - \ell_k) + c_2(X_k, I_k)] + c_2(X_N, I_N)\} \tag{2}$$

where $c_1(\cdot)$ is the communications cost of switching the fusion center between locations, $c_2(\cdot,\cdot)$ represents the tracking error. For our purposes, we choose the square estimation error

$$c_2(X_k, I_k) = (X_k - E[X_k|I_k])^T Q (X_k - E[X_k|I_k])$$

where Q is a weighting matrix that can be used to select the position entries in X_k, or any other desired weighting. The communications cost is chosen as $\| u \|_1$, as discussed later.

Note that, when $N > 1$, the resulting problem is a partially observed Markov decision problem with an underlying continuous state space (X_k), rendering the problem intractable unless the sensor observations have special structure (e.g. finite valued measurements). For the special case that the horizon $N = 1$, the problem reduced to evaluating the cost in (2) for each of the possible choices of fusion center locations at the next stage, $\ell_1 \in \mathcal{S}$. This can be done approximately using a Cramer-Rao bound approach or an extended Kalman filter (EKF) to approximate the error covariance of the estimate, as in [5], as follows. Let $\Sigma_{0|0}$ denote the error covariance $E[(X_0 - E[X_0|I_0])(X_0 - E[X_0|I_0)^T|I_0]$, and let $\hat{X}_0$ denote the estimate of the state at time 0 given I_0. Then, the EKF can be used to estimate the tracking error at time 1 given the choice of $\ell_1 = s_i$, as

$$\begin{aligned}
\hat{X}_{1|0} &= f(\hat{X}_0, 0) \\
\Sigma_{1|0} &= \frac{\partial}{\partial X} f(X,0)|_{\hat{X}_0} \Sigma_0 \frac{\partial}{\partial X} f(X,0)|_{\hat{X}_0}^T + \Sigma_w \\
C^j &= \frac{\partial}{\partial X} F(X, s_j)|_{\hat{X}_{1|0}} \\
\Sigma_{1|1}^{-1}(i) &= \Sigma_{1|0}^{-1} + \sum_{j \in N_i} (C^j)^T \Sigma_v^{-1} C^j \\
E[c_2(X_1, I_1)] &\approx Trace[\Sigma_{1|1}(i) Q]
\end{aligned} \tag{3}$$

The choice of ℓ_1 is selected as

$$\begin{aligned}
i^* &= \mathrm{arg} min_i \{Trace[\Sigma_{1|1}(i)Q] + Kc_1(s_i - s_{i_0}) \\
\ell_i^* &= s_{i^*}
\end{aligned} \tag{4}$$

The above algorithm can be used in a receding horizon manner to generate a decision policy $\mu_k(I_k)$ for each time k, based on a one-step lookahead horizon. However, extending the approach to a policy that is based on looking ahead for more than one time period is much more complex. In the next section, we present an approximate approach that allows us to develop decision policies based on multi-period horizons.

III. Continuous Approximation

We want to formulate a simpler problem, which captures the essential features of the original problem. Assume that we have a nearly uniform placement of identical sensors in a region, and that the quality of observations degrades substantially with increasing sensor distance to the target. We assume that the reporting sensors are located near the fusion center. Thus, we approximate the tracking error by the distance from the fusion center to the target. We also adopt a certainty equivalent perspective, decoupling the problem of estimation and control, by modeling the current target position as perfectly observed due to the sensor density. We refer to this approach as Certainty Equivalent Control (CEC).

In terms of modeling communication costs, single-hop communication energy is proportional to the square (or fourth power) of the distance of communication. However, in low-power sensor networks, the usual protocols are multi-hop, and the communication cost is dominated by the number of hops. In our context, the main communications cost is dominated by the switch of fusion center location. For a mesh network with multihop communications on a uniform grid, this is proportional to $\| \ell_1 - \ell_2 \|_1$, where ℓ_1 is the location of the current fusion node and ℓ_2 is the new location.

Let ℓ_k denote the location of the leader node at time k, and let p_k denote the position of the target at time k. The sensors measuring the target are located around the leader; since measurement signal power decays with distance, we model tracking error as proportional to $||p_k - \ell_k||^2$. The decision at time k is the location of the next active sensor, denoted by ℓ_{k+1}, based on the current state of the target. Let $u(t) = \ell(t+1) - \ell(t)$. The cost of communication is proportional to $||u_k||_1 = |u_x(t)| + |u_y(t)|$ for multi-hop networks.

A. Random Walk Model

We analyze first a simple dynamic system where the target motion is one-dimensional. The results extend to two-dimensional motion. as the minimization at each step is decoupled in different dimensions. Let the target dynamics be given by

$$x_{k+1} = x_k + w_k$$

where w_k is a zero mean, white noise independent of x_k. The state of the augmented system includes both the target and fusion center positions, as $\tilde{x}_k = [x_k, \ell_k]'$ with dynamics

$$\tilde{x}_k = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tilde{x}_k + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_k + \begin{bmatrix} 1 \\ 0 \end{bmatrix} w_k \qquad (5)$$

Defining a new state $z_k = x_k - \ell_k$ with dynamics

$$z_{k+1} = z_k - u_k + w_k$$

As discussed previously, the objective function is

$$\min_{u_1,..,u_{N-1}} \mathbf{E}\left(\sum_{k=1}^{N-1} \left[||u_k||_1 + ||z_k||^2\right] + ||z_N||^2\right)$$

For sake of exposition we consider two cases, viz., (a) no process noise, i.e., a stationary target and (b) a uniform process noise.

1) Noiseless Case: Assume $w_k = 0$, and perfect state observation. This is a deterministic control problem with non-zero initial condition. The cost-to-go at time $N-1$, with one step to go, is

$$J_{N-1}(z_{N-1}) = \min_u \left\{|u| + z_{N-1}^2 + (z_{N-1} - u)^2\right\} \qquad (6)$$

This is a convex function of $|u_{N-1}|$, which is not differentiable at 0. The subgradient set [6] of the right hand side of (6) at $u = 0$ is $[-1, 1] - 2z_{N-1}$. Hence, 0 is in the subgradient set if $|z_{N-1}| < 1/2$, which is a sufficient condition for $u = 0$ to be optimal [6]. At $u \neq 0$, the right hand side of (6) is differentiable, so the optimal control is $u = z_{N-1} - 0.5 \ \mathrm{sign}(z_{N-1})$. Then,

$$J_{N-1}(z_{N-1}) = \begin{cases} 2z_{N-1}^2 & \text{if } |z_{N-1}| \leq 1/2 \\ |z_{N-1}| - 1/4 + z_{N-1}^2 & \text{if } |z_{N-1}| > 1/2 \end{cases}$$

The interpretation of the control at N-1 is that we do not apply any control if the state z_{N-1} is in a "dead-zone" region $[-1/2, 1/2]$. Otherwise, the control brings the state to the edge of this region.

By induction, at time $N-k$ the cost to go function is given by

$$J_{N-k}(z_{N-k}) = \begin{cases} (k+1)z_{N-k}^2 & |z_{N-k}| \leq 1/2k \\ |z_{N-k}| - \frac{1}{2k} + z_{N-k}^2 + \frac{k}{4k^2} & |z_{N-k}| > 1/2k \end{cases}$$

and the optimal "dead-zone" region at stage $N-k$ is $[-1/2k, 1/2k]$. Note that the cost-to-go is a differentible convex function of z, and the dead-zone shrinks as the number of stages k increases.

2) Uniform Bounded Process Noise: Suppose now that the noise process w_k is not zero, so we have a stochastic control problem. Assume that the white noise process w_k is a uniform noise bounded between $[-\alpha, \alpha]$. Under perfect state observation, the cost-to-go at stage $N-1$ is

$$J_{N-1}(z_{N-1}) = \min_{u_{N-1}} (\lambda|u_{N-1}| + z_{N-1}^2 + \mathbf{E}_w(z_{N-1} - u_{N-1} + w_{N-1})^2)$$

Since the noise is zero mean and uncorrelated, the cost to go at stage $N-1$ is given by

$$J_{N-1}(z_{N-1}) = \begin{cases} 2z_{N-1}^2 + \sigma_w^2 & |z_{N-1}| \leq \lambda/2 \\ \lambda||z_{N-1}| - \lambda/2| + \lambda^2/4 + z_{N-1}^2 + \sigma_w^2 & |z_{N-1}| > \lambda/2 \end{cases}$$

By induction, we establish the following result.

Theorem 3.1: The optimal n-stage policy is a switching policy, i.e.,

$$u_{N-n} = 0, \ |z| \leq \Gamma_n$$
$$u_{N-n} = (|z| - \Gamma_n)\mathrm{sign}(z), \ |z| \geq \Gamma_n$$

where, $\Gamma_n > 0$. The corresponding n-stage cost-to-go is described by:

$$J_{N-n}(z) = \begin{cases} z^2 + \mathbf{E}(J_{N-n+1}(z+w)), \ |z| \leq \Gamma_n \\ \lambda\left(|z| - \Gamma_n\right) + z^2 + \mathbf{E}J_{N-n+1}(\Gamma_n + w) \ \text{else} \end{cases}$$

The switching point Γ_n is uniformly bounded from below, i.e.,

$$\Gamma_n \geq \Gamma_0 = \min\left(\frac{\alpha}{2}, \frac{\alpha\lambda}{\alpha + \lambda}\right)$$

Consequently, the infinite horizon policy is a switching policy as well.

The proof is outlined in the appendix. The theorem implies that the optimal stationary policy is described by a deadbeat zone around the current fusion center location that does not switch the fusion center until the track has moved far enough from the fusion center. Note that this strategy is a consequence of our choice of ℓ_1 penalty on communication costs, representative of multihop networks.

A key assumption in the above model is that the track position at each stage is perfectly observed, so x_k is known. In general, we only have an estimate of the target position with a non-zero error covariance. We can incorporate this estimation error heuristically into the solution by increasing the process noiseat each stage. Although the density of the noise w_k may no longer be uniformwith finite support, it will still be zero-mean and symmetric. Those properties are sufficient to establish that the optimal strategy and the cost-to-go will be of the form

Lemma 3.1: The n-stage optimal policy is

$$\begin{aligned} u &= 0,\ |z| \leq \Gamma_n \\ u &= (|z| - \Gamma_n)\mathrm{sign}(z),\ |z| \geq \Gamma_n \end{aligned}$$

The corresponding n-stage cost-to-go is given by:

$$J_{N-n}(z) = \begin{cases} z^2 + \mathbf{E}(J_{N-n+1}(z + w + w_0)) & |z| \leq \Gamma_n \\ \lambda(|z| - \Gamma_n) + z^2 + \mathbf{E}J_{N-n+1}(\Gamma_n + w + w_0) & |z| > \Gamma_n \end{cases}$$

The proof is same as that outlined in the appendix for proof of theorem 3.1. However, since the noise is not uniform, we cannot guarantee that $\Gamma_n > 0$.

In two dimensions the problem can be treated separately in each dimension, as the $N-1$ step cost-to-go function decouples as the sum of the errors in the two dimensions. Note that, for the case when the error covariances are different in different dimensions, we get *asymmetric* regions of switching in the two dimensions, i.e., the switching regions are *rectangular* rather than square.

B. Target motion with velocity dynamics

Now consider the target motion described by the following dynamical system

$$\mathbf{X}_{k+1} = \begin{bmatrix} x_{k+1} \\ v^x_{k+1} \\ y_{k+1} \\ v^y_{k+1} \end{bmatrix} = \mathbf{A} \begin{bmatrix} x_k \\ v^x_k \\ y_k \\ v^y_k \end{bmatrix} + \mathbf{G}\mathbf{w}_k$$

$$\mathbf{A} = \begin{bmatrix} 1 & dt & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & dt \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{G} = \mathbf{I}_{4\times 4}$$

where dt is the sampling time. The process noise $w(k)$ has variance given by

$$\mathbf{Q} = q \begin{bmatrix} dt^3/3 & dt^2/2 & 0 & 0 \\ dt^2/2 & dt & 0 & 0 \\ 0 & 0 & dt^3/3 & dt^2/2 \\ 0 & 0 & dt^2/2 & dt \end{bmatrix}$$

where q is a scaling constant. The above is a standard discrete time model derived from continuous time target dynamics. Assume that the leader node position is given by ℓ_x, ℓ_y. The joint leader node and the target dynamics are

$$\begin{bmatrix} \mathbf{X}_{k+1} \\ \ell^x_{k++1} \\ \ell^y_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{X}_k \\ \ell^x_k \\ \ell^y_k \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ u^x_k \\ u^y_k \end{bmatrix} + \begin{bmatrix} \mathbf{w}_k \\ 0 \\ 0 \end{bmatrix}$$

As in the previous case identify $z^1_k = x_k - \ell^x_k$ and $z^2_k = y_k - \ell^y_k$. Then,

$$\begin{bmatrix} z^1_{k+1} \\ v^x_{k+1} \\ z^2_{k+1} \\ v^y_{k+1} \end{bmatrix} = \mathbf{A} \begin{bmatrix} z^1_k \\ v^x_k \\ z^2_k \\ v^y_k \end{bmatrix} - \begin{bmatrix} u^x_k \\ 0 \\ u^y_k \\ 0 \end{bmatrix} + \mathbf{w}_k$$

Note that $\mathbf{w}_k$ is independent of z^1_k, z^2_k and u^x_k, u^y_k and is zero mean. The objective is

$$(Q) : \min_{\mathbf{u}_1,..,\mathbf{u}_{N-1}} \mathbf{E}\left(\sum_{k=1}^{N-1} \lambda||\mathbf{u}_k||_1 + ||\mathbf{z}_k||^2 + ||\mathbf{z}_N||^2\right)$$

where $\mathbf{u}_k = [u^x_k, u^y_k]'$ and $\mathbf{z}_k = [z^1_k, z^2_k]'$. The cost-to-go at stage $N-1$ is

$$J_{N-1}(\mathbf{z}, \mathbf{v}) = \min_{\mathbf{u}} \left\{\lambda||u||_1 + ||z||^2 + ||\mathbf{z} + \mathbf{v} - \mathbf{u}||^2\right\} + \sigma^2_{\mathbf{w}}$$

where $\sigma^2_{\mathbf{w}} = \mathbf{E}(w(1)^2 + w(2)^2 + w(3)^2 + w(4)^2)$; and $w(i)$ are the components of the noise vector $\mathbf{w}$. The optimization problem in the two dimensions is decoupled, so we will focus on describing the one-dimensional problem instead. For simplicity, we drop the use of superscripts and refer to variables z, v and u for a given coordinate. For the motion in x direction we have

$$J_{N-1}(z, v) - \tfrac{\sigma^2_{\mathbf{w}}}{2} = \min_u \left\{\lambda|u| + (z)^2 + (z + v - u)^2\right\}$$

As in the zero velocity case, the optimal control satisfies $u_{N-1} = 0$ if $|z_{N-1} + v_{N-1}| \leq 0.5\lambda$. Outside this region the optimal control is given by,

$$u_{N-1} = z_{N-1} + v_{N-1} - \tfrac{1}{2}\mathrm{sign}(z_{N-1} + v_{N-1})$$

A similar policy holds for the optimal control in the y direction. The resulting one-dimensional cost-to-go function, at stage $N-1$ is

$$J_{N-1}(z, v) = z^2 + \frac{1}{2}\sigma^2_{\mathbf{w}} + \begin{cases} |z + v| - 1/4 & |z + v| > 0.5 \\ (z + v)^2 & |z + v| \leq 0.5 \end{cases}$$

At stage $N-2$ the cost-to-go is

$$J_{N-2}(z, v) = \min_u [|u| + |z|^2 + \mathbf{E}_w J_{N-1}(z + v + w(1) - u, v + w(2))]$$

From the above minimization it is clear that the cost-to-go is jointly convex in z, v. Consider the cost to go function at stage $N-n$. By recursion, we get

$$J_{N-n}(z, v) = \min_u [|u| + z^2 + E_w J_{N-n+1}(z + v - u + w(1), v + w(2))]$$

Lemma 3.2: If the joint density function $p(w(1), w(2))$ is symmetric in $(w(1), w(2))$ and $J_{N-n+1}(z, v) = J_{N-n+1}(-z, -v)$, then $J_{N-n}(z, v) = J_{N-n}(-z, -v)$.

Proof: The minimization remains unaltered if we substitute for $u = -u$, $w(1) = -w(1)$ and $w(2) = -w(2)$ and noting that $J_{N-n+1}(z, v) = J_{N-n+1}(-z, -v)$. ■

Since $J_{N-1}(z, v) = J_{N-1}(-z, -v)$, lemma 3.2 implies that $J_{N-n}(z, v) = J_n(-z, -v)$ for all n.

Let $z + v - u = \eta$. Then at stage $N-n$, we have the subgradient set for zero control

$$\frac{\partial J_{N-n}}{\partial u}|_{u=0} = \left\{[-\lambda, \lambda] - \frac{\partial}{\partial \eta}\mathbf{E}_w J_{N-n+1}(\eta + w(1), v + w(2))\right\}_{u=0}$$

Since $J_{N-n+1}(.,.)$ is a convex function in both arguments, $\mathbf{E}_w J_{n-1}$ is also convex. For a convex function with left and the right derivatives almost everywhere, [7]. Furthermore, the derivative exists almost everywhere, so we can exchange the expectation and the derivative operation. Thus we have,

$$\frac{\partial J_{N-n}}{\partial u}|_{u=0} = [-\lambda, \lambda] - \mathbf{E}_w \left\{\frac{\partial}{\partial \eta} J_{N-n+1}(\eta + w(1), v + w(2))\right\}_{u=0}$$

The "deadzone" region, i.e. the region of zero control is given by,

$$\Gamma_n(v) = \left\{ \eta : |\mathbf{E}_w \left\{ \frac{\partial}{\partial \eta} J_{N-n+1}(\eta + w(1), v + w(2)) \right\}_{u=0} | < \lambda \right\}$$

Note that it is a function of the velocity v. Simplifying as in the Appendix can show that there is a sufficiently small threshold region, Γ_n such that the optimal policy is to switch the fusion center only when $|z+v| \geq \Gamma_n$ and not switch otherwise. Furthermore, the optimal control is when there is switching is $u = z + v - \Gamma_n \text{sign}(z+v)$.

IV. Experiments

To illustrate the performance of the CEC algorithms described above, we must map the solution of the continuous sensor field approximation to a discrete grid of sensors, and evaluate the relative performance of the algorithms as compared with alternatives such as the greedy algorithm described in Section II. We do so for two classes of dynamic target models: the random walk model and the model with velocity dynamics.

A. Random Walk

We consider a network of 225 identical sensors laid out in a uniform grid on integers coordinates in $[-7,7] \times [-7,7]$. We assume that the process noise is uniformly distributed on [-0.5,0.5]. The target is initially at [0,0] with small initial covariance in each direction. The observation model at each sensor is given by

$$z(x_p, y_p) = \frac{20}{1 + \sqrt{2}(((x-x_p)^2 + (y-y_p)^2))} + v$$

where v is Gaussian noise of variance 1 and x_p and y_p are the coordinates of the sensor. We collect measurements at time intervals $dt = 1$. The information provided by selected sensors is processed by an Extended Kalman Filter (EKF) to generate state estimates for 50 time steps.

To implement the CEC algorithm in this example, we use the EKF updated state estimate as corresponding to the true target position (certainty equivalence). The fusion center is constrained to be at one of the sensors. We use a threshold along both the x and y directions and compute the distances along each direction between the updated estimate and the current fusion center. If the difference exceeds the threshold, we compute the control u generated by the CEC control, and round the continuous fusion center location to the nearest integer point to select the next fusion location. The reporting sensors were selected as the four sensors that surrounded the continuous fusion center location. The threshold is a function of the weight λ on communication costs and the look-ahead horizon; varying this weight yields a performance curve in terms of communications cost versus tracking error.

As reference algorithms, we use two algorithms: the one-step lookahead combinatorial algorithm in (4) with a variable weight λ to obtain a similar tradeoff between tracking error and communications. The second algorithm was a non-causal algorithm corresponding to a hindsight policy where the set of four reporting sensors are selected based on the true future target position (and hence had optimal observability) and the fusion center was selected as the sensor in this set of four that had the smallest switching communication cost. This provides a bound on the best estimation accuracy obtainable.

Figure 1 shows the tradeoff between the communication cost and the rms tracking error for the two causal algorithms varying the weight on the communication cost in the performance objective, as well as the performance of the hindsight optimization algorithm, averaged over 50 Monte Carlo runs, tracking over 50 time steps per Monte Carlo run. The results illustrate the advantages of CEC control over a one-step combinatorial policy, achieving smaller tracking error for comparable communication costs.

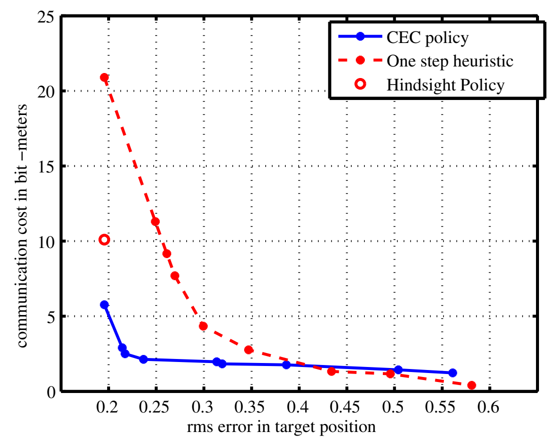

Fig. 1. Communication costs vs tracking error performance for different algorithms in random walk experiments

B. Target motion with velocity

In this case, target motion can cover additional area, so we need a larger sensor grid. For these experiments, we use a uniform two dimensional grid in $[-45,45] \times [-45,45]$, where the sensors are placed at a distance of $\rho = 3$ meters apart. The measurement model is given by

$$y(x_p, y_p) = \frac{20}{1 + (((X-x_p)^2 + (Y-y_p)^2))/10} + v$$

where v is a Gaussian noise of variance 1 and x_p and y_p are the x and y coordinates of the sensor. Note that each sensor has an effective sensing range that overlaps its adjacent sensors. For the target dynamics $q = 0.01$ and $dt = 1$. The target starts at $(0,0)$ with a velocity of 0.25 in both x and y directions.

In this experiment, 6 sensors around the leader node were selected to report their observations in each of the cases. Specifically, the 3 adjacent sensors north of the leader position and the 3 sensors south of the leader position report. The results averaged over 20 Monte Carlo runs and 50 time steps are shown in figure 2. As before, the CEC control achieves smaller RMS error for comparalble communication cost, with tracking performance close to the bound provided by the hindsight policy with significanlty smaller communications.

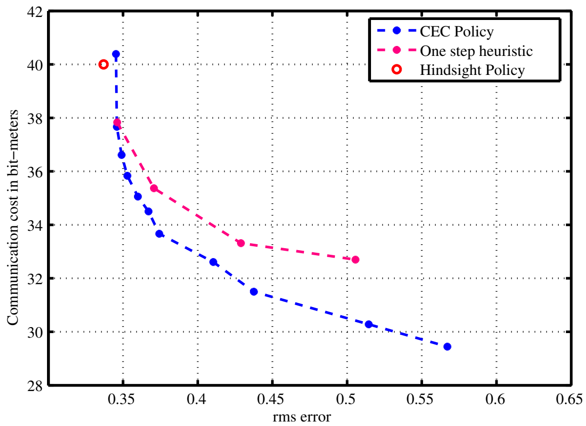

Fig. 2. Communication costs vs tracking error performance for different algorithms in velocity model experiments

V. Conclusions

In this paper we studied the problem of locating dynamically a fusion center in a sensor network for distributed target tracking that trades off communication costs and estimation error. For the case of regular networks with homogeneous sensors, we developed a continous approximation as a stochastic control problem and characterized the optimal policies for fusion center location in this approximation. The optimal policies have a dead-zone region where the fusion center does not change until the track's predicted position is far enough. The feedback solution of the continuous approximation can be used with a certainty equivalence approach to locate the fusion center in discrete fields of homogeneous sensors. In simulation, the resulting fusion center control provides improved performance versus a one-step policy based on discrete sensor locations, in terms of reduced communication costs for comparable accuracy. Furthermore, the on-line computation requirements for our feedback policy are minimal.

The main limitations of our results are that we assume the sensors are homogeneous, and that there is enough information for accurate tracking so that certainty equivalence control is adequate. Future work will explore to what extent these assumptions can be relaxed.

VI. Appendix

A. Outlilne of Proof of theorem 3.1

The complete proof and additional results can be found in the report [8]. We outline the major steps here: Let n denote the number of stages left. Then,

$$J_n(z) = \min_u \left\{\lambda|u| + z^2 + \mathbf{E}(J_{n-1}(z+w-u))\right\} \quad (7)$$

The following properties hold for the n-stage cost-to-go function:

Lemma 6.1: The n-stage cost-to-go function, $J_n(\cdot)$ is real valued, positive, convex, & symmetric about $z=0$ and monotonic for $z \geq 0$.

The proof follows by induction, and exploits the fact that the noise w has a symmetric distribution.

Lemma 6.2: The n-stage cost-to-go, $J_n(\cdot)$ is differentiable almost everywhere (with respect to lebesgue measure) and

$$\frac{\partial}{\partial \eta}\mathbf{E}J_n(\eta+w) = \mathbf{E}\frac{\partial}{\partial \eta}J_n(\eta+w) = \frac{1}{2\alpha}(J(\eta+\alpha)-J(\eta-\alpha))$$

This follows from the convexity property, and by direct computation.

Lemma 6.3: The n-stage optimal policy is a switching policy, i.e.,

$$\begin{aligned} u &= 0, \ |z| \leq \Gamma_n \\ u &= (|z|-\Gamma_n)\mathrm{sign}(z), \ |z| \geq \Gamma_n \end{aligned}$$

where, $\Gamma_n > 0$. The corresponding n-stage cost-to-go is described by:

$$J_n(z) = \begin{cases} z^2 + \mathbf{E}(J_{n-1}(z+w)) & |z| \leq \Gamma_n \\ \lambda\left(|z|-\Gamma_n\right) + z^2 + \mathbf{E}J_{n-1}(\Gamma_n + w) & |z| > \Gamma_n \end{cases}$$

The necessary and sufficient condition for optimality is that the sub-gradient set contain the element zero [6]. The optimality of switching policy follows from the fact that zero is an interior point of the sub-gradient of the argument in the minimization problem of Equation 7.

To establish that Γ_n is bounded away from zero independent of n, let $\Gamma_* = \inf_n \Gamma_n$.

Lemma 6.4:

$$\Gamma_* \geq \Gamma_0 = \min\left(\frac{\alpha}{2}, \frac{\alpha\lambda}{\alpha+\lambda}\right)$$

Assume there is an n such that, $\Gamma_n \leq \alpha/2$. If not then the result is established. Consider, for $0 \leq z \leq \Gamma_0$

$$\begin{aligned} &\tfrac{1}{2\alpha}(J_n(\alpha+z) - J_n(\alpha-z)) \\ &\overset{(a)}{=} \tfrac{1}{2\alpha}\left(\lambda|\alpha+z| + (\alpha+z)^2 - \lambda|\alpha-z| - (\alpha-z)^2\right) \\ &= z\tfrac{\alpha+\lambda}{\alpha\lambda} \leq \lambda, \ \forall\, z \leq \Gamma_0 \end{aligned}$$

where, (a) follows from Lemma 6.3 and noting the fact that $\alpha - \Gamma_n \geq \Gamma_n$. Therefore, whenever $z \leq \Gamma_0$ the sub-gradient set always contains the zero element for the zero policy, i.e., $u = 0$. This implies that $\Gamma_n \geq \Gamma_0$.

References

[1] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration for tracking applications," *IEEE Signal Processing Magazine*, vol. 19, no. 2, pp. 61–72, March 2002.

[2] J. Liu, J. Reich, and F. Zhao, "Collaborative in-network processing for target tracking," *EURASIP Journal on Applied Signal Processing*, vol. 4, pp. 378–391, March 2003.

[3] C. Kreucher, A. Hero, K. Kastella, and D. Chang, "Efficient methods of non-myopic sensor management for multitarget tracking," in *The Proceedings of the 43rd IEEE Conference on Decision and Control*, vol. 1, Dec 2004, pp. 722–727.

[4] J. Williams, J. F. III, and A. Willsky, "An approximate dynamic programming approach for communication constrained inference," *Proc. IEEE Workshop on Statistical Signal Processing*, July 2005.

[5] A. S. Chhetri, D. Morrell, and A. Papandreou-Supapola, "Sensor scheduling using a 0-1 mixed integer programming framework," in *Proc. 2006 IEEE Workshop on Sensor Array and Multichannel Processing*, July 2006.

[6] R. T. Rockafellar, *Convex Analysis*. Princeton University Press, Princeton NJ, 1970.

[7] H. Royden, *Real Analysis*, 3rd ed. Prentice Hall, 1988.

[8] S. Aeron, V. Saligrama, and D. A. Castañón, "Energy efficient policies for distributed target tracking in multi-hop sensor networks," ECE Department, Boston University, Tech. Rep., September 2006.