Automatica 50 (2014) 1135–1142

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Sensor management for multi-target tracking via multi-Bernoulli filtering ☆

Hung Gia Hoang, Ba Tuong Vo

Department of Electrical and Computer Engineering, Curtin University, Perth WA 6845, Australia

ARTICLE INFO

Article history:
Received 15 December 2012
Received in revised form
14 October 2013
Accepted 23 December 2013
Available online 26 February 2014

Keywords:
Multi-target tracking
Sensor control
Random finite sets
Sequential Monte Carlo method

ABSTRACT

In multi-object stochastic systems, the issue of sensor management is a theoretically and computationally challenging problem. In this paper, we present a novel random finite set (RFS) approach to the multi-target sensor management problem within the partially observed Markov decision process (POMDP) framework. The multi-target state is modelled as a multi-Bernoulli RFS, and the multi-Bernoulli filter is used in conjunction with two different control objectives: maximizing the expected Rényi divergence between the predicted and updated densities, and minimizing the expected posterior cardinality variance. Numerical studies are presented in two scenarios where a mobile sensor tracks five moving targets with different levels of observability.

1. Introduction

Multi-target sensor control/management is essentially an optimal non-linear control problem. The goal of multi-target sensor management is to "direct the right sensor on the right platform to the right target at the right time" (Mahler, 2003b). However, the multi-target sensor control problem differs from the classical control problem in that it deals with highly complex multi-object stochastic systems. In multi-object stochastic systems, not only do the number of objects vary randomly in time, but the measurements are subject to missed detections and false alarms. This means that the multi-target state and multi-target observation are inherently finite-set-valued. Consequently, standard optimal control techniques are not directly applicable (Mahler, 2004). Nonetheless, the multi-target sensor scheduling problem can still be cast in the framework of partially observed Markov decision processes (POMDPs), where the states and observations are instead finite-set-valued, and control vectors are drawn from a set of admissible sensor actions based on the current information states, which are then judged against the values of an objective function associated with each action (Castanón & Carin, 2008).

A unified approach to characterizing systems with finite-set-valued states is the multi-object systems framework, where uncertainty is described by multi-object probability density functions, and formalized via point process theory (Daley & Vere-Jones, 1988; Stoyan, Kendall, & Mecke, 1995), or equivalently by random finite set (RFS) theory through Mahler's finite set statistics (FISST) (Mahler, 2007b). The key advantage of the RFS based approach is that of a principled framework for modelling, estimation and control of multi-object systems. In this paper, we formulate the sensor control problem as a POMDP with an information-theoretic objective function as well as finite-set-valued states and measurements. In essence, our approach can be summarized by three basic steps:

(1) Modelling the sensor and targets as a multi-object stochastic system, i.e. the multi-target states and multi-target observations as RFSs;
(2) propagating the multi-object posterior density recursively in time, or alternatively a tractable approximation to the posterior;
(3) at each time, determining the control action based on optimization of the reward function over a set of admissible actions.

In the context of single-target tracking, the work in Doucet, Vo, Andrieu, and Davy (2002) is the first to propose a practical particle implementation based on the Kullback–Leibler divergence, and the approach in Singh, Kantas, Vo, Doucet, and Evans (2007)

☆ The first author is partially supported by Curtin's matching funding to the Australian Research Council Future Fellowship FT0991854 and the second author is supported by Discovery Early Career Researcher Award DE120102388. The material in this paper was partially presented at the First International Conference on Control, Automation and Information Sciences (ICCAIS 2012), November 26–29, 2012, Ho Chi Minh City, Viet Nam. This paper was recommended for publication in revised form by Associate Editor Giancarlo Ferrari-Trecate under the direction of Editor Ian R. Petersen.

E-mail addresses: hung.hoang@curtin.edu.au (H.G. Hoang), ba-tuong.vo@curtin.edu.au (B.T. Vo).

http://dx.doi.org/10.1016/j.automatica.2014.02.007

further considers the issue of observer trajectory planning. In the more difficult multi-target context, there are a handful of works falling within the RFS framework. Using the Rényi divergence as the reward function, in Ristic and Vo (2010) the particle multi-object Bayes filter (Vo, Singh, & Doucet, 2005) is used to propagate the multi-object posterior, while in Ristic, Vo, and Clark (2011) the particle probability hypothesis density (PHD) filter (Vo et al., 2005) is used to propagate the first moment of the multi-object posterior.

This paper adopts an information theoretic approach for multi-target sensor control, similar to the approach in Ristic and Vo (2010) and Ristic et al. (2011), except that the Cardinality Balanced Multi-Target Multi-Bernoulli (CB-MeMBer) filter (Vo, Vo, & Cantoni, 2009) is used to propagate a parametrized approximation to the multi-object posterior. The proposed approach is attractive in that it is applicable to general non-linear non-Gaussian models, and when coupled with a particle implementation further reduces the computational load significantly. Propagating an approximate multi-Bernoulli posterior as proposed is drastically cheaper than propagating the full multi-object posterior as in Ristic and Vo (2010), and thus the computation of any associated cost function using a multi-Bernoulli approximation is generally cheaper than using the full posterior. While the proposed use of the CB-MeMBer filter incurs the same complexity as the use of the PHD filter for the same purpose in Ristic et al. (2011), performing state estimation with the former is more efficient and reliable than the latter because the need for clustering is eliminated. The work in Ristic and Vo (2010) and Ristic et al. (2011) also demonstrates that the Rényi divergence can be used as a reward function for multi-target sensor control. In the same regard, the use of the CB-MeMBer filter equally allows the Rényi divergence to be used as a reward function, and further allows a new type of the reward function to be developed. Since the variance of the estimated cardinality of a multi-Bernoulli posterior can be evaluated in closed form (Mahler, 2007b), minimizing the cardinality variance can be used as the control objective, thereby enabling direct control of the cardinality estimation error.

The main contribution of this paper is a computationally efficient sensor control algorithm for multiple targets, using the CB-MeMBer filter, as well as the numerical assessment of two types of control objectives. Our preliminary result, in particular the idea of using the CB-MeMBer filter, has been reported in the conference paper (Hoang, 2012). The current paper provides full details of the algorithm, an alternative cheaper control objective, and more complete numerical studies.

The organization of the paper is as follows. In Section 2 we review RFS modelling of multi-object systems and the approximation of the multi-object posterior density using the CB-MeMBer filter. The two reward functions are discussed in Section 3 while sequential Monte Carlo (SMC) implementation is described in Section 4. Section 5 presents simulation results and finally, Section 6 concludes the paper.

2. Cardinality balanced MeMBer filter

In this section, we summarize the CB-MeMBer filter, the main tool that will be used throughout the paper. The filter was originally introduced in Vo et al. (2009) to account for the cardinality bias of the MeMBer filter in Mahler (2007b).

2.1. General system model

In contrast with single-object systems where the states and observations are modelled by random vectors, the states and observations of a multi-object system are random finite sets of vectors in the single-object state space $\mathcal{X} \subseteq \mathbb{R}^n$ and single-object observation space $\mathcal{Z} \subseteq \mathbb{R}^m$, respectively:

$$\mathbf{X}_k = \{\mathbf{x}_1^k, \dots, \mathbf{x}_n^k\} \in \mathcal{F}(\mathcal{X}); \tag{1}$$

$$\mathbf{Z}_k = \{\mathbf{z}_1^k, \dots, \mathbf{z}_m^k\} \in \mathcal{F}(\mathcal{Z}). \tag{2}$$

Here $\mathcal{F}(\mathcal{X})$ and $\mathcal{F}(\mathcal{Z})$ denote the spaces of all finite subsets of $\mathcal{X}$ and $\mathcal{Z}$. The system is described by the following probabilistic state space model:

$$\mathbf{X}_k \sim \pi_{k|k-1}(\mathbf{X}_k|\mathbf{X}_{k-1}) \tag{3}$$

$$\mathbf{Z}_k \sim g_k(\mathbf{Z}_k|\mathbf{X}_k) \tag{4}$$

where $\mathbf{X}_k$ and $\mathbf{Z}_k$ respectively are the state and observation of the system at time k. Eq. (3) describes the system dynamics encapsulating all aspects of object birth, death and transition while Eq. (4) encapsulates all aspects of sensor detection and false alarms.

Given the system model (3)–(4), the objective is to determine at each time step k the multi-object posterior probability density $f_k(\mathbf{X}_k|\mathbf{Z}_{1:k})$. In the Bayesian filtering framework, $f_k(\mathbf{X}_k|\mathbf{Z}_{1:k})$ is obtained through two steps: time prediction and measurement update (Mahler, 2007b). The predicted density at time k, denoted as $f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}_{1:k-1})$, is computed by the multi-object Chapman–Kolmogorov equation:

$$f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}_{1:k-1}) = \int \pi_{k|k-1}(\mathbf{X}_k|\mathbf{X}_{k-1}) f_{k-1}(\mathbf{X}_{k-1}|\mathbf{Z}_{1:k-1}) \delta\mathbf{X}_{k-1} \tag{5}$$

where $f_{k-1}(\mathbf{X}_{k-1}|\mathbf{Z}_{1:k-1})$ is the posterior density from the previous time step $k-1$. When new observations arrive at the sensor(s), the new posterior density is computed via the multi-object Bayes rule:

$$f_k(\mathbf{X}_k|\mathbf{Z}_{1:k}) = \frac{g_k(\mathbf{Z}_k|\mathbf{X}_k) f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}_{1:k-1})}{\int g_k(\mathbf{Z}_k|\mathbf{X}_k) f_{k|k-1}(\mathbf{X}_k|\mathbf{Z}_{1:k-1}) \delta\mathbf{X}_k}. \tag{6}$$

Notice that the integrals in (5) and (6) are not ordinary integrals, but are set integrals, and that the recursion (5) and (6) has no analytic solution in general. A sequential Monte Carlo (SMC) implementation of the Bayes multi-object filter is given in Vo et al. (2005). However, this technique is computationally prohibitive which at best is able to accommodate a small number of targets. The multi-target sensor scheduling algorithm proposed in Ristic and Vo (2010) employs this SMC implementation of the multi-object Bayes filter.

Since propagation of the full posterior density given by (6) is in general intractable, several alternatives have been proposed, which propagate only summary statistics or important parameters in place of the full posterior density. For example, the PHD and Cardinalized PHD (CPHD) filters (Mahler, 2003a, 2007a; Vo & Ma, 2006; Vo et al., 2005; Vo, Vo, & Cantoni, 2007) propagate the intensity or first order moment of the posterior density, and were employed by the multi-target sensor scheduling approach in Ristic et al. (2011). An alternative is the CB-MeMBer filter (Mahler, 2007b; Vo et al., 2009), which propagates a parametrized multi-Bernoulli approximation of the multi-object posterior density. The main advantage of the CB-MeMBer approach is its direct applicability to non-linear non-Gaussian models, which when coupled with an SMC implementation, obviates the need for the clustering of the particle population in order to extract estimates.

2.2. CB-MeMBer recursion

We now summarize the recursion for the CB-MeMBer filter. A Bernoulli RFS $\mathbf{X}$ has realizations either as the empty set or a singleton and is characterized jointly by a probability of existence $r \in [0, 1]$ and a probability density p. That is, the Bernoulli RFS takes on a singleton value with probability r, and conditional upon existence, the value of the singleton is distributed according to