

Mutual Information Methods with Particle Filters for Mobile Sensor Network Control [†]

Gabriel M. Hoffmann, Steven L. Waslander, and Claire J. Tomlin

Department of Aeronautics and Astronautics

Stanford University, Stanford, CA 94305

{gabeh, stevenw, tomlin}@stanford.edu

Abstract—This paper develops a set of methods enabling an information-theoretic distributed control architecture based on particle filters to facilitate search by a mobile sensor network, permitting the use of nonlinear and non-Gaussian sensor models. Given a particular configuration sensors, this technique exploits the structure of the probability distributions of the target state and of the sensor measurements to compute the control inputs to the mobile sensors leading to future observations that minimize, in expectation, the future uncertainty of the target state. We compute the mutual information using a particle set representation of the posterior distribution. In order to control a large number of mobile sensors as a network, single-node and pairwise-node approximation schemes are presented, with analytically bounded error, making the approach scalable to increasing network sizes, while still planning cooperatively. The methods are applied in simulation to bearings-only sensing, and to localizing an avalanche rescue beacon of a buried victim, using transceivers on quadrotor aircraft to measure the magnetic field. Monte Carlo simulations also demonstrate that as network size increases, the sensors more quickly localize the target, and the pairwise-node approximation results in superior performance to the single-node approximation.

I. INTRODUCTION

Mobile sensor network technology is becoming increasingly available, making it possible to exploit these platforms to efficiently and automatically search for targets with unprecedented speed, safety, and reliability. Mobile sensor network search is complicated by low prior information about the target’s state, the frequently nonlinear mapping between sensor observations and the physical world, and non-Gaussian target state representation.

This paper addresses these challenges by developing a set of methods enabling an information-theoretic distributed control architecture based on particle filters. We derive methods to compute information-theoretic quantities using particle filters, and approximations that enable scalability in network size for real-time control. Given a particular configuration of sensors, this technique exploits the structure of the probability distributions of the target state and of the sensor measurements to compute control inputs leading to future observations that minimize, in expectation, the future uncertainty of the target state. The approximation schemes presented have polynomial time complexity in the network size, and analytically bounded error. The *single-node* approximation leads to cooperative control due to common knowledge of the target state distribution, whereas

the *pairwise-node* approximation makes improved cooperation possible by approximating the effect of other sensing nodes’ planned observations. By using particle filters, we are able to numerically approximate the target state probability distribution, without assumptions about the family of the distribution or about the linearity of the dynamics[1].

In previous literature, one method minimizes metrics of the expected estimation covariance to compute trajectories to help estimation [2], [3]. Though computationally efficient, the approach linearizes measurement models, and relies on a Gaussian noise assumption and a guessed initial solution. This can lead to underestimation of the covariance and divergence of the filter [1], [4]. Another method, using grid cell discretization, is developed in [5], enabling use of nonlinear, non-Gaussian models. These techniques suffer poor computational time complexity with the dimension of the state, particularly for information theoretic quantities. Information theoretic costs have been used to manage sensors [6], and in [7] the expected alpha-divergence of a particle filter distribution was used for sensor management, and specialized to select modes for binary sensors, though scalability in sensor network size is not addressed.

We proceed by showing equivalence of searching for a target and maximizing mutual information. Then, algorithms to compute mutual information using particle filters are then derived, and used in a distributed control algorithm to cooperatively control the vehicles. The methods are applied in simulation to localization using bearings-only sensors on airplanes, and to localization of a buried avalanche victim, using rescue beacons on quadrotor aircraft from the STARMAC project [8]. The results show successful localization under all circumstances. The pairwise-node approximation leads to faster localization than using the single-node approximation, and increasing the size of the network speeds localization. Successful localization of the rescue beacon’s magnetic field source exemplifies the utility of these techniques.

II. PROBLEM FORMULATION

A. Mobile Sensor Motion and Measurement Models

Consider a set of n_v vehicles carrying sensors to locate a target in the search domain, Θ . The location of the target, $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$, is unknown to the vehicles. A motion model is used for non-stationary targets [9]. A prior distribution, $p_0(\theta)$, is provided, using any information available *a priori*.

The state of the i^{th} vehicle, at time index t , is $\mathbf{x}_t^{(i)} \in \mathbb{R}^{n_s}$, with n_s vehicle states. The discrete time dynamics are

$$\mathbf{x}_{t+1}^{(i)} = f_t^{(i)}(\mathbf{x}_t^{(i)}, \mathbf{u}_t^{(i)}) \quad (1)$$

[†] This research was supported by ONR under the CoMotion MURI contract N00014-02-1-0720, and by the NASA Joint University Program.

where $\mathbf{u}_t^{(i)} \in U^{(i)} \subset \mathbb{R}^{n_u}$ is the set of n_u control inputs, $U^{(i)}$ is their domain, and the time step is Δ . A minimum separation of \bar{d} is required between all vehicles. Let $x_t^{(i)}$ be the subset of the i^{th} vehicle's states that correspond to its position. The collision avoidance constraint is,

$$\|x_t^{(i)} - x_t^{(j)}\| \geq \bar{d} \quad \forall j \in \{1, \dots, n_v : j \neq i\} \quad (2)$$

Sensor measurements for the i^{th} vehicle, $\mathbf{z}_t^{(i)} \in Z^{(i)} \subset \mathbb{R}^{n_z}$, are taken at rate $\frac{1}{\Delta}$. The domain of observations is $Z^{(i)}$, with dimension n_z . The measurement model is,

$$\mathbf{z}_t^{(i)} = h_t^{(i)}(\mathbf{x}_t^{(i)}, \theta, \eta_t^{(i)}) \quad (3)$$

provided *a priori* to all vehicles. The observation noise is $\eta_t^{(i)} \in \mathbb{R}^{n_\eta}$ with assumed distribution, $\eta_t^{(i)} \sim p_t(\eta^{(i)})$. The noise distribution need not be Gaussian. The problem formulation admits a broad class of measurement models, as $h_t^{(i)}$ could be a nonlinear or discontinuous mapping of the states and measurement noise onto the observation space.

The vehicles are equipped with communication devices such that each vehicle can incorporate the measurements made by all other vehicles in the estimation process at each time step. One such reliable technology is demonstrated in the 802.11b network used in the STARMAC project.

B. Information Seeking

Each vehicle, i , maintains its own minimum mean square error (MMSE) estimate of the current target state, $\hat{\theta}_t^{(i)} \in \mathbb{R}^{n_\theta}$, and of the posterior distribution, $p_t(\theta^{(i)})$. The posterior distribution is the prior distribution conditioned on all new observations that vehicle has made or received, i.e., $p_{t-1|t}(\theta^{(i)}|\mathbf{z})$.¹ These distributions, at each aircraft, based on all shared observations, can be assumed nearly identical.

The main goal of a multi-vehicle search team is to minimize the total time to localize a target with a given likelihood, by making observations at some fixed rate. Therefore, the goal is to *control sensor locations to minimize the expected number of future observations needed to ascertain the target's state*. A set of observations can be interpreted, in an information-theoretic sense, as a code word, with an alphabet comprised of all possible quantized outputs of the n_z sensors. These encode the target state, which is represented computationally by an alphabet of a finite number of symbols, such as 64 bits in a double precision floating point data type. Therefore, to minimize the expected number of remaining observations is to maximize the expected log-likelihood of the posterior distribution with each observation of the vehicles, as derived in [10]. With constraints on control authority, the optimal observation positions may not be reachable in one time step, leading to a greedy algorithm that moves the vehicles toward the best reachable locations.

In order to compute the posterior probability distribution following an observation, assuming that the target state represents the complete state of the target and the system is Markov, Bayes' Rule can be used,

¹The subscript of probability distribution functions indicates the time steps of the random variable arguments. For instance, $p_{t+1|t,t+1}(a|b,c;d)$ is the distribution of random variable a at time $t+1$ conditioned on random variables b at time t and c at time $t+1$, parameterized by d .

$$p_{t+1}(\theta^{(i)}) = \frac{p_t(\theta^{(i)})p_{t+1|t}(\mathbf{z}|\theta^{(i)})}{\int_{\Theta} p_t(\theta^{(i)})p_{t+1|t}(\mathbf{z}|\theta^{(i)})d\theta^{(i)}} \quad (4)$$

Taking the log-likelihood of each side yields

$$H(\theta_t^{(i)}|\mathbf{z}_{t+1}) = H(\theta_t^{(i)}) - I(\mathbf{z}_{t+1};\theta_t^{(i)}) \quad (5)$$

where

$$H(\theta_t^{(i)}) = - \int_{\theta \in \Theta} p_t(\theta^{(i)}) \log p_t(\theta^{(i)}) d\theta \quad (6)$$

$$H(\theta_t^{(i)}|\mathbf{z}_{t+1}) = - \int_{\substack{\theta \in \Theta \\ \mathbf{z} \in Z}} p_{t,t+1}(\theta^{(i)}, \mathbf{z}) \log p_{t|t+1}(\theta^{(i)}|\mathbf{z}) d\theta d\mathbf{z} \quad (7)$$

$$I(\mathbf{z}_{t+1};\theta_t^{(i)}) = \int_{\substack{\theta \in \Theta \\ \mathbf{z} \in Z}} p_{t,t+1}(\theta^{(i)}, \mathbf{z}) \log \frac{p_{t,t+1}(\theta^{(i)}, \mathbf{z})}{p_{t,t+1}(\theta^{(i)})p_{t+1}(\mathbf{z})} d\theta d\mathbf{z} \quad (8)$$

$H(\theta_t^{(i)})$ is the *entropy* of the target state distribution, $I(\mathbf{z}_{t+1};\theta_t^{(i)})$ is the *mutual information* between the distributions of the target state and the sensors, and $H(\theta_t^{(i)}|\mathbf{z}_{t+1})$ is the *conditional entropy* of the distribution — the expected entropy of the target state when conditioning with \mathbf{z}_{t+1} [11]. The entropy of a probability distribution is a metric of the uncertainty of that distribution. The mutual information is a metric of the expected divergence (Kullback-Liebler) between the independent and joint distributions of $\theta_t^{(i)}$ and \mathbf{z}_{t+1} . It is large when two distributions have strong interdependence, and zero when they are independent.

The control inputs, \mathbf{u}_t , and vehicle states, \mathbf{x}_t , influence the observations, \mathbf{z}_{t+1} , through Equations (1) and (3). To minimize the expected future uncertainty of the target state distribution, with respect to \mathbf{u}_t , is to minimize Equation (5). The actual uncertainty depends on the real measurement made of \mathbf{z}_{t+1} . The prior uncertainty is independent of the future control inputs, so to minimize the expected posterior uncertainty, one must maximize the observation information with respect to the control inputs. It expands as,

$$\underbrace{I(\mathbf{z}_{t+1};\theta_t^{(i)})}_{\text{observation information}} = \underbrace{H(\mathbf{z}_{t+1})}_{\text{observation uncertainty}} - \underbrace{H(\mathbf{z}_{t+1}|\theta_t^{(i)})}_{\text{conditional observation uncertainty}} \quad (9)$$

Thus, minimizing the expected posterior uncertainty is equivalent to maximizing the difference between the uncertainty that any particular observation will be made, and the uncertainty of the measurement model. In order to seek information, the network computes its control inputs by maximizing the *mutual information utility function*,

$$V^{(i)}(\mathbf{x}_t, \mathbf{u}_t, \theta_t^{(i)}) = I(\mathbf{z}_{t+1};\theta_t^{(i)}) \quad (10)$$

as described in the next section.

III. METHODOLOGY

A. Particle Filter

Each vehicle approximates $p_t(\theta^{(i)})$ with an onboard particle filter, incorporating the observations shared by all vehicles, with a set of N particles, $(\tilde{\theta}_{t,k}^{(i)}, \mathbf{w}_{t,k}^{(i)})$, indexed by k , where $\tilde{\theta}_{t,k}^{(i)} \in \Theta$ is the state of the particle, and $\mathbf{w}_{t,k}^{(i)} \in \mathbb{R}_+$

is the likelihood weight.² The particles approximate $p_t(\theta^{(i)})$ by the probability mass function,

$$\hat{p}_t(\theta^{(i)}) = \sum_{k=1}^N \mathbf{w}_{t,k}^{(i)} \delta(\theta^{(i)} - \tilde{\theta}_{t,k}^{(i)}) \quad (11)$$

where $\delta(\cdot)$ is the Dirac delta function.

By maintaining a set of particles locally aboard each vehicle, only the observations need to be communicated, as opposed to the values of the entire set of particles. The particle filter iteratively incorporates new observations by predicting the state of each particle, updating the likelihood weights with the likelihood of new observations, and then resampling the particles, as described in detail in [12], [9]. A low variance sampler [4] is used for this work, with time complexity of $\mathcal{O}(N)$. The MMSE estimate is

$$\hat{\theta}_t^{(i)} = \int_{\Theta} \theta^{(i)} p_t(\theta^{(i)}) d\theta^{(i)} \approx \sum_{k=1}^N \mathbf{w}_{t,k}^{(i)} \tilde{\theta}_{t,k}^{(i)} \quad (12)$$

B. Determining Mutual Information from Particle Sets

To evaluate the mutual information utility function, Equation (10), with the particle filter representation, expand it according to Equation (9). Consider the first term,

$$H(\mathbf{z}_{t+1}) = - \int_{\mathcal{Z}} p_{t+1}(\mathbf{z}; \mathbf{u}_t, \mathbf{x}_t) \log p_{t+1}(\mathbf{z}; \mathbf{u}_t, \mathbf{x}_t) d\mathbf{z} \quad (13)$$

This cannot be directly evaluated, because $p_{t+1}(\mathbf{z}; \mathbf{u}_t, \mathbf{x}_t)$ is not available as a continuous function. Rather, the observation distribution must be determined from the particle set and sensor model. First, expand the distribution as,

$$\begin{aligned} p_{t+1}(\mathbf{z}; \mathbf{u}_t, \mathbf{x}_t) &= \int_{\Theta} p_{t+1|t}(\mathbf{z}|\theta^{(i)}; \mathbf{u}_t, \mathbf{x}_t) p_t(\theta^{(i)}) d\theta^{(i)} \\ &= \int_{\Theta} \left(p_t(\theta^{(i)}) \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \right) d\theta^{(i)} \end{aligned} \quad (14)$$

where the second step expands the joint distribution using the assumption that the likelihoods of sensor observations are conditionally independent given the target state distribution. This is exact when the sensors lack correlated noise. Then, Monte Carlo integration techniques can be used[13]. Substituting Equation (11) into Equation (14) yields the particle filter approximation at the i^{th} vehicle,

$$\begin{aligned} p_{t+1}(\mathbf{z}; \mathbf{u}_t, \mathbf{x}_t) &\approx \\ &\sum_{k=1}^N \left(\mathbf{w}_{t,k}^{(i)} \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}=\tilde{\theta}_{t,k}^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \right) \end{aligned} \quad (15)$$

Substituting Equation (15) back into Equation (13) yields the observation entropy of the distribution represented by the particle filter approximation at the i^{th} vehicle,

$$\begin{aligned} H(\mathbf{z}_{t+1}) &\approx \\ &- \int_{\mathcal{Z}} \left\{ \left(\sum_{k=1}^N \left(\mathbf{w}_{t,k}^{(i)} \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}=\tilde{\theta}_{t,k}^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \right) \right) \right. \\ &\cdot \log \left(\sum_{k=1}^N \left(\mathbf{w}_{t,k}^{(i)} \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}=\tilde{\theta}_{t,k}^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \right) \right) \left. \right\} d\mathbf{z} \end{aligned} \quad (16)$$

²The second subscript of any variable denotes the index of the particle to which the variable belongs.

which can then be evaluated by numerical quadrature [14].

Next, similar methods are applied to compute

$$H(\mathbf{z}_{t+1}|\theta_t) = - \int_{\mathcal{Z}, \Theta} p_{t+1,t}(\mathbf{z}, \theta^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \log p_{t+1|t}(\mathbf{z}|\theta^{(i)}; \mathbf{u}_t, \mathbf{x}_t) d\mathbf{z} d\theta^{(i)} \quad (17)$$

The joint distribution can be expanded using the chain rule and assuming conditional independence,

$$p_{t+1,t}(\mathbf{z}, \theta^{(i)}; \mathbf{u}_t, \mathbf{x}_t) = p_t(\theta^{(i)}) \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \quad (18)$$

Substituting Equation (18) into Equation (17) and applying the approximation of Equation (11), yields the conditional observation entropy of the distribution represented by the particle filter approximation at the i^{th} vehicle,

$$\begin{aligned} H(\mathbf{z}_{t+1}|\theta_t^{(i)}) &\approx \\ &- \int_{\mathcal{Z}} \sum_{k=1}^N \left\{ \mathbf{w}_{t,k}^{(i)} \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}=\tilde{\theta}_{t,k}^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \right. \\ &\cdot \log \left. \prod_{j=1}^{n_v} p_{t+1|t}(\mathbf{z}^{(j)}|\theta^{(i)}=\tilde{\theta}_{t,k}^{(i)}; \mathbf{u}_t, \mathbf{x}_t) \right\} d\mathbf{z} \end{aligned} \quad (19)$$

Thus, the mutual information utility function, Equation (10), can be found by using Equations (16), (19), and (9). This optimization remains highly coupled between the n_v vehicles. The degree of cooperation between the vehicles is analyzed to determine a scalable control strategy.

C. Approximately Decoupling Mutual Information

The mutual information between the random variables $\theta_t^{(i)}$ and \mathbf{z}_{t+1} quantifies the expected reduction in uncertainty. However, the computational complexity of using a particle set representation to evaluate this quantity grows exponentially n_z , due to integration over each dimension.

We present two different approximations to mutual information that allow it to be evaluated in polynomial time with respect to the number of sensors. This makes the network scalable, yet capable of exploiting the descriptiveness of the particle filter. First, in the single-node approximation, a sensing node's utility function uses the previous observations of all sensing nodes, but only considers its own future observations. This is a generalization of an increasingly popular technique (e.g., [2]). Second, in the pairwise-node approximation, a sensing node additionally considers the effect of plans of each other vehicle, pairwise, on the utility of its own future observation.

Proposition 1: The mutual information utility function can be approximated using the *single-node approximation*:

$$V_s^{(i)}(\mathbf{x}_t, \mathbf{u}_t, \theta_t^{(i)}) = I(\mathbf{z}_{t+1}^{(i)}; \theta_t^{(i)}) \quad (20)$$

with an error from the true value of

$$\epsilon_s^{(i)} = c_s^{(i)} + \sum_{\substack{j=2 \\ j \neq i}}^{n_v} \left(I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(i)}, \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) \right) \quad (21)$$

where $c_s^{(i)}$ encompasses the terms that are constant with respect to the i^{th} mobile sensor's control inputs.

Proof: Without loss of generality, consider the case of approximating the mutual information from the perspective of vehicle $i = 1$. The mutual information can be expanded using the chain rule [11], and then rewritten by exchanging conditioning variables (see Appendix) to yield

$$\begin{aligned} I(\theta_t^{(1)}; \mathbf{z}_{t+1}) &= \sum_{j=1}^{n_v} I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(j)} | \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) \\ &= \sum_{j=1}^{n_v} \left(I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(j)}) - I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) \right. \\ &\quad \left. + I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \mathbf{z}_{t+1}^{(j)}) \right) \end{aligned}$$

Exchanging conditioning variables on the latter two terms, and canceling terms that sum to zero,

$$\begin{aligned} I(\theta_t^{(1)}; \mathbf{z}_{t+1}) &= V_s^{(1)}(\mathbf{x}_t, \mathbf{u}_t, \theta_t^{(1)}) + c_s^{(1)} \\ &\quad + \sum_{j=2}^{n_v} \left(I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \theta_t^{(1)}) \right. \\ &\quad \left. - I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) \right) \end{aligned}$$

with constant $c_s^{(i)} = - \sum_{\substack{j=1 \\ j \neq i}}^{n_v} I(\mathbf{z}_{t+1}^{(j)}; \theta_t^{(i)})$.

Applying the assumption that observations are conditionally independent given the target state, the first term in the summation is zero. Thus, generalizing to the i^{th} vehicle, the mutual information utility function can be evaluated using Equation (20) with error given by Equation (21). ■

The computational complexity of the single-node approximation to the mutual information utility function is constant with respect to n_v . To improve on this approximation, considering the pairwise interactions of all vehicles leads to a utility function approximation that more accurately captures the effect of group control inputs on mutual information, with computational complexity linear in n_v .

Proposition 2: The mutual information utility function can be approximated using the *pairwise-node approximation*:

$$\begin{aligned} V_p^{(i)}(\mathbf{x}_t, \mathbf{u}_t, \theta_t^{(i)}) &= (2 - n_v) \left(I(\mathbf{z}_{t+1}^{(i)}; \theta_t^{(i)}) \right) \\ &\quad + \sum_{\substack{j=1 \\ j \neq i}}^{n_v} \left(I(\mathbf{z}_{t+1}^{(i)}, \mathbf{z}_{t+1}^{(j)}; \theta_t^{(i)}) \right) \end{aligned} \quad (22)$$

when $n_v \geq 2$, with an error from the true value of

$$\epsilon_p^{(i)} = \sum_{\substack{j=2 \\ j \neq i}}^{n_v} \left(I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \mathbf{z}_{t+1}^{(i)}) \right) \quad (23)$$

Proof: Without loss of generality, consider the case of approximating the mutual information from the perspective of vehicle $i = 1$. The mutual information can be expanded using an application of the chain rule, separating the first term in the summation, and applying the chain rule again,

$$\begin{aligned} I(\theta_t^{(1)}; \mathbf{z}_{t+1}) &= I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(1)}) + \\ &\quad \sum_{j=2}^{n_v} \left(I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(1)}, \mathbf{z}_{t+1}^{(j)} | \mathbf{z}_{t+1}^{(2)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) \right. \\ &\quad \left. - I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(1)} | \mathbf{z}_{t+1}^{(2)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) \right) \end{aligned}$$

Exchanging conditioning variables in the summation (see Appendix), canceling the resulting terms that sum to zero, and splitting the remaining summation, yields

$$\begin{aligned} I(\theta_t^{(1)}; \mathbf{z}_{t+1}) &= V_p^{(1)}(\mathbf{x}_t, \mathbf{u}_t, \theta_t^{(1)}) \\ &\quad + \sum_{j=3}^{n_v} \left(I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(2)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \mathbf{z}_{t+1}^{(1)}, \mathbf{z}_{t+1}^{(j)}) \right. \\ &\quad \left. - I(\theta_t^{(1)}; \mathbf{z}_{t+1}^{(2)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \mathbf{z}_{t+1}^{(1)}) \right) \end{aligned}$$

Exchanging conditioning variables inside the summation, canceling terms summing to zero, and assuming conditional independence of observations given the target state yields,

$$I(\theta_t^{(1)}; \mathbf{z}_{t+1}) = V_p^{(1)}(\mathbf{x}_t, \mathbf{u}_t, \theta_t^{(1)}) + \epsilon_p^{(1)}$$

Thus, the mutual information utility function can be evaluated using Equation (22) with error given by Equation (23). ■

Now we examine the effect that this added computational complexity has on the error terms, as a function of the value of the i^{th} vehicle's control inputs.

Proposition 3: The magnitude of the error terms, that vary with the i^{th} vehicle's control inputs, in the single-node approximation, is greater than or equal to the magnitude of the pairwise-node approximation error terms, and equal only when the vehicle's observations are independent of all other vehicles. That is,

$$|\epsilon_s^{(i)} - c_s^{(i)}| \geq |\epsilon_p^{(i)}| \quad (24)$$

Proof: Subtract from the single-node approximation error, Equation (21), the terms that do not vary with the i^{th} vehicle's control inputs, $c_s^{(i)}$, and apply the chain rule for mutual information [11],

$$\begin{aligned} I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(i)}, \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) &= \\ I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(i)} | \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) &+ I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \mathbf{z}_{t+1}^{(i)}) \end{aligned} \quad (25)$$

Mutual information is always greater than or equal to zero, and is equal only if the distributions are independent. So,

$$\begin{aligned} I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(i)}, \mathbf{z}_{t+1}^{(1)}, \dots, \mathbf{z}_{t+1}^{(j-1)}) &\geq \\ I(\mathbf{z}_{t+1}^{(j)}; \mathbf{z}_{t+1}^{(2)}, \dots, \mathbf{z}_{t+1}^{(j-1)} | \mathbf{z}_{t+1}^{(i)}) &\end{aligned} \quad (26)$$

The magnitudes of the sum of the left and right sides of this equation, from $j = 1$ to $j = n_v$, are equal to, respectively, the left and right sides of Equation (24). ■

The pairwise-node approximation yields an estimate of the mutual information gain as good or better than the single-node approximation. When the vehicles' sensors make uncorrelated measurements, the single-node approximation computes faster and yields the same result as the pairwise-node approximation. However, if the observations are correlated, as is more frequently true, then the pairwise-node approximation yields a closer estimate. Although the magnitude of the pairwise-node approximation error is less, it is not possible to guarantee that the optimization is not skewed by some systematic error between the exact solution and the single-node error. However, using the pairwise-node approximation still yields an approximate expected mutual information surpassing what seemed possible using the single-node approximation, and in experiments, the pairwise-node approximation yields better results.

D. Mobile Sensor Network Control

The mobile sensor network control is structured as a set of local optimal control problems for each sensing node, coupled through interconnecting constraints. The local optimization problem is formed holding the actions of the other vehicles fixed, and an iterative algorithm is implemented based on [15] ensuring convergence to ϵ -feasible solutions that satisfy the necessary conditions for Pareto optimality.

An inexact penalty function is defined for each vehicle as,

$$P(\mathbf{x}_t^{(i)}, \mathbf{u}_t^{(i)} | \mathbf{x}_t^{(-i)}, \mathbf{u}_t^{(-i)}) = \sum_{m=1}^{n_c^{(i)}} \max(0, g^{(i,m)}(\mathbf{x}_t^{(i)}, \mathbf{u}_t^{(i)} | \mathbf{x}_t^{(-i)}, \mathbf{u}_t^{(-i)}) \gamma) \quad (27)$$

where m indexes the set of $n_c^{(i)}$ interconnecting inequality constraints, $g^{(i,m)}$ that affect vehicle i , and $-i$ is the set of all vehicles other than i . The penalty functions are zero wherever the constraints are satisfied, and are differentiable for all $\gamma \geq 2$, $\gamma \in \mathbb{R}_+$. The $n_v - 1$ collision avoidance constraints from Equation (2) are rewritten as,

$$g^{(i,m)} = \bar{d} - \|x_{t+1}^{(i)} - x_{t+1}^{(m)}\| \leq 0 \quad \forall m \in \{1, \dots, n_v : m \neq i\} \quad (28)$$

The penalty function is added to the individual vehicle cost scaled by penalty parameter β , varying the tradeoff between constraint violation and the information-theoretic cost, with an update $\beta := \alpha\beta$ at each iteration, where $\alpha \in (0, 1)$ is a design parameter. The local optimization problem based on the single-node approximation, Equation (20), is

Single-Node Local Optimization Program:

minimize $-V_s^{(i)}(\mathbf{x}_t^{(i)}, \mathbf{u}_t^{(i)}, \theta_t^{(i)} | \mathbf{x}_t^{(-i)}, \mathbf{u}_t^{(-i)}) + \frac{1}{\beta} P(\mathbf{x}_t^{(i)}, \mathbf{u}_t^{(i)} | \mathbf{x}_t^{(-i)}, \mathbf{u}_t^{(-i)})$

subject to $\mathbf{x}_{t+1}^{(i)} = f_t^{(i)}(\mathbf{x}_t^{(i)}, \mathbf{u}_t^{(i)})$
 $\mathbf{z}_{t+1}^{(i)} = h_t^{(i)}(\mathbf{x}_{t+1}^{(i)}, \theta_t^{(i)}, \eta_t^{(i)})$

(29)

The decentralized penalty method requires costs depending only on local variables, as in the single-node approximation. For the pairwise-node approximation, Equation (22), slack variables are added to decouple the local sensor costs, resulting in additional interconnecting constraints to include in the penalty function. Define the slack variable, $\tilde{\mathbf{u}}_t^{(i)}$, as the vector of all sensors' control inputs computed by the i^{th} sensor. Agreement among mobile sensing nodes on $\tilde{\mathbf{u}}_t^{(i)}$ is realized through the penalty function enforced constraint,

$$\tilde{\mathbf{u}}_t^{(i)} = \tilde{\mathbf{u}}_t^{(j)} \quad \forall i, j \in \{1, \dots, n_v : j \neq i\} \quad (30)$$

The pairwise-node local optimization program is defined analogously to the single-node program, with optimization instead over the entire control vector, $\tilde{\mathbf{u}}_t^{(i)} \in U$.

The distributed algorithm uses an iterative approach where interim solutions of control inputs are communicated amongst the vehicles in between local optimizations. This process can be hierarchical, synchronous, or asynchronous. Detailed algorithms and proof of convergence are in [15].




Fig. 1. Mean and quartile bars of the probability that the true target state is within 1 unit of the MMSE estimate, for sets of 1000 trials. The vehicles each make one observation at each time step. (a) Results for $n_v = 4$. Single-node approximation is dashed line, pairwise-node approximation is solid line. (b) Results using pairwise-node approximation for varying number of search vehicles: 2, 3, 4 and 10.

IV. APPLICATIONS

A. Bearings-Only Sensing

A bearings-only sensor, such as a camera, measures the bearing to the target only. The measurement equation is,

$$h_b^{(i)}(\mathbf{x}_t^{(i)}, \theta, \eta_t^{(i)}) = \delta_b^{(i)} + \eta_t^{(i)}$$

$$\mathbf{x}_t^{(i)} = \begin{pmatrix} x_t^{(i)} \\ y_t^{(i)} \end{pmatrix}, \theta = \begin{pmatrix} x_s \\ y_s \end{pmatrix}, \eta_t^{(i)} \sim \mathcal{N}(0, \sigma_s^2)$$

where $\delta_b^{(i)}$ is the angle of the ray from the i^{th} vehicle's position, $(x_t^{(i)}, y_t^{(i)})$, to the target's position, (x_s, y_s) . The additive Gaussian noise has variance σ_s^2 . The search domain, Θ , spans a unitless 40 by 40 square. The prior distribution is uniform over the domain. The vehicles are modeled as airplanes, with fixed speeds, limited turning rates, and initially close proximity.

The algorithm results in the rapid localization of the target, despite complete prior uncertainty, consistently in a large number of trials. Use of the pairwise-node approximation, Figure 1a, reduced time-to-convergence compared to the single-node approximation, on average. The pairwise-node approximation yielded more consistent performance compared to the single-node approximation, shown by narrower error bands. When applied to varying sizes of vehicle fleets, the algorithm successfully exploits the additional availability of agents, as shown in Figure 1b. The time-to-convergence is reduced, on average, as vehicles are added to the fleet.

B. Rescue Beacon Sensing

To rescue a victim buried in snow due to an avalanche, each vehicle uses a beacon receiver to measure the local orientation an electromagnetic (EM) field transmitted by the standard avalanche rescue beacon, carried by the victim. The position and orientation of the beacon are unknown.

The magnetic field, \mathbf{H} , of an EM source is modulated to minimize interference by snow and rocks. Measurements




Fig. 2. Four quadrotor aircraft search for a rescue beacon using the pairwise-node approximation. The particles are shown as short lines indicating position and orientation, with darkness indicating the particles' weights. The MMSE estimate is the black X, the red dots are the vehicle waypoint history, and the red square is the true target location. Its true orientation is $\frac{\pi}{4}$. The plots (a) through (d) are at time 1, 2, 5, and 20.

are made of the near-field at spherical coordinates (r, ψ, ϕ) relative to the antenna, with components[16],

$$H_r = \frac{\beta^3}{2\pi} Il \left[\frac{j}{(\beta r)^2} + \frac{1}{(\beta r)^3} \right] \cos(\psi) e^{-j\beta r} + \eta_{t,r}$$

$$H_\psi = -\frac{\beta^3}{4\pi} Il \left[\frac{1}{\beta r} - \frac{j}{(\beta r)^2} - \frac{1}{(\beta r)^3} \right] \sin(\psi) e^{-j\beta r} + \eta_{t,\psi}$$

where $\beta = \frac{2\pi}{\lambda}$ is a constant of the signal modulation, I is the antenna current, l is the equivalent length, and η_t is measurement noise. Measurements of signal strength from two orthogonal antennas in the receiver give a measured direction, δ , of the field line. The measurement model is

$$h_{beac}^{(i)}(\mathbf{x}_t^{(i)}, \theta, \mathbf{w}_t^{(i)}) = \delta(\mathbf{x}_t^{(i)}, \theta, \eta_t^{(i)})$$

$$\mathbf{x}_t^{(i)} = \begin{pmatrix} x_t^{(i)} \\ y_t^{(i)} \end{pmatrix}, \theta = \begin{pmatrix} x_a \\ y_a \\ \psi_a \end{pmatrix}, \eta_t^{(i)} \sim \mathcal{N}(0, \Sigma_r)$$

where $(x_t^{(i)}, y_t^{(i)})$ is the i^{th} vehicle's position, and (x_a, y_a) is the beacon antenna's position, with rotation angle ψ_a in the horizontal plane. The covariance of the Gaussian noise is Σ_r . The search domain, Θ , spans a unitless 40 by 40 square, with $\psi_a \in [0, 2\pi)$. The prior distribution is uniform over the domain. The search vehicle is modeled as a quadrotor aircraft [8]. They accelerate laterally proportional to their tilt, with constraints on speed, acceleration, and proximity.

The proposed method quickly localizes the target, as shown in Figure 2. The quadrotors move to locations that reinforce one another's measurements, in a behavior substantially more complicated than that required for range or bearing sensors. The posterior distribution, visualized by the particles, demonstrates the ability of this method to handle complicated posterior beliefs. It successfully exploits the structure of the distribution to reduce uncertainty.

V. CONCLUSION

A set of methods were developed to enable information-theoretic distributed control of mobile sensor network search,

using particle filters. Formulae were derived to compute information-theoretic quantities using particle filters, and single-node and pairwise-node approximations were derived to enable scalability in network size. The methods were applied in simulation. The results of search using bearings-only sensing demonstrate the utility of the techniques, and the results of search using the avalanche rescue beacon demonstrate the ability to handle problems that would pose significant hurdles to previous strategies. The techniques open the door to a variety of future applications.

REFERENCES

- [1] Gordon, N., Salmond, D., and Smith, A., "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *Radar and Signal Processing, IEE Proceedings F*, Vol. 140(2), Apr 1993, pp. 107–113.
- [2] Grocholsky, B., Makarenko, A., and Durrant-Whyte, H., "Information-Theoretic Coordinated Control of Multiple Sensor Platforms," *Proceedings of the IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, Sept 2003, pp. 1521–1526.
- [3] Chung, T. H., Gupta, V., Burdick, J. W., and Murray, R. M., "On a Decentralized Active Sensing Strategy using Mobile Sensor Platforms in a Network," *43rd IEEE Conference on Decision and Control*, December 2004.
- [4] Thrun, S., Burgard, W., and Fox, D., *Probabilistic Robotics*, MIT Press, Cambridge, MA, 2005.
- [5] Bourgault, F., Furukawa, T., and Durrant-Whyte, H. F., "Decentralized Bayesian Negotiation for Cooperative Search," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robotics and Systems*, Sendai, Japan, September 2004, pp. 2681–2686.
- [6] Zhao, F., Shin, J., and Reich, J., "Information-Driven Dynamic Sensor Collaboration," *IEEE Signal Processing Magazine*, March 2002, pp. 61–72.
- [7] Kreucher, C., Kastella, K., and III, A. O. H., "Information Based Sensor Management for Multitarget Tracking," *Proceedings of SPIE*, Vol. 5204, Bellingham, WA, August 2003, pp. 480–489.
- [8] Waslander, S. L., Hoffmann, G. M., Jang, J. S., and Tomlin, C. J., "Multi-Agent Quadrotor Testbed Control Design: Integral Sliding Mode vs. Reinforcement Learning," *In Proceedings of the IEEE/RSJ International Conference on Intelligent Robotics and Systems 2005*, Edmonton, Alberta, August 2005, pp. 468–473.
- [9] Gustafsson, F., Gunnarsson, F., Bergman, N., Forsell, U., Jansson, J., Karlsson, R., and Nordlund, P.-J., "Particle Filters for Positioning, Navigation and Tracking," *IEEE Transactions on Signal Processing*, Vol. 50, February 2002, pp. 425–437.
- [10] Shannon, C. E., "A Mathematical Theory of Communication," *The Bell System Technical Journal*, Vol. 27, July, October 1948, pp. 379–423, 623–656.
- [11] Cover, T. M. and Thomas, J. A., *Elements of Information Theory*, John Wiley & sons, 1991.
- [12] Doucet, A., de Freitas, N., and Gordon, N., editors, *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, Springer, New York, NY, 2001.
- [13] Bergman, N., *Recursive Bayesian Estimation: navigation and tracking applications*, PhD thesis 579, Linköpings Universitet, 1999.
- [14] Moin, P., *Fundamentals of Engineering Numerical Analysis*, Cambridge University Press, 2001.
- [15] Inalhan, G., Stipanovic, D. M., and Tomlin, C. J., "Decentralized Optimization, with Application to Multiple Aircraft Coordination," *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, December 2002.
- [16] Narayana Rao, N., *Elements of Engineering Electromagnetics*, Prentice-Hall, Upper Saddle River, New Jersey, 5th ed., 2000.

APPENDIX

To exchange conditioning variables in mutual information, use random variables, a , b , and c . Then,

$$I(a; b) - I(a; b|c) = H(a) - H(a|b) - H(a|c) + H(a|b, c)$$

$$= I(a; c) - I(a; c|b)$$

therefore, $I(a; b|c) = I(a; b) - I(a; c) + I(a; c|b)$.