Proceedings of the 2003 IEEE/RSJ
Intl. Conference on Intelligent Robots and Systems
Las Vegas, Nevada · October 2003

Coordinated Decentralized Search for a Lost Target in a Bayesian World

Frédéric Bourgault
Australian Centre for Field Robotics
The University of Sydney
Sydney, NSW 2006, Australia
f.bourgault@acfr.usyd.edu.au

Tomonari Furukawa
School of Mech. and Manuf. Eng.
The University of New South Wales
Sydney, NSW 2052, Australia
t.furukawa@unsw.edu.au

Hugh F. Durrant-Whyte
Australian Centre for Field Robotics
The University of Sydney
Sydney, NSW 2006, Australia
hugh@acfr.usyd.edu.au

***Abstract*— This paper describes a decentralized Bayesian approach to coordinating multiple autonomous sensor platforms searching for a single non-evading target. In this architecture, each decision maker builds an equivalent representation of the target state PDF through a Bayesian DDF network enabling them to coordinate their actions without exchanging any information about their plans. The advantage of the approach is that a high degree of scalability and real time adaptability can be achieved. The effectiveness of the approach is demonstrated in different scenarios by implementing the framework for a team of airborne search vehicles looking for a stationary, and a drifting target lost at sea.**

I. Introduction

"Yacht *Grimalkin* capsized in position thirty miles north-west of Land's End..."[1]

When rescue authorities receive a distress signal time becomes critical. Survival expectancy decreases rapidly when lost at sea and a rescue mission's primary goal is to search for and find the castaways as diligently and efficiently as possible. The search, based on some coarse estimate of the target location, must often be performed in low visibility conditions and despite strong winds and high seas causing the location estimate to grow even more uncertain as time goes by. Keeping these time and physical constraints in mind, and given a large team of heterogeneous platforms such as high flying long range aircrafts, helicopters and ships equipped with different sensors, what should be the optimal search strategy?

This paper presents a decentralized Bayesian approach to the target detection problem as described in [8] (Chapter 9). It expands the single vehicle framework proposed in [2] to an arbitrary number of sensing platforms by integrating a fully decentralized Bayesian data fusion (DDF) technique with a decentralized coordinated control scheme that was first proposed by Grocholsky [6]. Scalability, modularity and real-time adaptability are the advantages of the decentralized approach. At any time, new rescue vehicles can join, or momentarily quit for refuelling, the search effort and the system should seemly and robustly adapt to the change.

The breakdown of the paper is as follows. Firstly, the decentralized Bayesian filtering algorithm that accurately maintains and updates the target state probability distribution is described in the next section. Then section III describes the searching problem, and section IV describes the utility function selected and formulates the decentralized control optimization problem. Then, in section V the effectiveness of the approach is demonstrated for multiple airborne search vehicles in three different scenarios for stationary, and drifting targets, and in one instance, the optimal cooperative solution is compared with the coordinated one. Finally, conclusions and ongoing research directions are highlighted in the last section.

[1]Coastguard broadcast during the desastrous 1979 *Fastnet* yacht race, August 14, 1979 [9]

II. Bayesian Filtering

This section presents the mathematical foundations of the Bayesian decentralized data fusion algorithm that keeps track of the target state PDF. The Bayesian approach is particularily suitable for combining in a rational manner heterogeneous non-gaussian sensor observations with other sources of quantitative and qualitative information [11][1].

In Bayesian analysis any quantity that is not known is considered a random variable. The state of knowledge about such a random variable is expressed in the form of a probability density function (PDF). Any new information in the form of a probabilistic observation is combined with the previous PDF using the Baye's theorem in order to update the state of knowledge and form the new *a posteriori* PDF. That PDF forms the quantitative basis on which all inferences, or control decisions (Sec.IV) are made.

In the searching problem, the unknown variable is the target state vector $\mathbf{x}^t \in \mathscr{X}^t$ which in general describes the target location but could also include its attitude, velocity, etc. The analysis starts by determining the *a priori* PDF of $\mathbf{x}^t$, $p(\mathbf{x}_0^t|\mathbf{z}_0) \equiv p(\mathbf{x}_0^t)$, which combines all available information including past experience. For example, this *a priori* PDF could be in the form a gaussian distribution representing the prior coarse estimate of the parameter of interest. If noting but the bounds is known about the parameter, the least biased approach is to represent this knowledge by a uniform PDF over the bounded region of the space. Then, once the prior distribution has been established, the PDF of the target state at time step k, $p(\mathbf{x}_k^t|\mathbf{z}_{1:k})$, can be constructed recursively, provided the sequence $\mathbf{z}_{1:k} = \{\mathbf{z}_j^i : i = 1,...,N_s, j = 1,...,k\}$ of all the observations made from the N_s sensors on board the search vehicles, $\mathbf{z}_j^i$ being the observation from the i^{th} sensor at time step j. This recursive estimation is performed in two stages: prediction and update.

A. Prediction

A prediction stage is necessary in Bayesian analysis when the PDF of the state to be evaluated is evolving with time i.e. the target is in motion or the uncertainty about its location is increasing. Suppose we are at time step $k-1$ and the latest PDF update, $p(\mathbf{x}_{k-1}^t|\mathbf{z}_{1:k-1})$ is available. Then the predicted target state PDF at time step k is obtained by the following Chapman-Kolmogorov equation

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

where $p(\mathbf{x}_k^t|\mathbf{x}_{k-1}^t)$ is a probabilistic Markov motion model. If the motion model is invariant over the target states, then the above integral is simply a convolution. Practically, this convolution is performed numerically by a discretization of the two PDF's on a grid, followed by the multiplication of their Fast Fourier Transforms (FFT)'s, and an inverse FFT of the produce to retrieve the result.

B. Update

At time step k a new set of observations $\mathbf{z}_k = \{\mathbf{z}_k^1, \ldots, \mathbf{z}_k^{N_s}\}$ becomes available and the update is performed using the Bayes rule where all the observations are assumed to be independent. In other words, the update is performed by multiplying the prior PDF (posterior from the prediction stage) by all the individual conditional observation likelihoods $p(\mathbf{z}_k^i|\mathbf{x}_k^t)$ as in the following

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

where the normalization coefficient K is given by

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

Practically, the multiplication of equation 2 is performed numerically by multiplying together the corresponding elements of a grid.

C. Bayesian DDF

In an information gathering task such as searching, if each sensor is connected to a processing unit called a node, then it is possible through communication and fusion of the information to reconstruct at each node the global information state of the world, e.g. the target state PDF. Figure 1 depicts how the update and prediction equations are integrated in the Bayesian DDF node of fully connected network.

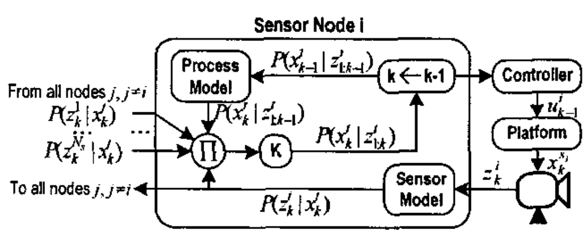

Fig. 1. Bayesian DDF node of a fully connected network.

III. The Searching Problem

This section describes the equations for computing the probability of detection of a lost object referred to as the target. For further details on the searching problem see [10] and [8].

If the target detection likelihood (observation model) from the i^{th} sensor at time step k is given by $p(\mathbf{z}_k^i|\mathbf{x}_k^t)$ where $\mathbf{z}_k^i = D_k^i$, and D_k^i represents a "detection" event by sensor i at t_k, then the likelihood of "no detection" by the same sensor is given by its complement $p(\overline{D}_k^i|\mathbf{x}_k^t) = 1 - p(D_k^i|\mathbf{x}_k^t)$.

The combined 'no detection' likelihood for all the sensors at time step k is simply a multiplication of the individual no detection likelihoods

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

where $\overline{D}_k = \overline{D}_k^1 \cap \ldots \cap \overline{D}_k^{N_s}$ represents the event of a 'no detection' observation by every sensor at time step k.

At time step k, given all the previous observations $\mathbf{z}_{1:k-1}$, the conditional probability of a combined 'no detection' event ($\mathbf{z}_k = \overline{D}_k$) to occur, noted $p(\overline{D}_k|\mathbf{z}_{1:k-1}) = q_k$, depends on how the combined 'no detection' likelihood (Eq. 4), and the latest target PDF (from the prediction stage 1) overlap. In fact, q_k is given by the reduced volume (i.e. < 1) of the target PDF after having been carved out (multiplied) by the 'no detection' likelihood in the update stage equation (Eq. 2) and before applying the normalization coefficient K to it.

$$p(\overline{D}_k|\mathbf{z}_{1:k-1}) = \int p(\overline{D}_k|\mathbf{x}_k^t)p(\mathbf{x}_k^t|\mathbf{z}_{1:k-1})d\mathbf{x}_k^t = q_k \tag{5}$$

Notice that q_k is exactly the inverse of the normalization factor K for a 'no detection' event (i.e. $q_k = 1/K$ for $\mathbf{z}_k = \overline{D}_k$ in equation 3), and is always smaller than 1.

Hence, if q_k represents the conditional probability of failing to detect the target for a specific observation step, then the joint probability of failing to detect the target in all of the steps from 1 to k, noted $Q_k = p(\overline{D}_{1:k})$, is obtained by the product of all the q_k's as follows

$$Q_k = \prod_{i=1}^{k} p(\overline{D}_i|\overline{D}_{1:i-1}) = \prod_{i=1}^{k} q_i = Q_{k-1}q_k \tag{6}$$

where $\overline{D}_{1:i-1}$ is the set $\mathbf{z}_{1:i-1}$ of observations where every observation is 'no detection' ($\overline{D}_i, \forall i$). From the above it can be deduced that the probability that the target *has* gotten detected in k steps, noted P_k, is given by $P_k = 1 - Q_k$.

Another way of obtaining P_k is to first compute p_k, the probability that the target gets detected for the first time on time step k as follows

$$\begin{aligned} p_k &= \prod_{i=1}^{k-1} p(\overline{D}_i|\overline{D}_{1:i-1})\left[1 - p(\overline{D}_k|\overline{D}_{1:k-1})\right] \\ &= \prod_{i=1}^{k-1} q_i\left[1 - q_k\right] = Q_{k-1}\left[1 - q_k\right] \end{aligned} \tag{7}$$

Assuming no false detection from the sensors, the probability of detection P_k is given by the cumulative sum of the p_k's

$$P_k = \sum_{i=1}^{k} p_i = P_{k-1} + p_k \tag{8}$$

For this reason we will refer to P_k as the 'cumulative' probability of detection to distinguish it from the conditional probability of detection at time k which is equal to $1-q_k$. Notice that as k goes to infinity, P_k increases towards one. With k increasing, the added probability of detection p_k gets smaller and smaller as the conditional probability of detection $(1-q_k)$ in Eq. 7 gets discounted by a continuously decreasing Q_{k-1}.

The mean time to detection (MTTD) is the expectation of the number of steps required to detect the target

$$E[k] = \sum_{k=1}^{\infty} k p_k = \text{MTTD} \tag{9}$$

To summarize, the goal of a searching strategy could either be to maximize the chances of finding the target given a restricted amount of time by maximizing P_k over the time horizon, or to minimize the expected time to find the target by minimizing the MTTD. The difficulty though in evaluating the MTTD lies in the fact that one must in theory evaluate p_k for all k's up to infinity. Although in practice MTTD could be evaluated approximatively over a sufficiently long interval (i.e. P_k must be close to 1).

IV. Planning

A. Optimal Trajectory

Optimality is always defined in relation to an objective, or utility function [12]. For a multiple sensor platforms control solution to be optimal (i.e. cooperative), it must be the negotiated jointly optimal group decision. For the searching problem there are two suitable candidates to evaluate a trajectory utility, namely the the MTTD (Eq. 9), and the cumulative probability of detection P_k (Eq. 8). In this paper, the later objective function is used.

Hence, for an action sequence $\mathbf{u} = \{\mathbf{u}_1, \ldots \mathbf{u}_{N_k}\}$ over a time horizon of length $T = N_k dt$, we have the following utility function

$$J(\mathbf{u}, N_k) = \sum_{i=k}^{k+N_k} p_i = P_{k+N_k} - P_k \tag{10}$$

where $\mathbf{u}$ is a $N_s \times N_k$ matrix where N_s and N_k are the number of sensors/vehicles and the number of lookahead steps respectively. The optimal control strategy $\mathbf{u}^*$ is the sequence that maximizes that utility subject to $\mathbf{u}_{LB} \leq \mathbf{u} \leq \mathbf{u}_{UB}$.

$$\mathbf{u}^* = \{\mathbf{u}_1^*, \ldots, \mathbf{u}_{N_k}^*\} = \arg\max_{\mathbf{u}} J(\mathbf{u}, N_k) \tag{11}$$

For the searching problem, because early actions strongly influence the utility of subsequent actions, the longer the time horizon, the more likely the computed trajectory is to be globally optimal. However, the computational cost follows the "curse of dimensionality" and with increasing lookahead depth becomes intractable. In practice only solutions for very restricted number of lookahead steps are possible. One way to increase the lookahead without increasing the cost of the solution too much is to have a piecewise constant (see [7] and [3]) control sequence where each control parameters is maintained over a specified number of time steps, and to recompute the planned trajectory at short intervals. Such control solutions are said to be 'quasi-optimal' as they compromise the global optimality of the control solution for a lower computation cost, but nevertheless, depending on the problem at hand, often provide better trajectories than the ones computed with the same number of control parameters but with shorter time horizons.

B. Coordinated Planning

A coordinated is different than a cooperative control solution. In a coordinated control architecture, decision makers plan individually based on their current knowledge of the world (e.g. target state PDF) and only exchange observed information via the DDF network which ensures that each platform share a common global image of the world [6]. There is no mechanism to reach a negotiated outcome, but the information exchanged between the decision makers influence each others subsequent decisions by altering the prior on which these local decisions are made. Hence coordinated trajectories are realized simply by activating a DDF network with independent control laws implemented internally at each sensor/vehicle node (fig.1).

Coordinated solutions are suboptimal, but they have the advantage of being completely decentralized. As such, because the individual planning computation costs do not increase with the number of platforms, they offer tremendous scalability potential limited only by the communication medium. Although it can be implemented for longer lookahead, the simplest form of coordinated control is for one-step lookahead. As will be demonstrated in the results section (Sec. V-B), this greedy form of coordinated searching strategy provides very sensible control solutions at very low computational costs.

V. Application

Ultimately, the goal of the ongoing research effort is to demonstrate the autonomous search framework on ACFR's fleet of unmanned air vehicles (UAV's) (fig. 2a). A stepping stone towards this goal is to use ACFR's high fidelity simulator (fig. 2b) of the UAV's hardware, complete with different sensor models, and DDF communication protocols, on which the flight software can be tested before being implemented on board the platform almost without any modifications [4].

The rest of this section describes the implementation of the coordinated Bayesian searching framework for a single lost target that could either be stationary or mobile by multiple airborne vehicles. However, the method is readily applicable to searching problems of all kinds, let it be on land, underwater or airborne search for bushfires, lost hikers, enemy troops in the battlefield, or prospection for ore and oil, or even to search for water or evidence of life on another planet.

A. Problem description

The problem chosen for the illustration of the framework involves the search for a life-raft lost at sea by a

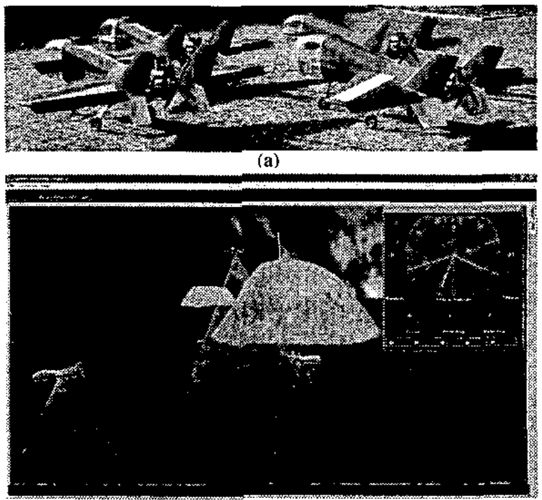

Fig. 2. (a) The fleet of Brumby Mark-III uav's been developed at ACFR as part of the ANSER project. These flight vehicles have a payload capacity of up to 13.5 kg and operational speed of 50 to 100 knots. (b) Display of the high fidelity multi-UAV simulator.

group of N_s airborne sensor platform $i = 1,...,N_s$ equipped with GPS receivers (assuming perfect localization), and a searching sensor (downward looking millimeter wave radar) that can be modelled by likelihood functions (over range and bearing) hence relating the control actions to the probability of finding the target. Each vehicle is moving in the xy plane at constant velocity V_i where the single control parameter u_k^i is the heading rate and is maintained over the time interval dt. The maximum heading rate amplitude ($u_{max} = \pm 1.1607$ rad/s) corresponds to a $6g$ acceleration, the UAV's manoeuvre limit at $V = 50$ m/s (100 knots). There is one observation (full scan) made once every second by each sensor. The sensors are assumed to have perfect discrimination i.e. no false target detection. However, they may fail to call a detection when the target is present i.e. miss contact. The omnibearing sensors' maximum range (400m) is much smaller than the size of the searching area (2km x 2km). Drift current and winds (of up to 30 knots) affect the target distribution over time in a probabilistic way through the process model. The target PDF is of general form and is evaluated and maintained on a discrete grid. As the length of the search is limited by the vehicles fuel autonomy, the search objective consists in maximizing the cumulative probability of finding the target in a fixed amount of time (Eq.10).

For details of the vehicles and the target motion models, as well as the sensors detection likelihood,the readers are referred to [2].

B. Results

For all the results presented in this section, except the last example, the initial target PDF is assumed to be a symmetric Gaussian distribution centered at the origine with a standard deviation of 500m, and except for the heterogeneous case, the searching vehicles are all flying at an altitude of 250m.

1) Stationary Target: Figure 3 shows the resulting coordinated 'greedy' (1-step lookahead) search trajectories for 2 vehicles and the corresponding 3D views of the target PDF evolution at different stages as the search progresses from 0 to 180 seconds. The fact that each vehicle build an equivalent representation of the target state PDF through DDF enables them to coordinate their actions without exchanging any information about their plans. Although this solution is very cheap computationally, it produces efficient plans that correspond to maximizing locally the individual payoff gradients. However because of the myopic planning, the vehicles fail to detect higher payoff values outside their sensor range. Figure 3e displays in solid line

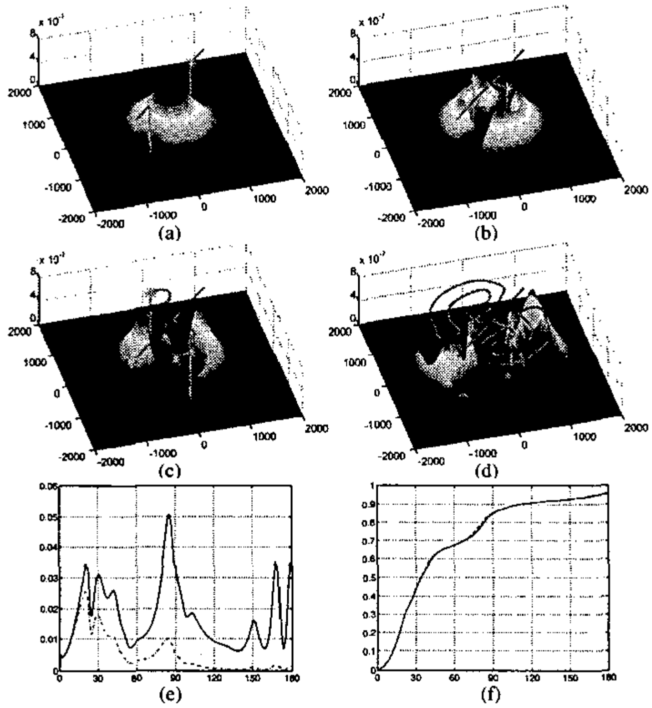

Fig. 3. Coordinated greedy search for a static target: (a) 3D view of the prior (Gaussian) target distribution and the platform locations (time t_k = 10), (b) to (d) Views of the platform trajectories and the updated target pdf at time t_k = 30, 60 and 180 respectively, (e) Conditional (solid line) and 'discounted' (dashed dotted line) probability of detecting the target on time step k ($p(\mathbf{D}_k|\mathbf{z}_{1:k}) = 1 - q_k$, and $p_k = Q_k(1-q_k)$), and (f) cumulative probability probability of detection P_k for the coordinated (solid line), and cooperative (dashed line) 1-step solutions.

the conditional probability of detection $(1 - q_k)$ obtained at every time step t_k. The dashed line represents the actual probability that the target gets detected on that time step which is the same as the solid line, but discounted by Q_{k-1} which corresponds to the payoff function p_k. The peaks in both functions are happening when the search vehicles flyby over high probability regions in the target PDF. Figure 3f shows the 'cumulative' probability P_k that the the target as been detected by time step t_k for the coordinated superimposed with the cooperative solution. This along with the computed trajectories (not shown) confirms that for very short lookahead depths, both solutions are very similar. Another phenomenon to notice about the greedy search is the fact that because the volume under the PDF is always equal to one, as the vehicles traverse a mode of the function (e.g. when they both converge to the original PDF mode for the first time (figure 3a), it has the effect of pushing away the probability mass hence increasing the entropy of the distribution, consequently making it harder and harder to increase the utility as time passes. The phenomenon will be referred to as the scattering effect. Intuitively, for a given fixed trajectory length one could imagine that, instead of rushing to the PDF's peak as in the greedy solution, the optimal strategy would be to circle around the peak but without flying over it, in such a manner as to plow the probability mass towards the peak, effectively concentrating it (reducing the entropy), in order to increase the payoff of the last observations. In fact, as shown on figure 4a, this is exactly what happen. The piecewise constant cooperative 'optimal' control solution with 6 parameters per trajectory, for a 60s plan, shows the paths spiraling in instead of spiraling out. The comparison between the utility function evolution (figure 4b) shows what one would anticipate. The greedy solution first gets a head start as each vehicle go straight to the peak to finish with $P_{60} = .673$, but the 'quasi-optimal' solution progresses steadily to ultimately finish with $P_{60} = .757$, a 12.5% increase.

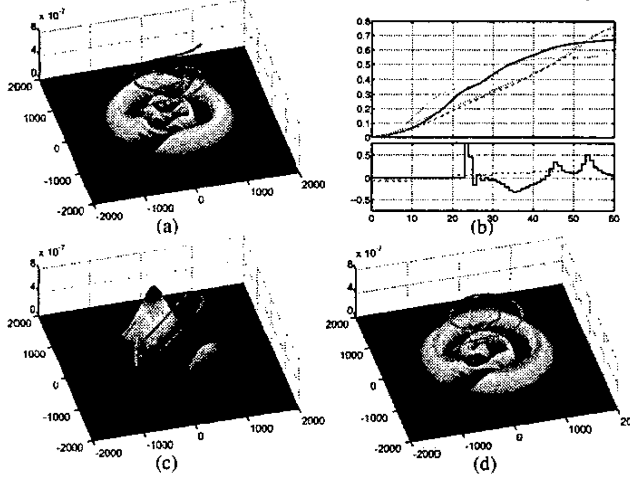

Fig. 4. Trajectory optimization: (a) Quasi-optimal cooperative trajectories for a 60s search (6 control parameters per trajectory maintained for 10s each), and (b) comparison between P_k evolution (top), and control selections $u(k)$'s (bottom) for the coordinated 1s lookahead (greedy) solution (solid line), and the 6 parameters piecewise constant solution (dashed line), (c) and (d) Greedy and quasi-optimal trajectories (12 parameters) respectively for 1 vehicle over 120s. The corresponding P_k's compressed on 60s are the dotted lines shown in (b).

Also shown in figure 4b are the single vehicle 'greedy' and 'quasi-optimal' P_k's (dotted lines) for which the corresponding trajectories are illustrated in figure 4c and d. In order to compare the cumulative probabilities for the same number of observations, the P_k's of the single vehicle are actually the results of 120s long plans displayed on the 60s long graph. Hence one can see that when the optimal trajectories are computed, two vehicles are performing about twice as fast as one vehicle, but with a very small loss in efficiency due to interference. For the 'greedy' (1-step lookahead) case the coordinated solution is also very similar to the single vehicle case, but is not necessarily worse than a single vehicle going twice as fast.

2) Drifting Target – Heterogeneous Vehicles: This section demonstrates the method for heterogeneous vehicles searching for a drifting target. A slower vehicle (V_2 = 40m/s), flying at an altitude of 600m is teaming with two faster vehicles ($V_{1,3}$ = 55m/s) flying at lower altitude (250m). Because vehicle number 2 is flying higher, it has a lower resolution (i.e. lower detection likelihood), but it has a larger field of view (800m vs 400m of ground radius). The optimization technique is the same used as for the static target, but the computational costs are increased by a few fold as the convolution operation needed for the target prediction stage is a costly operation. This is also compounded by the fact that because the target PDF is moving, a larger grid is necessary, making it even more costly to perform the convolution and the optimization. Nevertheless, the coordinated greedy solution is still very effective. The 3D plots of the search evolution are shown on figure 5. Once again, the coordinated solution shows quite reasonable trajectories terminating with $P_{150} = 95.5\%$ (figure 5d).

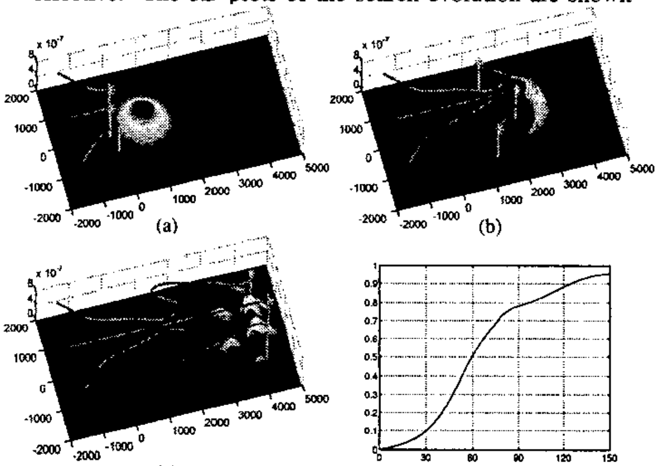

Fig. 5. Coordinated (1s lookahead) search for a drifting target with 3 heterogeneous vehicles: (a) to (c) 3D views of the searching vehicles trajectories and updated target PDF at time t_k = 30, 90, and 150 respectively, and (d) Cumulative probability of detection P_k.

3) Scalability: In this example the real strength of the coordinated control strategy is demonstrated for 10 vehicles searching for a stationary target without increasing the computation cost at each node. Figure 6 illustrates the evolution of the coordinated trajectories. By allowing a more efficient

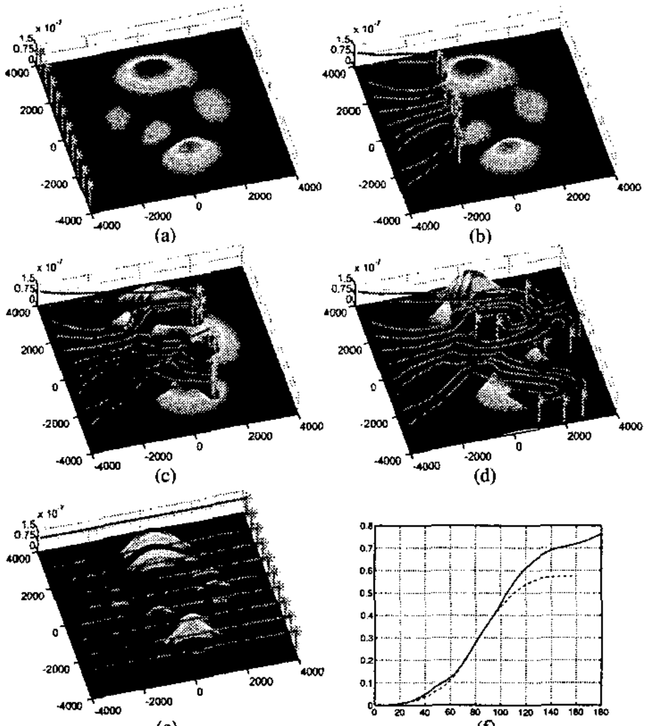

Fig. 6. Coordinated search for a static target with 10 vehicles:(a) to (d) 3D views of the target PDF and the coordinated trajectories evolution at time t_k = 1, 60, 120 and 180s respectively, (e) Straight pattern search at t_k = 160s, and (f) P_k for the coordinated (solid line) vs. the flight formation (dashed line) search.

allocation of the search effort, the coordinated approach compares advantageously to the simple scanning search strategy shown on figure 6e which is somewhat reminiscent of current searching strategies followed by coastguards and Navies. In fact after 160s, the time needed for the formation to traverse the searching area, P_k reaches a value of only $P_{160} = .575$ vs. $P_{160} = .718$ for the coordinated solution, a 24.9% increase (fig. 6f).

VI. Summary and Ongoing Work

This paper addressed the problem of coordinating multiple possibly heterogeneous sensing platforms performing a search mission for a single target in a dynamic environment. The general decentralized Bayesian framework presented explicitly considers the search vehicles kinematics, the sensors detection function, as well as the target arbitrary motion model. It was demonstrated to adaptively find efficient coordinated search plans in a completely decentralized way. A major appeal of the approach is that nodal computation costs are kept constant regardless of their number thus offering a high potential for scalability.

Because of the nature of the search problem, it is quite important to accurately keep track of the very non-Gaussian target state PDF. However, any grid based approach such as the one presented is intrinsically subject to the "curse of dimensionality", and as soon as one needs to increase the search area, the resolution of the grid, or the number of dimensions in the state-space, computational costs tend to get out of hand. As part of the ongoing research effort, techniques such as Monte Carlo methods, or particle filters [5], as well as the so called kernel methods are being investigated to overcome the computational limitations.

Another limitation of the technique as presented comes from the assumption that every DDF node transmits and receives every single observation without a miss via broadcasting. Beyond the obvious bandwidth limitations, such assumptions are not quite practical in real life since communication systems are plagued by delays and intermittent transmissions. To overcome this problem, work in progress also involves developing a channel filter to allow the Bayesian DDF network to be tree connected and hence reduce drastically the communication loads that are incurred in a fully connected network, as well as allowing intermittent burst communications.

Beyond the demonstration of the approach on a team of UAV's, the ultimate objective of this research is to eventually have multiple platforms participating in actual search and rescue (SAR) missions with real-time cooperative planning and fully integrated human in the loop inputs. As shown by the results presented, the technique as the potential to greatly improve on current SAR protocols, which in turn could be critical in saving human lives.

VII. Acknowledgments

The authors wish to thank Ali Göktoğan from ACFR, the developer of the RMUS simulator, for his assistance with the simulation implementation.

VIII. References

[1] J.O. Berger. *Statistical decision theory and Bayesian analysis*. Springer series in statistics. Springer-Verlag, New York, 2nd edition, 1985.

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

[3] T. Furukawa. Time-subminimal trajectory planning for discrete nonlinear systems. *Engineering Optimization*, 34:219–243, 2002.

[4] A.H. Goktogan, E. Nettleton, M. Ridley, and S. Sukkarieh. Real time multi-uav simulator. In *IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, 2003.

[5] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings-F*, 140(2):107–113, April 1993.

[6] B. Grocholsky, A. Makarenko, and H. Durrant-Whyte. Information-theoretic coordinated control of multiple sensor platforms. In *IEEE Int. Conf. on Robotics and Automation (ICRA'03)*, 2003.

[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] J.S. Przemieniecki. *Mathematical Methods in Defense Analyses*. AIAA Education Series. American Institute of Aeronautics and Astronautics, Inc., Washington, DC, 2nd edition, 1994.

[9] John Rousmaniere. *Fastnet Force 10*. Nautical Publishing Co Ltd, Lymington, Hamsphire, GB, 1980.

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

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

[12] K.L. Teo, C.J. Goh, and K.H. Wong. *A Unified Computational Approach to Optimal Control Problems*. Longman Scientific and Technical, 1991.