Proceedings of the 2005 IEEE
International Conference on Robotics and Automation
Barcelona, Spain, April 2005

Multi-vehicle Bayesian Search for Multiple Lost Targets

El-Mane Wong†, Frédéric Bourgault††, Tomonari Furukawa†
ARC Centre of Excellence in Autonomous Systems
†School of Mechanical and Manufacturing Engineering, The University of New South Wales, Sydney 2052 Australia
††Australian Centre for Field Robotics, The University of Sydney, Sydney 2006 Australia
Email: {f.bourgault, t.furukawa}@cas.edu.au

***Abstract*— This paper presents a Bayesian approach to the problem of searching for multiple lost targets in a dynamic environment by a team of autonomous sensor platforms. The probability density function (PDF) for each individual target location is accurately maintained by an independent instance of a general Bayesian filter. The team utility for the search vehicles trajectories is given by the sum of the 'cumulative' probability of detection for each target. A dual-objective switching function is also introduced to direct the search towards the mode of the nearest target PDF when the utility becomes too low in a region to distinguish between trajectories. Simulation results for both clustered and isolated targets demonstrate the effectiveness of the proposed search strategy for multiple targets.**

I. Introduction

Search theory is the discipline that studies the problem of how best to search for an object when the search effort resources are limited and only probabilities of the possible location of the object are given [6]. This research discipline was initiated by B. O. Koopman and his colleagues in the Antisubmarine Warfare Operations Research Group (ASWORG) during World War II [13] and further generalized by Stone [11], [12].

This paper is concerned with the problem of optimally controlling a team of autonomous sensor platforms searching for a known number of mobile targets with uncertain locations. In this approach, different instances of a Bayesian filter fuse the sensor observations and track the evolution of each target probability density function (PDF) as the search progresses. The search vehicles use these PDFs to plan trajectories with high probability of finding the targets. This paper presents a multi-target generalization of the single target decentralized Bayesian search framework published in [1] which built on the single vehicle framework introduced in [2]. A new switching objective function is also introduced to prevent the search vehicles getting trapped in regions of low probability density.

In [11], search plans for static single targets with optimal effort allocation have been discussed with the restrictive assumption of exponential detection functions. In [16], Bayesian techniques are used to update discrete probability grids and quasi-optimal paths are found to search for static targets. A distributed search framework which also uses Bayesian techniques to update the target probability density can be found in [10] and [5]. In [15] a pursuit-evasion game approach is used to deal with evading targets. In [9], a sensor placement problem is solved using distributed heuristics to place the sensor in view of the targets. In this paper, unlike for the above references, the entire target PDF is accurately maintained for each target throughout the search. It is argued that no adequate search control decision may be made without it.

The paper is organized as follows. First, Sec. II describes the Bayesian filtering algorithm that accurately predicts and updates the PDF of each target. Section III presents the cooperative search control problem and the objective functions used to planned the multi-target search trajectories. Section IV demonstrates the efficacy of the proposed search strategy through numerical examples and discussed the relevant issues. Finally, conclusions and on going research are presented in the last section.

II. Multi-Target Bayesian Search

This section reviews the mathematical formulation of the Bayesian filtering algorithm that accurately maintains the essential information about the targets throughout the search. The Bayesian approach is particularly suitable for maintaining the highly non-Gaussian general PDF of the targets state by taking into account all sources of uncertainty, including non-linear process models and heterogeneous non-Gaussian sensor observations [14].

A. Bayesian Search Filter

For the multi-target search problem, with a number N_s of search sensors and N_t of targets, the random variable of interest is the combined targets state vector, denoted $\mathbf{x}_k^t = \{\mathbf{x}_k^{t_j} : j = 1, ..., N_t\}$ where $\mathbf{x}_k^{t_j} \in \Re^{n_x}$ represents the individual state vector for the j^{th} target at time step k. In general, $\mathbf{x}_k^t$ describes the targets location but could also include their attitude, velocity and other properties. In this paper, the superscripts t_j and s_i indicate a relationship to the target j and the sensor onboard the search vehicle i respectively. The subscripts are used to indicate the time index.

The purpose of the filter is to produce an estimate for the targets joint probability density, $p(\mathbf{x}_k^t|\mathbf{z}_{1:k}) = p(\mathbf{x}_k^{t_1}, ..., \mathbf{x}_k^{N_t}|\mathbf{z}_{1:k})$, given the sequence of all the observations made up to time step k by the N_s search sensors, $\mathbf{z}_{1:k} = \{\mathbf{z}_m^i : i = 1, ..., N_s, m = 1, ..., k\}$, with $\mathbf{z}_m^i$ being the observation from the i^{th} vehicle at the time step m.

When the different target states are highly correlated, there is a definite advantage of maintaining the entire joint PDF as the observation of one target may contribute information about the location of the other targets. However powerful, this sort of approach is intractable for large numbers of targets. The computational cost and memory usage rapidly become prohibitive as they increase exponentially with the number of targets and the number of states for each target [14]. In this paper, in an effort to limit the complexity of maintaining such a high dimensional distribution, the targets individual densities are assumed to be independent of each other. This implies that a different Bayesian filter may be instantiated to maintain a separate independent PDF for each target.

Fortunately, in real life scenarios it is often the case that the targets are completely unrelated, e.g. two independent hikers lost in the bush. The targets may also be loosely coupled by sharing the same process, e.g. drifting life-rafts exposed to the same wind environment. In this later case, it is still reasonable to assume independence between the individual target densities as the induced estimation error is conservative and often considered negligible.

B. Individual Bayesian Filter for Each Target

The individual filter for each target j is initialized by determining a prior PDF, denoted $p(\mathbf{x}_0^{t_j}|\mathbf{z}_0) \equiv p(\mathbf{x}_0^{t_j})$, for its state at time 0, given all available prior information including past experience and domain knowledge. If nothing other than the initial bounds on the target state vector is known, then a least informative uniform PDF is used as the prior. Once the prior distribution has been established, the PDF at time step k, $p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k})$, can be produced recursively by using the following prediction and update equations alternatively.

1) **Prediction:** Suppose the system is at time step $k-1$ and the latest update for the j^{th} target, $p(\mathbf{x}_{k-1}^{t_j}|\mathbf{z}_{1:k-1})$, is available. This prior PDF is predicted forward to time step k by using the following Chapman-Kolmogorov equation

$$p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k^{t_j}|\mathbf{x}_{k-1}^{t_j})p(\mathbf{x}_{k-1}^{t_j}|\mathbf{z}_{1:k-1})d\mathbf{x}_{k-1}^{t_j} \quad (1)$$

where $p(\mathbf{x}_k^{t_j}|\mathbf{x}_{k-1}^{t_j})$ is a probabilistic Markov motion model. Also referred to as the process model, it describes the probability of transition of the target states form a given prior state, $\mathbf{x}_{k-1}^{t_j}$, to a destination state, $\mathbf{x}_k^{t_j}$. Deriving the process model from the equations of motion of the target and the probability distribution on their inputs is out of the scope of this paper. Ref. [4], however, provides some examples of realistic process models with constraints.

2) **Update:** At time step k, a new set of observations $\mathbf{z}_k = \{\mathbf{z}_k^1, ..., \mathbf{z}_k^{N_s}\}$ becomes available. For each sensor i, the mapping of the target state observation probability, $\mathbf{z}_k^i \in \Re^{n_z}$, for each given target state, $\mathbf{x}_k^{t_j} \in \Re^{n_x}$, is denoted $p(\mathbf{z}_k^i|\mathbf{x}_k^{t_j})$ and will be referred to as observation likelihood for a fixed $\mathbf{z}_k^i$. It is reasonable to assume all observations to be conditionally independent given the knowledge of the current state. Combining this assumption with the Bayes rule leads to the "independent opinion pool" used to update the predicted target PDF:

$$p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k}) = K\, p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k-1}) \prod_{i=1}^{N_s} p(\mathbf{z}_k^i|\mathbf{x}_k^{t_j}) \quad (2)$$

where K is the normalization factor given by

$$K = 1/\int \left[p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k-1}) \prod_{i=1}^{N_s} p(\mathbf{z}_k^i|\mathbf{x}_k^{t_j})\right] d\mathbf{x}_k^{t_j} \quad (3)$$

C. Search Performance – The Cumulative Prob. of Detection

This section describes how, using the output from the above filter equations, the performance of a multi-vehicle search plan may be assessed by determining the 'cumulative' probability of detection for each target.

Let the detection likelihood for target j by sensor i at time step k be given by $p(\mathbf{z}_k^i = D_k^i|\mathbf{x}_k^{t_j})$ where D_k^i represents a 'detection' event by the sensor on vehicle i at time k. The likelihood of 'miss' by the same sensor is given by its complement $p(\mathbf{z}_k^i = \overline{D}_k^i|\mathbf{x}_k^{t_j}) = 1 - p(D_k^i|\mathbf{x}_k^{t_j})$. The combined 'miss' likelihood for all the vehicles at time step k is simply the multiplication of the individual 'miss' likelihoods for that target

$$p(\overline{D}_k|\mathbf{x}_k^{t_j}) = \prod_{i=1}^{N_s} p(\overline{D}_k^i|\mathbf{x}_k^{t_j}) \quad (4)$$

where $\overline{D}_k = \overline{D}_k^1 \cap ... \cap \overline{D}_k^{N_s}$ represents the event of a 'miss' observation by every sensor at time step k.

If the normalization factor K is neglected, the update equation (2) can be rewritten as the following pseudo-update equation.

$$p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k})' = p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k-1})' \prod_{i=1}^{N_s} p(\mathbf{z}_k^i|\mathbf{x}_k^{t_j}) \quad (5)$$

The advantage of not normalizing the target PDF is that the joint probability of failing to detect (miss) the target in all of the steps from 1 to k, denoted $Q_k^{t_j} = p(\overline{D}_{1:k})$ with $\overline{D}_{1:k}$ representing the series of 'miss' observations $\mathbf{z}_k = \overline{D}_k, \forall k$, can be obtained directly from its integration as in

$$Q_k^{t_j} = \int p(\mathbf{x}_k^{t_j}|\overline{D}_{1:k})' d\mathbf{x}_k^{t_j} = \int p(\mathbf{x}_k^{t_j}|\overline{D}_{1:k-1})' p(\overline{D}_k|\mathbf{x}_k^{t_j})\, d\mathbf{x}_k^{t_j} \quad (6)$$

It corresponds to the volume left under the surface of the pseudo-density function after update. It represents the residual probability that the target is still present at time step k despite the search effort expended up to that time.

As shown in [1], the probability that a target gets detected for the first time at time step k, namely the probability of detection at time step k, is denoted $p_k^{t_j} = p(\overline{D}_{1:k-1}, D_k)$. It corresponds to the reduction in volume under the pseudo-density function $(-\Delta Q_k^{t_j})$ when it is updated with the combined 'detection' likelihood, denoted $p(D_k|\mathbf{x}_k^{t_j}) = \left[1 - p(\overline{D}_k|\mathbf{x}_k^{t_j})\right]$, with $p(\overline{D}_k|\mathbf{x}_k^{t_j})$ given in (4), and is obtained as follows

$$p_k^{t_j} = \int p(\mathbf{x}_k^{t_j}|\overline{D}_{1:k-1})' \left[1 - p(\overline{D}_k|\mathbf{x}_k^{t_j})\right] d\mathbf{x}_k^{t_j} = Q_{k-1}^{t_j} - Q_k^{t_j} \quad (7)$$

Assuming no false detection from the sensors, the probability that the target j *has* been detected in k steps, denoted $P_k^{t_j}$, is obtained from the cumulative sum of the $p_k^{t_j}$'s as in

$$P_k^{t_j} = \sum_{m=1}^{k} p_m^{t_j} = P_{k-1}^{t_j} + p_k^{t_j} \quad (8)$$

For this reason $P_k^{t_j}$ will be referred to as the 'cumulative' probability of detection to distinguish it from the payoff probability of detection function $p_k^{t_j}$.

III. Search Control

This section describes the cooperative search control problem and the necessary utility functions needed to compute efficient cooperative search trajectories.

A. Cooperative Search Control Problem

The cooperative search control problem consists in determining the control sequence that will maximize the team search utility over the duration of the search mission. At time step k, the team utility for a given planning horizon depth of N_k steps is denoted $J_k(\mathbf{u}, N_k)$ with $\mathbf{u} = \mathbf{u}_{k:k+N_k-1} = \{\mathbf{u}_{k:k+N_k-1}^{s_i} : i = 1, ..., N_s\}$ being the control action sequence of all the vehicles starting at step k. Given the search vehicles dynamic models and sensors observation models, the optimal control trajectory $\mathbf{u}^*$ is the sequence that maximizes that utility subject to the control bounds $\mathbf{u}_{LB} \leq \mathbf{u} \leq \mathbf{u}_{UB}$ and the kinematic and dynamic constraints $\mathbf{g}(\mathbf{u}, \mathbf{x}, N_k) \leq \mathbf{0}$, where $\mathbf{x} = \mathbf{x}^{s_{1:N_s}}_{k:k+N_k-1}$.

$$\mathbf{u}^* = \{\mathbf{u}^{s_1*}_{k:k+N_k-1}, ..., \mathbf{u}^{s_{N_s}*}_{k:k+N_k-1}\} = \arg\max_{\mathbf{u}} J_k(\mathbf{u}, N_k) \quad (9)$$

B. Multi-Target Team Utility

The probability of detecting target j, given a series of observations generated by the control sequence over the planning horizon starting at time step k, is given by the engendered net variation in cumulative probability of detection

$$\Delta P^{t_j}_k = \sum_{l=k+1}^{k+N_k} p^{t_j}_l = P^{t_j}_{k+N_k} - P^{t_j}_k \quad (10)$$

with $p^{t_j}_l$ obtained from (7). Notice that $\Delta P^{t_j}_k$ can also be directly obtained from the corresponding reduction in volume, $Q^{t_j}_k - Q^{t_j}_{k+N_k}$, under the pseudo-density function caused by the observations. This measure was used in [1] as the team utility function for the single target search problem. For the multi-target search problem, this paper proposes the following weighted sum of the above measure

$$J_k(\mathbf{u}, N_k) = \sum_{j=1}^{N_t} w^{s_i}_k \Delta P^{t_j}_k, \quad \text{with} \quad \sum_{j=1}^{N_t} w^{s_i}_k = 1 \quad (11)$$

where the weights $w^{s_i}_k$'s correspond to the relative priority for each target at time step k.

C. Solving the Search Control Problem

To obtain truly optimal search trajectories, the cooperative search control problem in (9) should be solved for the entire duration of the search mission. In that case, the team utility function from (11) reduces to the weighted sum of the cumulative probability of detection for each target. However, with increasing lookahead depth and number of agents, the solution becomes intractable due to the "curse of dimensionality". This is why, in this paper, a *rolling time horizon* solution approach is adopted where the planned trajectory is recomputed at short intervals to keep the lookahead constant as the agents progress forward.

To further reduce the computational complexity of the problem, a decentralized "coordinated" control solution was proposed in [1]. In that approach each decision maker builds an equivalent estimate of the target PDF by communicating and fusing the information from the other sensors and replans its trajectory at short interval based on its local greedy utility function, denoted $J^{s_i}_k(\mathbf{u}^{s_i}, N_k)$, without considering the effects on the PDF of the other vehicle predicted observations. The same approach is followed in this paper with the individual greedy version of the team utility in (11) given by

$$J^{s_i}_k(\mathbf{u}^{s_i}, N_k) = \sum_{j=1}^{N_t} w^{s_i}_k \sum_{l=k+1}^{k+N_k} \int p(\mathbf{x}^{t_j}_l | \overline{D}^i_{1:l-1})' [1 - p(\overline{D}^i_l | \mathbf{x}^{t_j}_l)] \, d\mathbf{x}^{t_j}_l \quad (12)$$

which for a time-horizon of $N_k = 1$, reduces to the weighted sum of each target probability of detection, as in $J^{s_i}_k(\mathbf{u}^{s_i}, 1) = \sum_{j=1}^{N_t} w^{s_i}_k p^{t_j}_k$, with $p^{t_j}_k$ given in (7). For longer time-horizons, piecewise constant control solutions [7] are obtained for the individual optimization problems by using a constrained nonlinear programming technique called Sequential Quadratic Programming (SQP) [8].

For a cooperative control solution that is jointly optimal for the group, the reader is referred to [3]. A decentralized negotiation algorithm is presented that enables each vehicle to iteratively converge to the control solution that is good for them, and the rest of the team, i.e. Nash solution.

D. Distance to PDF Mode as Another Objective Function

One practical issue related to the control optimization problem comes from the optimization routine sometimes terminating before reaching a proper solution. This often occurs due to the difficulty for the routine to distinguish between the utility of different trajectories when the search vehicle is in a region of the space where the target probability density is very low. This effect is also made worse when the lookahead depth is very short.

A quick fix to this problem consists in increasing the optimizer sensitivity by reducing the termination tolerance. However, this has the undesirable secondary effect of significantly increasing the computation time and the optimizer might still fail to converge when the target density is very low.

To alleviate this issue when the utility in (12) it is too low, i.e. below a certain threshold, it is proposed to direct the vehicle towards the mode of the nearest predicted target PDF. The distance $d^{s_i t_j}_k$ between the position estimate for vehicle i, denoted $\hat{\mathbf{x}}^{s_i}_k = [\hat{x}^{s_i}_k, \hat{y}^{s_i}_k, \hat{\theta}^{s_i}_k]^T$, and the mode for the j^{th} target predicted PDF, as defined by

$$\hat{\mathbf{x}}^{t_j}_k = [\hat{x}^{t_j}_k, \hat{y}^{t_j}_k]^T = \arg\max_{\mathbf{x}^{t_j}_k} p(\mathbf{x}^{t_j}_k | \mathbf{z}_{1:k-1}) \quad (13)$$

is obtained as follows

$$d^{s_i t_j}_k = \sqrt{(\hat{x}^{s_i}_k - \hat{x}^{t_j}_k)^2 + (\hat{y}^{s_i}_k - \hat{y}^{t_j}_k)^2} \quad (14)$$

Hence, the distance from vehicle i to the nearest target mode at time step k is given by

$$d^{s_i}_k = \max_j d^{s_i t_j}_k \quad (15)$$

The proposed alternative objective function is defined as

$$G^{s_i}_k(\mathbf{u}^{s_i}, N_k) = -d^{s_i}_{k+N_k} \cdot (1 - J^{s_i}_k) \quad (16)$$

where $J^{s_i}_k$, determined in (12), is the 'cumulative' probability of detection of any target by sensor i over the time horizon. This objective function corresponds to a tradeoff between the distance to the nearest target and the detection reward along the path. It will take the vehicle on the shortest path towards the nearest mode while trying to maximize the latter.

E. Switching Objective Function

The two objective functions $J^{s_i}_k$ and $G^{s_i}_k$, defined in (12) and (16) respectively, can be merged into a dual-objective switching function as follows

$$U^{s_i}_k(\mathbf{u}^{s_i}, N_k) = \gamma^{s_i}_k J^{s_i}_k(\mathbf{u}, N_k) + (1 - \gamma^{s_i}_k) G^{s_i}_k(\mathbf{u}, N_k) \quad (17)$$

where $\gamma^{s_i}_k \in [0, 1]$ is the adaptive switching parameter for vehicle i at time step k. By default $\gamma^{s_i}_k$ is set to one. When $J^{s_i}_k$ falls under a threshold that is proportional to the planning horizon length, the number of targets and the sum of $P^{s_i}_k$'s, $\gamma^{s_i}_k$ switches to zero so the distance utility $G^{s_i}_k$ is optimized.

IV. Simulation Results

This section presents simulation results for the multi-target multi-vehicle Bayesian search approach described in this paper and discusses the relevant issues. The chosen scenarios involve a team of unmanned air vehicles (UAVs), each equipped with a downward looking millimeter wave radar and searching for multiple lost targets, i.e. life-rafts, drifting with the wind at sea over an area of about 8x8km. The prior target PDFs are Gaussian densities with varying standard deviations in each direction. The wind process is modelled as a beta-Gaussian function as described in [2]. The UAVs are flying at the constant speed of 50 m/s, which is considerably faster than the targets, and communicate via wireless broadcasting. All vehicles carry their own filters that maintain equivalent estimates of the targets PDFs and plan for their own greedy actions. The planning horizon for all the examples is only one step ahead. More about the implementation details of the framework including the sensor observation model, vehicle model, wind process model and the decentralized coordinated control framework implemented can be found in [2] and [1].

A. Single Objective Multi-Target Bayesian Search

Figure 1 results demonstrate the multi-vehicle multi-target search algorithm for a team of five homogeneous search vehicles performing a decentralized coordinated search for three drifting targets. The planning is done for the single objective function, $J_k^{s_i}$, with a lookahead of one step, which corresponds to the sum of the probability of detection $p_k^{t_j}$'s as mentioned in Sec. III-C.

Figure 2 shows the corresponding search performance measures for the trajectories shown in Fig. 1. As the search progress and k increases towards infinity, $P_k^{t_j}$ levels off towards one as it becomes harder to generate additional observation payoff, $p_k^{t_j}$, from hardly any probability mass left in the PDF. As seen on Fig. 2a, the final 'cumulative' probability of detection reached for each target after searching for eighty seconds, are $P_{80}^{t_1} = .41$, $P_{80}^{t_2} = .83$, and $P_{80}^{t_3} = .69$, which, as seen on Fig. 2b, sum up to $\Sigma_j P_k^{t_j} = 1.93$, over a possibility of 3.0, corresponding to 64.3%. These values confirm that target #1 has been neglected and that it would be rewarding for the vehicles to allocate resources to that target. As confirmed by Fig. 1h, there is still much probability mass left in that PDF as only one vehicle has searched for it.

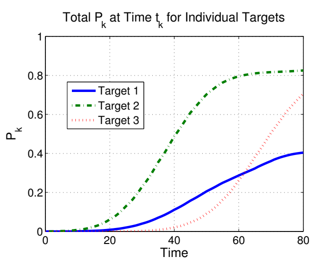

(a) $P_k^{t_j}$'s vs. k

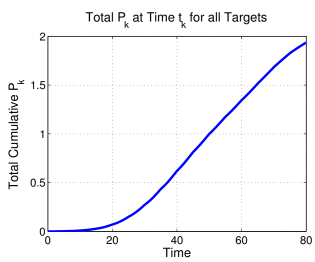

(b) $\Sigma_j P_k^{t_j}$ vs. k

Fig. 2. Search performance for the single objective control: (a) Evolution of the 'cumulative' probability of detection, $P_k^{t_j}$, for each target; (b) The sum of all targets 'cumulative' probability of detection.

B. Search with Switching Objective

Figure 3 illustrates the advantage of introducing the dynamically switching objective function. In Fig. 3b, it can be seen that the switching function steers the search vehicles back towards the PDF peaks, even at the cost of crossing over low probability density regions, instead of drifting away like it occurs with the single objective function in Fig. 3c. It can be seen that at the same when the switching parameters in Fig. 3e start being active, the cumulative probability sum start diverging in Fig. 3f. The final values reached are 2.90, 2.71, and 1.93 for the switching objective, single objective and parallel search respectively. This represents a 7% increase in utility for the switching, over the single objective search, and more than 50% improvement over the area coverage search.

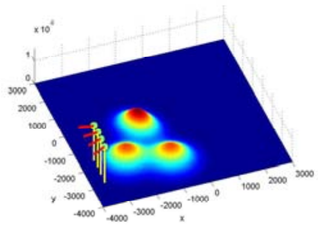

(a) Switch. obj., $k = 10$s

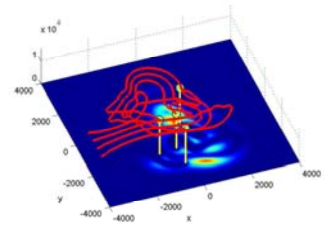

(b) Switch. obj., $k = 250$s

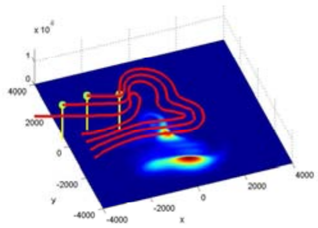

(c) Single obj., $k = 190$s

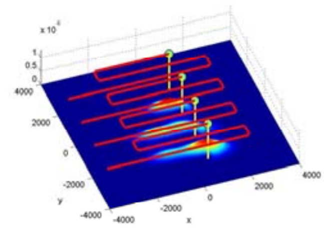

(d) Area Cov., $k = 250$s

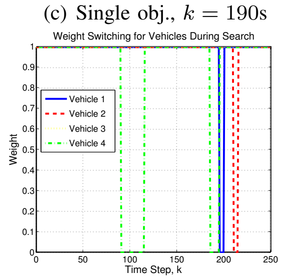

(e) $\gamma_k^{s_i}$'s vs. k

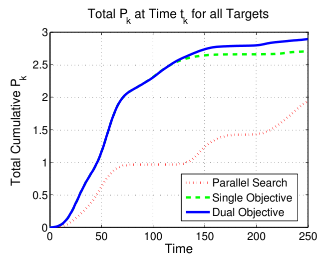

(f) $\Sigma_j P_k^{t_j}$ vs. k

Fig. 3. Coordinated search for 3 drifting targets by 4 vehicles using the switching objective function, $U_k^{s_i}$: (a)-(b) 3D views of the pseudo-PDFs sum, $\Sigma_j p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k})$, at time step $k = 10$ and 250s respectively and the corresponding search trajectories obtained using the switching objective function, $U_k^{s_i}$; (c) Incomplete search result, at $k = 190s$, obtained using the single objective function $J_k^{s_i}$; (d) Area coverage search with flight with parallel flight patterns formation search result at $k = 250s$; (e) Evolution of the vehicles switching function parameters, $\gamma_k^{s_i}$'s; (f) Comparison of the combined 'cumulative' probability of detection, $\Sigma_j P_k^{t_j}$, for the switching objective search (blue solid), single objective (green dashed), and parallel patterns area coverage search (red dotted).

C. Simulating a 'Detection' Event

To simulate the overall search problem, actual targets are introduced randomly according to their prior distribution. In Fig. 4, once a target is detected, its PDF is removed and replaced by an arrow sign. It is no more considered in the search planning. For the simulation purposes, a target is

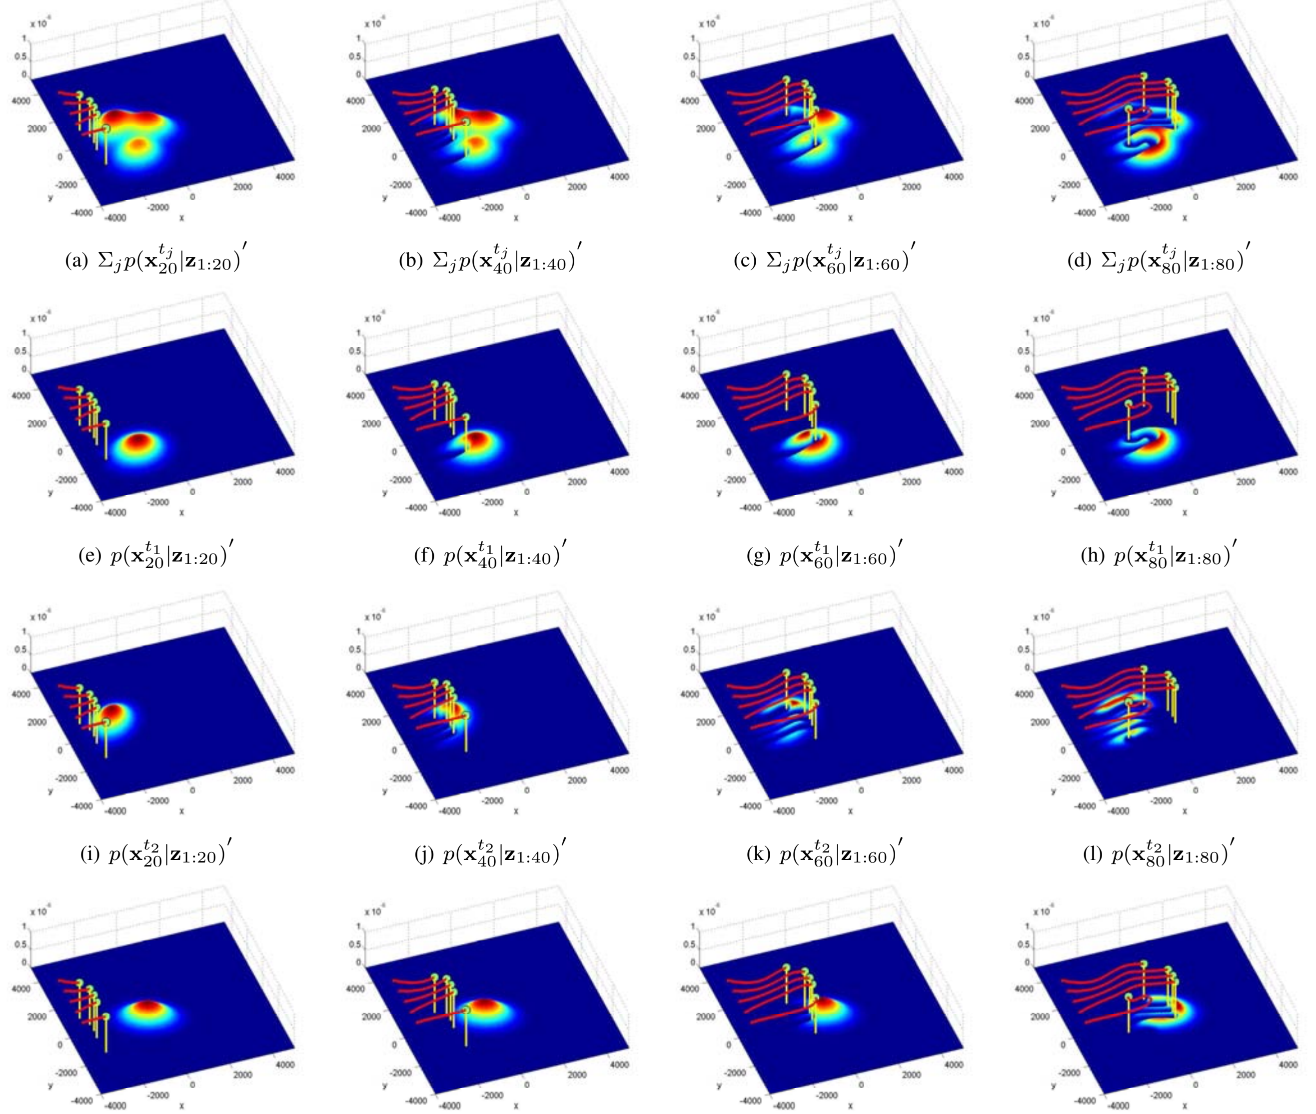

(a) $\Sigma_j p(\mathbf{x}_{20}^{t_j}|\mathbf{z}_{1:20})'$ (b) $\Sigma_j p(\mathbf{x}_{40}^{t_j}|\mathbf{z}_{1:40})'$ (c) $\Sigma_j p(\mathbf{x}_{60}^{t_j}|\mathbf{z}_{1:60})'$ (d) $\Sigma_j p(\mathbf{x}_{80}^{t_j}|\mathbf{z}_{1:80})'$

(e) $p(\mathbf{x}_{20}^{t_1}|\mathbf{z}_{1:20})'$ (f) $p(\mathbf{x}_{40}^{t_1}|\mathbf{z}_{1:40})'$ (g) $p(\mathbf{x}_{60}^{t_1}|\mathbf{z}_{1:60})'$ (h) $p(\mathbf{x}_{80}^{t_1}|\mathbf{z}_{1:80})'$

(i) $p(\mathbf{x}_{20}^{t_2}|\mathbf{z}_{1:20})'$ (j) $p(\mathbf{x}_{40}^{t_2}|\mathbf{z}_{1:40})'$ (k) $p(\mathbf{x}_{60}^{t_2}|\mathbf{z}_{1:60})'$ (l) $p(\mathbf{x}_{80}^{t_2}|\mathbf{z}_{1:80})'$

(m) $p(\mathbf{x}_{20}^{t_3}|\mathbf{z}_{1:20})'$ (n) $p(\mathbf{x}_{40}^{t_3}|\mathbf{z}_{1:40})'$ (o) $p(\mathbf{x}_{60}^{t_3}|\mathbf{z}_{1:60})'$ (p) $p(\mathbf{x}_{80}^{t_3}|\mathbf{z}_{1:80})'$

Fig. 1. Coordinated 1-step lookahead search for 3 drifting targets by 5 vehicles using the single objective function, $J_k^{s_i}$: The first row of figures shows snapshots of the vehicles trajectories and the corresponding 3D views for the targets pseudo-PDFs sum, $\Sigma_j p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k})$, at time step $k = 20, 40, 60$ and $80s$ respectively. Rows 2, 3 and 4 show the snapshots for the individual target pseudo-PDFs, $p(\mathbf{x}_k^{t_j}|\mathbf{z}_{1:k})$'s, for target $j = 1, 2$ and 3 respectively.

automatically detected when the detection likelihood at its location gets above 5%. The data association problem is out of the scope in this paper. The search continues until all the targets are detected.

As a result of removing the target PDF upon detection, the probability density in the surrounding area may drop significantly. When this occurs, some of the vehicles accordingly switch objective and head towards the nearest target peak, as shown in Figs. 4a to 4d. When using the single objective function however the detection utility may fall below the optimization tolerance for some of the vehicles causing them to keep a steady heading that takes them out of the search area, as depicted in Fig. 4e. In Fig. 4f, this problem is resolved by considerably lowering the termination tolerance of the optimization routine. This, however, costs additional computation time and does not always guarantee convergence to a solution. For example, vehicle #1 incurs about 25% more computational time when using the single objective function with low termination tolerance than with the switching function. Despite this additional cost, it still becomes trapped in a low density region as illustrated by the lowest of the trajectories in Fig. 4f. Figure 4h shows when the utility function switches to direct the vehicle towards the target #4 after the detection of target #3.

V. Conclusion and Ongoing Work

An multi-target Bayesian search approach to coordinating multiple mobile sensor platforms was presented. The targets were assumed independent, thus allowing the use of different instances of a Bayesian filter to accurately maintain their PDFs. For the one-step lookahead coordinated control solutions presented the individual search utility was reduced to the sum of the probability of detecting each target. In

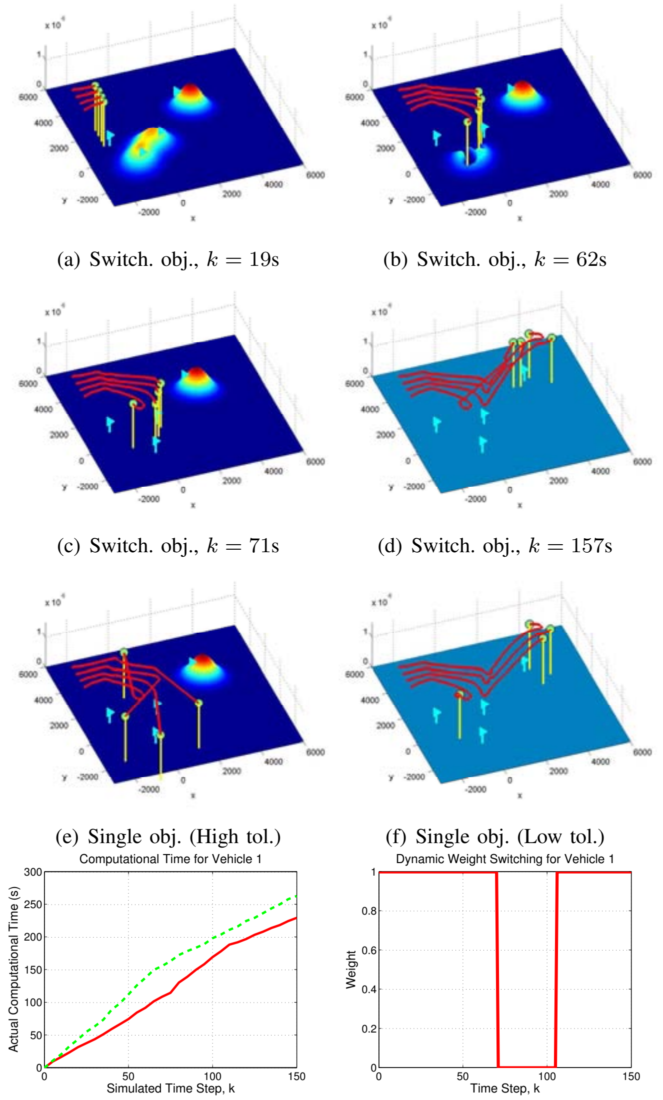

(a) Switch. obj., $k = 19$s (b) Switch. obj., $k = 62$s

(c) Switch. obj., $k = 71$s (d) Switch. obj., $k = 157$s

(e) Single obj. (High tol.) (f) Single obj. (Low tol.)

(g) Computation time (h) γ_k^{s1} vs. k

Fig. 4. Coordinated search results for 4 targets with 4 vehicles: (a)-(d) Snapshots of the search trajectories and 3D views of the pseudo-PDFs sum obtained using the switching objective function, at time step $k = 19, 62, 71$ and 157s, corresponding to the respective moment when each target gets detected; (e)Results obtained using the single objective function with high terminal tolerance at $k = 122$s; (f) Improved search trajectories using the single objective function with a low termination tolerance; (g) Computational time required with the switching objective (bottom red line) versus the single objective with low termination tolerance (top green line); (h) Evolution of the switching parameter value for vehicle 1.

multi-target search problems, the targets may often be far apart, or their densities may be well dispersed. To improve the search performance for this sort of conditions where the control optimization routine may fail to provide a solution, an additional objective function was introduced that directs the vehicles towards regions of high probability density. The approach was proven to find effective search trajectories even with a very short planning horizon.

Some of the current work emphasis is now on the coordination problem at the mission level. The goal is to solve the dual problem of searching and, at the same time, allocating some of the resources to rescuing the targets that have already been detected. This problem is very important in real-life search and rescue scenarios and can become quite complex when resources are limited and some of the vehicles can only conduct search, some only rescue, and some others can perform both parts of the mission with various degrees of efficacy.

Another aspect of the ongoing research effort relates to the problem of human-robot interactions. Humans may be operators giving orders or contributing their own observations to the network thus affecting the team control decisions. They may also be active team members cooperating with the network of autonomous robots. These are crucial aspects if such robotic systems are going to be deployed and contribute to time critical missions where human lives may be at stake.

Acknowledgment

This work is partly supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

References

[1] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte. Coordinated decentralized search for a lost target in a Bayesian world. In *Proceedings of the 2003 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'03)*, pages 48–53, Las Vegas, Nevada, October 2003.

[2] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte. Optimal search for a lost target in a Bayesian world. In *The 4th Int. Conf. on Field and Service Robotics (FSR'03)*, Lake Yamanaka, Japan, July 2003.

[3] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte. Decentralized Bayesian negotiation for cooperative search. In *Proceedings of the 2004 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'04)*, pages 2681–2686, Sendai, Japan, September 2004.

[4] F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte. Process model, constraints, and the coordinated search strategy. In *Proceedings of the 2004 IEEE Int. Conf. on Robotics and Automation (ICRA'04)*, pages 5256–5261, New Orleans, LA, April 2004.

[5] M. Flint, M. Polycarpou, and E. Fernandez-Gaucherand. Cooperative control for multiple autonomous uav's searching for targets. In *Proc. 41st IEEE Conference on Decision and Control*, pages 2823–2828, Las Vegas, USA, December 2002.

[6] J.-P. Le Cadre. Searching tracks. *IEEE Trans. Aero. Electr. Sys.*, 36(4):1149–1166, 2000.

[7] H.J.W. Lee, K.L. Teo, V. Rehbock, and L.S. Jennings. Control parametrization enhancing technique for time optimal control problems. *Dyn. Sys. and Appl.*, 6(2):243–262, April 1997.

[8] G.L. Nemhauser, A.H.G. Rinnooy Kan, and M.J. Todd. *Vol. 1: Optimization*. Handbooks in Operations Management Science. Elsevier Science Publishers B.V., Amsterdam, 1989.

[9] L.E. Parker. Distributed algorithms for multi-robot observation of multiple moving targets. *Autonomous Robots*, 12(3):231–255, 2002.

[10] M.M. Polycarpou, Y. Yang, and K.M. Passino. A cooperative search framework for distributed agents. In *Proc. IEEE Int. Conference on Intelligent Control*, pages 1–6, Mexico City, Mexico, September 2001.

[11] L.D. Stone. *Theory of Optimal Search*, volume 118 of *Mathematics in Science and Engineering*. Academic Press, New York, 1975.

[12] L.D. Stone. What's happened in search theory since the 1975 Lanchester prize? *Operations Research*, 37(3):501–506, 1989.

[13] L.D. Stone. *Search and Screening: General Principle with Historical Applications (B.O. Koopman)*. MORS Heritage Series, Alexandria, VA, 1999.

[14] L.D. Stone, C.A. Barlow, and T.L Corwin. *Bayesian Multiple Target Tracking*. Mathematics in Science and Engineering. Artech House, Boston, 1999.

[15] R. Vidal, S. Rashid, C. Sharp, O. Shakernia, J. Kim, and S.S. Sastry. Pursuit-evasion games with unmanned ground and aerial vehicles. In *Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA'01)*, pages 2948–2955, Seoul, Korea, May 2001.

[16] I. Yan and G.L. Blankenship. Numerical methods in search path planning. In *Proceedings of the 27th IEEE Conference on Decision and Control*, pages 1563–1569 vol.2, Astin, Texas, December 1988.