Signal Processing 119 (2016) 115–127

Contents lists available at ScienceDirect

Signal Processing

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

Algorithm design for parallel implementation of the SMC-PHD filter

Tiancheng Li [a,b,*], Shudong Sun [b], Miodrag Bolić [c], Juan M. Corchado [a,d]

[a] The BISITE group, Faculty of Science, University of Salamanca, 37008 Salamanca, Spain
[b] School of Mechanical Engineering, Northwestern Polytechnical University, 710072 Xi'an, PR China
[c] School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, Ontario, Canada K1N 6N5
[d] Osaka Institute of Technology, Asahi-ku Ohmiya, 535-8585 Osaka, Japan

ARTICLE INFO

Article history:
Received 28 November 2014
Received in revised form
7 July 2015
Accepted 16 July 2015
Available online 4 August 2015

Keywords:
Multi-target tracking
Particle filter
Probability hypothesis density filter
Parallelization

ABSTRACT

The sequential Monte Carlo (SMC) implementation of the probability hypothesis density (PHD) filter suffers from low computational efficiency since a large number of particles are often required, especially when there are a large number of targets and dense clutter. In order to speed up the computation, an algorithmic framework for parallel SMC-PHD filtering based on multiple processors is proposed. The algorithm makes full parallelization of all four steps of the SMC-PHD filter and the computational load is approximately equal among parallel processors, rendering a high parallelization benefit when there are multiple targets and dense clutter. The parallelization is theoretically unbiased as it provides the same result as the serial implementation, without introducing any approximation. Experiments on multi-core computers have demonstrated that our parallel implementation has gained considerable speedup compared to the serial implementation of the same algorithm.

1. Introduction

Multi-target tracking (MTT) involves the joint estimation of the number of multiple targets and their states in the presence of spontaneous birth/spawn/death of targets and clutter. MTT has a long history of research over a half of century, with many applications in both military and commercial realms [1]. Apart from handling the noises in the state dynamics and observation models, one has to contend with many more challenges, such as: (i) The number of targets is unknown and time varying because of the spontaneous birth, spawn and death of targets; (ii) clutter exists and can be significant; (iii) targets can be miss-detected; (iv) most challenging, data association between observations and targets in the presence of clutter that is required in traditional trackers is difficult.

The states of targets and observations in such an environment are finite-set-valued random variables that are random in both the number of elements and the values of the elements. The idea of modelling the states and observations as random finite set (RFS) is natural and it allows for overcoming the difficulty of data association [2,3] in the filtering stage. With the incorporation of RFS and point process theory in the MTT problem, the probability hypothesis density (PHD) filter provides a concise and tractable alternative to the optimal Bayesian filter that works by propagating the first-order moment associated with the multi-target Bayesian posterior and is essentially a density estimator.

* Correspondence to: Edificio I+D+I, C/ Espejo, s/n, 37008 Salamanca, Spain (Office). Tel.: +34 923 012 045; mobile: +34 655 24 8188; fax: +34 923 294 514.
E-mail addresses: tiancheng.li1985@gmail.com, t.c.li@usal.es (T. Li), sdsun@nwpu.edu.cn (S. Sun), mbolic@site.uottawa.ca (M. Bolić), corchado@usal.es (J.M. Corchado).

http://dx.doi.org/10.1016/j.sigpro.2015.07.013
0165-1684/

The PHD filter is attracting increasing attention, motivating different derivations, interpretations and implementations. It is found that the PHD filter asymptotically behaves as a mixture of Gaussian components, whose number is the true number of targets, and whose peaks collapse in the neighbourhood of the classical maximum likelihood estimates, with a spread ruled by the Fisher information [4]. A connection between the PHD recursion and spatial branching processes is established in [5], which gives a generalized Feynman–Kac systems interpretation of the PHD recursive equations and enables the derivation of mean-field implementations. A physical interpretation of the PHD filter based on the bin model is given in [6] which gives a more intuitive understanding of the PHD filter. The PHD filter has also been extended to solve more MTT-centric complex problems such as distributed sensor localization [7], parameter estimation [8] and mobile robot simultaneous localization and mapping [9], to name a few. Presently, the PHD filter has been implemented in forms of weighted particles [10], finite Gaussian mixtures (GM) [11] or their hybrids [12]. In particular, the sequential Monte Carlo (SMC) implementation that is often referred to the SMC-PHD filter is gaining special attention [13–15], which has also prompted new developments of SMC based on the RFS framework e.g. [16].

The advantage of SMC over GM is that it is free of linear and Gaussian assumptions. However, to maintain a sufficient approximation accuracy, a large number of particles are usually required causing a heavy computational burden. The situation gets worse in the SMC-PHD filter where the computational requirements also grow with the number of targets/observations [17]. Thence, fast computing techniques such as gating [18,19] and parallel processing [20–23] appear as promising approaches to ensure real-time performance, which however are often based on significant approximations. With the fast development of multi-core and multi-threading compatible hardware and software, the parallelization becomes increasingly attractive and even necessary. To the best of our knowledge, a fully parallel implementation of the SMC-PHD filter that is able to provide the same result as the serial implementation is still lacking. Such a parallel SMC-PHD filter is the focus of this paper.

There are four main steps in the implementation of particle filters (PFs) including state updating, weight updating, resampling and estimate extraction. The essence of parallelization is distributing calculation operations to different processing elements (PEs) for parallel computing. In the conventional parallel implementation of particle filters, particles are distributed among PEs. However, even for the PF targeted for single target tracking (referred to the basic PF hereafter), the resampling prevents direct parallel processing due to joint processing of all particles. The parallelization technique for resampling achieved in these particle filters e.g. [24–25] may be applied in the MTT-based particle PHD filter as they have been done in [20–22] at the price of inevitable considerable communication overhead. However, there are more challenging operations that consist of joint processing of particles in the SMC-PHD filter in addition to the resampling. First, the weight updating of each particle in the PHD updater requires the weight information of all the other particles thereby preventing the parallelization that distributes particles among PEs. Secondly, what used in [20–23] for extracting multiple estimates in the SMC-PHD filter is still the (k-means) clustering method which is computationally intensive and is less suitable for parallelization. It is unclear how the k-means clustering can be parallelized without introducing significant communication overhead. These operations together with the resampling step form the primary challenges for parallelization of the SMC-PHD filter.

To overcome the parallel computing difficulty in resampling, weight updating and estimate extraction, these steps are parallelized in a novel manner that is based on distributing observations instead of distributing particles while the prediction is carried out on the distributed resampling particles in each PE. In particular, new resampling and estimate extraction methods that are suitable for parallel processing are proposed. Significantly different from [20–23], our parallelization is able to obtain the same estimation result as the serial implementation. All four steps of the SMC-PHD filter are fully paralleled and significant speedup is achieved. The parallel algorithm is described with regard to the multi-core computer that consists of one central unit (CU) and several PEs.

The paper is arranged as follows. A novel insight of the SMC-PHD filter that partitions the weight of particles into components with respect to observations is provided in Section 2, which forms the foundation of our approach. Related works are also briefly described in this section. The technical details of our approach are presented in Section 3. Qualitative analysis and quantitative experiments are given in Sections 4 and 5 respectively. We conclude in Section 6.

2. Background and related work

2.1. An observation-based view of the PHD equation

To model the discrete filtering problem, the state is assumed to follow a Markov process in the state space $\chi \subseteq \mathbb{R}^{n_x}$, with a transition density $f_{k|k-1}(\cdot|\cdot)$. That is, a given state x_{k-1} at time $k-1$ will either die with probability $1-p_{S,k}(x_{k-1})$ or continue to exist at time k with survival probability $p_{S,k}(x_{k-1})$ and move to a new state with a transition probability density $f_{k|k-1}(x_k|x_{k-1})$. The Markov process is partially observed in the observation space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$. That is, at time k, a given target $x_k \in X_k$ is either miss-detected with probability $1-p_{D,k}(x_k)$ or detected with detection probability $p_{D,k}(x_k)$ and generates an observation $z_k \in Z_k$ with likelihood $g_k(z_k|x_k)$.

The collections of target states and observations at scan k can be represented as finite sets $X_k = \{x_{k,1}, \ldots, x_{k,N_k}\} \in F(\chi)$ and $Z_k = \{z_{k,1}, \ldots, z_{k,M_k}\} \in F(\mathcal{Z})$, where N_k and M_k are the number of targets and the number of observations respectively, while $F(\chi)$ and $F(\mathcal{Z})$ are the collections of all finite subsets of targets and observations, respectively.

Based on the finite set statistics, the PHD filter derived by Mahler estimates jointly the number and the state of targets by propagating in time the intensity function (namely the PHD function) [2]. The following assumptions are required in the PHD filter: (A.1) each target is assumed to evolve and generate observations independently of others; (A.2) the clutter distribution is assumed to be