3-2015

Tracking of Rhythmical Biomedical Signals Using the Maximum A Posteriori Adaptive Marginalized Particle Filter

Sunghan Kim
Portland State University

Lars Andreas Holmstrom
Portland State University

James McNames
Portland State University

Citation Details

Kim, Sunghan, Lars Holmstrom, and James McNames. "Tracking of Rhythmical Biomedical Signals Using the Maximum A Posteriori Adaptive Marginalized Particle Filter." British Journal of Health Informatics and Monitoring 2.1 (2015).

RESEARCH ARTICLE

The full electronic version of this research article with all related policies and conditions can be found online at: http://www.bjhim-online.org.

Received: 30 October 2014
Accepted: 04 February 2015
Published: 31 March 2015

TRACKING OF RHYTHMICAL BIOMEDICAL SIGNALS USING THE MAXIMUM A POSTERIORI ADAPTIVE MARGINALIZED PARTICLE

[a]Sunghan Kim, PhD, [b]Lars Holmstrom, PhD, [c]James McNames, PhD

[a]Biomedical Concentration, Department of Engineering, East Carolina University, Greenville, NC, USA
[a,b,c]Biomedical Signal Processing Laboratory, Department of Electrical and Computer Engineering, Portland State University, Portland, OR, USA
[a]kims@ecu.edu, [b]larsh@pdx.edu, [c]mcnames@pdx.edu

Abstract

Biomedical signals are often rhythmical and their morphologies change slowly over time. Arterial blood pressure and electrocardiogram signals are good examples with such property. It is of great interest to extract clinically useful information such as the instantaneous frequency (i.e. heart rate) and morphological changes (e.g. pulse pressure variation) from these signals. Conventional filtering methods such as the Kalman filter are not suitable for estimating the instantaneous frequency of quasi-periodic signals due to the non-Gaussian multi-modal property of its posterior distribution. One possible alternative is particle filters that are increasingly used for nonlinear systems and non-Gaussian posterior state distributions. However, canonical particle filters suffer from the problems of sample degeneracy and sample impoverishment and are not well suited to non-Gaussian multi-modal distributions. This paper describes two new algorithms that integrate the marginalized particle filter and maximum a-posterior particle filter and demonstrates challenging cases where the proposed algorithms outperform the conventional marginalized particle filter using both synthetic and real electrocardiogram signals.

ISSN 2057-0058 (Online)

Keywords: Extended Kalman filter, a-posterior distribution, maximum a-posterior estimation, marginalized particle filter, quasi-periodic

1 INTRODUCTION

Biomedical signals contain quasi-periodic rhythms with slowly changing morphologies. Given arterial blood pressure (ABP) and electrocardiogram (ECG) signals, for example, extracting clinically useful information such as the instantaneous frequency (i.e. heart rate) and morphological changes (e.g. pulse pressure variation) is of great interest to medical doctors and researchers [16]. Previously, we have proposed the Kalman filter based approach to tracking the rhythmicity of various biomedical signals such as ABP, ECG, and microelectrode recording (MER) [17, 18, 19]. As Sivia pointed out, however, the Kalman filter based methods are not suitable for estimating the instantaneous frequency of quasi-periodic signals [26]. The reason is that the posterior distribution of the instantaneous frequency of quasi-periodic signals exhibits the non-Gaussian multi-modal property, which imposes a challenge for the Kalman filter based methods since they rely on local linearization of the posterior distribution. This multimodal property of the frequency's posterior distribution not only raises a theoretical issue for the Kalman filter based methods, but does cause real problems for them in a way that once they lose the track of the instantaneous frequency, they can hardly regain it [19].

Monte Carlo methods are a possible alternative [21]. They can be used to estimate the point statistics of an unknown or analytically unsolvable posterior distribution, $p(\boldsymbol{x}_{0:n}|\boldsymbol{y}_{0:n})$, up to a normalizing constant based on a sequence of sufficient random samples drawn from the distribution, where the vector, $\boldsymbol{y}_{0:n}$, represents a sequence of observations, $\{y_0, \ldots, y_n\}$. Gordon et al. proposed a sequential version of Monte Carlo methods, which is better known as particle filters (PFs) [11]. PFs have been applied to various applications such as fault detection [1], harmonic tracking for audio signals [7], computer vision [25], speech recognition [28], and target localization [30]. In general, a large number of particles yield more accurate approximation results since the approximation of PFs converges to the true value asymptotically as the number of particles, N_p, becomes large enough. However, this is problematic since PFs are plagued by two curses of dimensionality: the number of particles, N_p, needed for reasonable approximations scales with both the state dimension, ℓ, and the duration of the observation sequence, n.

The marginalized particle filter (MPF) in [3, 6, 14, 20] can partly overcome the first curse, in special cases, where the state space can be partitioned into a high-dimensional linear portion and a low-dimensional nonlinear portion. The Kalman filter can handle the high-dimensional linear portion of the state space efficiently, which eases the overall computational burden. The marginalization of the state space, however, does not solve the second curse of dimensionality; the number of particles required to accurately estimate an expected value of the distribution still increases exponentially with the duration of the observations n at the time the estimate is needed. Practically this

ISSN 2057-0058 (Online)

causes a phenomenon called sample degeneracy. Resampling schemes have been proposed to address this problem, which essentially use a bootstrap approach to sample the posterior distribution with replacement. However, this approach causes a sample impoverishment problem referring to a phenomenon where many particles of only the most probable state trajectory are duplicated and the coverage of all possible state trajectories lessens. If the true posterior distribution makes an abrupt change or develops a new prominent mode in a region not covered by the particle population, a typical PF may be slow to adapt to the abrupt change or fail to lock on to the new prominent mode. This issue can be well addressed with by incorporating maximum a posteriori (MAP) estimation into particle filtering recursions. MAP-PFs produce an estimate of the most likely state trajectory instead of the mean state trajectory, which maximizes the posterior distribution,

$$\widehat{\boldsymbol{x}}_{0:n} \approx \operatorname{argmax}_{\boldsymbol{x}_{0:n}}\{p(\boldsymbol{x}_{0:n}|\boldsymbol{y}_{0:n})\}. \quad (1)$$

MAP estimation is suitable for multi-modal posterior distributions because it essentially selects the state estimates corresponding to the tallest mode of the distribution. MAP estimates minimize the most probable error, whereas mean estimates minimize mean square error (MSE).

Several research groups have used the term “marginal MAP estimation” to describe algorithms for various applications such as fault diagnosis for autonomously operating systems [12], state estimation of jump Markov linear systems [5], tempo tracking and rhythm quantization in music [4], and detection for Orthogonal Frequency Division Modulation (OFDM) systems [29]. Doucet et al. describe an algorithm that obtains the marginal MAP estimate of the state of a Jump Markov Linear System based on Markov chain Monte Carlo (MCMC) methods [5]. Cemgil et al. discussed the possibility of computing the MAP trajectory after integrating out (Rao-Blackwellizing) the hidden variables based on the SMC methods [4]. However, they were aware that Rao-Blackwellization causes coupling between all possible particle trajectories, and that the Viterbi algorithm does not find the actual MAP trajectory in this case. In simple words, combining the MPF and MAP-PF methods introduces new problems that are not encountered in using those two methods independently.

The objectives of this paper are to introduce an optimal algorithm and a fast approximate algorithm for combining the advantages of marginalized particle filters (MPF) and MAP particle filters (MAP-PF) and to demonstrate challenging cases where the conventional marginalized particle filters (MPF) fail to track the instantaneous frequency while the proposed methods succeed using synthetic and real ECG signals. It should be noted that the present work is a part of our continuous effort to develop robust tracking methods that can be applied to quasiperiodic multi-harmonic biomedical signals and readers are encouraged to refer to our previous works as necessary [17, 18].

ISSN 2057-0058 (Online)

2 METHODOLOGY

2.1 Notation

We have adopted the notation used in [2] with minor modifications. We use boldface to denote random processes, normal face for deterministic parameters and functions, upper case letters for matrices, lower case letters for vectors and scalars, superscripts in parenthesis for particle indices, uppercase superscripts for nonlinear/linear indication, and subscripts for time indices. For example, the nonlinear portion of the state vector for the i^{th} state trajectory (i.e., particle) is denoted as $\boldsymbol{x}_n^{\mathrm{N},(i)}$ where n represents the discrete time index and (i) denotes the i^{th} particle. The unnormalized particle weights are denoted as $\widetilde{w}^{(i)}$ and the normalized particle weights as $w^{(i)}$. The state trajectories before resampling are denoted as $\widetilde{\boldsymbol{x}}_n^{(i)}$ and as $\boldsymbol{x}_n^{(i)}$ after resampling.

2.2 Recursions for typical particle filtering methods

2.2.1 Standard resampling particle filter (PF)

The standard PF algorithm incorporates the stratified resampling scheme in which resampling is performed when the estimated number of effective particles drops below a user-defined threshold, N_t [13]. The number of effective particles is given by

$$\widehat{\boldsymbol{N}}_{\mathrm{e}} = \frac{1}{\sum_{i=1}^{N_{\mathrm{p}}} \left(\boldsymbol{w}_n^{(i)}\right)^2} \tag{2}$$

where $\boldsymbol{w}_n^{(i)}$ is a normalized importance weight of the i^{th} particle and N_{p} is the number of particles. The normalized importance weights $\boldsymbol{w}_n$ are calculated at each time index from the unnormalized importance weights $\widetilde{\boldsymbol{w}}_n$, which are themselves derived from the normalized importance weights from the previous time index. We assume that the importance density has been chosen such that it can be factored as follows,

$$q(\boldsymbol{x}_{0:n}|\boldsymbol{y}_{0:n}) = q_n(\boldsymbol{x}_n|\boldsymbol{x}_{0:n-1}, \boldsymbol{y}_{0:n}) q(\boldsymbol{x}_{0:n-1}|\boldsymbol{y}_{0:n-1}) \tag{3}$$

so that the particle weights can be calculated recursively.

Then, the weight update recursion can be written as,

$$\widetilde{w}_n^{(i)} = \widetilde{w}_{n-1}^{(i)} \frac{p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{(i)}\right) p\left(\boldsymbol{x}_n^{(i)} \middle| \boldsymbol{x}_{n-1}^{(i)}\right)}{q_n\left(\boldsymbol{x}_n^{(i)} \middle| \boldsymbol{x}_{n-1}^{(i)}, \boldsymbol{y}_{0:n}\right)} \tag{4}$$

ISSN 2057-0058 (Online)

where $\widetilde{w}_n^{(i)}$ represents an unnormalized importance weight at time n and $\widetilde{w}_{n-1}^{(i)}$ a normalized importance weight at time $n-1$. The most common choice for the marginal importance density is the prior probability given by the process model,

$$q_n(\boldsymbol{x}_n|\boldsymbol{x}_{n-1}, \boldsymbol{y}_{0:n}) = p(\boldsymbol{x}_n|\boldsymbol{x}_{n-1}) \tag{5}$$

The weight update recursion, then, is simplified as follows,

$$\widetilde{w}_n^{(i)} = \widetilde{w}_{n-1}^{(i)} \frac{p\left(\boldsymbol{y}_n\middle|\boldsymbol{x}_n^{(i)}\right)p\left(\boldsymbol{x}_n^{(i)}\middle|\boldsymbol{x}_{n-1}^{(i)}\right)}{q_n\left(\boldsymbol{x}_n^{(i)}\middle|\boldsymbol{x}_{n-1}^{(i)}, \boldsymbol{y}_{0:n}\right)} = \widetilde{w}_{n-1}^{(i)} p\left(\boldsymbol{y}_n\middle|\boldsymbol{x}_n^{(i)}\right) \tag{6}$$

where the current weight is the previous weight multiplied by the likelihood function, $p\left(\boldsymbol{y}_n\middle|x_n^{(i)}\right)$.

2.2.2 Maximum a posteriori particle filter (MAP-PF)

The MAP-PF approach based on the Viterbi algorithm provides a MAP estimate that avoids the sample degeneracy and sample impoverishment problems of the standard PF [2, 9, 27]. The algorithm is memory-efficient because at each sample time the Viterbi algorithm discards ($N_{\mathrm{p}}^2 - N_{\mathrm{p}}$) possible trajectories and only retains the Np most probable ones. Unlike the standard particle filter (PF), the MAP-PF does not calculate or track importance weights for each particle. These are unnecessary because the MAP-PF does not estimate the mean or other moments of the state posterior distribution. Instead the MAP-PF simply tracks the posterior probability of each state trajectory. These are represented by the coefficients $\boldsymbol{\alpha}_n^{(i)}$, which are computed recursively for each particle instead of the importance weights.

2.2.3 The marginalized particle filter (MPF)

The marginalized particle filter (MPF) can be applied to special state space models in which a portion of the state space is nonlinear and the other portion can be modelled as a linear process if the nonlinear portion of the state vector is known. When this partition can be performed, the linear portion of the state can be sequentially estimated using the Kalman filter, and particle filtering can be used to estimate the nonlinear portion of the state vector. The MPF reduces the variance of the posterior distribution estimation by providing optimal estimates for the linear portion of the state space while reducing the dimensionality of the nonlinear portion of the state that is estimated with a particle filter.

ISSN 2057-0058 (Online)

2.3 Optimal MAP adaptive marginalized particle filter (Optimal MAM-PF)

One can apply the Viterbi algorithm to obtain the MAP state trajectory within the canonical particle filter framework [2, 9]. However, the Viterbi algorithm does not guarantee the true MAP state trajectory when the state is marginalized [4]. Estimation of the MAP state trajectory with the marginalized state has not been previously described.

The MAM-PF is a hybrid particle filtering method which leverages the advantages of the MAP-PF and MPF algorithms. The MAP-PF portion of the algorithm requires that the likelihood function $p\left(\boldsymbol{y}_n \middle| \boldsymbol{y}_{0:n-1}, x_n^{\mathrm{N},(i)}\right)$ be evaluated for particles $x_n^{\mathrm{N},(i)}$, whose values may be far away from probable values. In this case some of the variation in $\boldsymbol{y}_n$ caused by the true underlying state would not be accounted for and the residual variance $\boldsymbol{y}_n - \tilde{\boldsymbol{y}}_{n|0:n-1}$ will be much larger than the measurement noise $R_{\boldsymbol{v},n}$ or the prediction or the prediction error $R_{\boldsymbol{e},n}$ provided by the Kalman filter recursions. This underestimation of the prediction error causes the likelihood function to have a distribution that is too narrow, which ultimately leads to suboptimal particle selection. This is a critical problem that has not been addressed previously and only occurs when attempting to use both marginalization and MAP estimation with the Viterbi algorithm.

One elegant solution to this problem is to continuously estimate the prediction error covariance from the residuals for each particle. We have adopted and modified the adaptive covariance estimation method proposed in [23]. To ensure that the estimated covariance matrix is positive semi-definite, we propose using an eigenvalue decomposition of the covariance matrix and eliminating all non-positive eigenvalues. We denote this operation as $[\boldsymbol{R}]_+$.

A second critical issue that occurs in merging marginalization and MAP estimation is that the likelihood function $p\left(\boldsymbol{y}_n \middle| \boldsymbol{y}_{0:n-1}, x_n^{\mathrm{N},(i)}\right)$ must be handled carefully because it is only conditioned on the nonlinear portion of the state vector. This distribution can be obtained from the Kalman filter recursions, as was done for the MPF

$$p\left(\boldsymbol{y}_n \middle| \boldsymbol{y}_{0:n-1}, x_n^{\mathrm{N},(i)}\right) = p\left(\boldsymbol{y}_n \middle| \hat{\boldsymbol{x}}_{n|0:n-1}^{\mathrm{L},(i)}, \boldsymbol{x}_n^{\mathrm{N},(i)}\right) \sim \mathcal{N}\left(\hat{\boldsymbol{y}}_{n|0:n-1}^{(i)}, \boldsymbol{R}_{e,n}^{(i)}\right) \tag{7}$$

However, during the maximization over all past trajectories, it is crucial to recognize that this includes the linear portion of the state space, $\hat{\boldsymbol{x}}_{n|0:n-1}^{\mathrm{L},(i)}$, unlike the MAP-PF.

The following algorithm gives a complete account of the Optimal MAMPF recursions. The covariance coefficient β is a user-specified parameter that controls the memory of the recursion for first order recursive estimation of the adaptive signal prediction error covariance, $\boldsymbol{R}_{\boldsymbol{e},n}$. The algorithm is as follows:

ISSN 2057-0058 (Online)

Initialization

for $i = 1, \dots, N_p$ do

Sample $\boldsymbol{x}_0^{\mathrm{N},(i)} \sim \pi_0\left(\boldsymbol{x}_0^{\mathrm{N}}\right)$ & $\boldsymbol{x}_{0:-1}^{\mathrm{L},(i)} = \mathrm{E}\left[\boldsymbol{x}_0^{\mathrm{L},(i)} | \boldsymbol{x}_0^{\mathrm{N},(i)}\right]$

end for

for $i = 1, \dots, N_p$ do

$\alpha_0^{(i)} = \pi_0\left(\boldsymbol{x}_0^{\mathrm{N},(i)}\right) p\left(\boldsymbol{y}_0 \middle| \boldsymbol{x}_0^{\mathrm{N},(i)}, \boldsymbol{x}_0^{\mathrm{L},(i)}\right)$ & $\boldsymbol{z}_0^{(i)} = \boldsymbol{x}_0^{(i)}$

end for

$i^* = \mathrm{argmax}_i \alpha_0^{(i)}$ & $\widehat{\boldsymbol{x}}_0 = \boldsymbol{x}_0^{i^*}$

for $i = 1, \dots, N_{\mathrm{T}}$ do

for $i = 1, \dots, N_{\mathrm{p}}$ do

Particle Propagation

$\boldsymbol{x}_n^{\mathrm{N},(i)} \sim q_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\mathrm{N},(i)}, \boldsymbol{y}_n\right)$

Marginalized Sequential Estimation

for $i = 1, \dots, N_{\mathrm{p}}$ do

$\widehat{\boldsymbol{y}}_{n|0:n-1} = H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) \widehat{\boldsymbol{x}}_{n|0:n-1}^{\mathrm{L},(k)}$ & $\boldsymbol{e}_n = \boldsymbol{y}_n - \widehat{\boldsymbol{y}}_{n|0:n-1}$

$\boldsymbol{R}_{v,n} = \left[\boldsymbol{e}_n \boldsymbol{e}_n^{\mathrm{T}} - H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) \boldsymbol{C}_{n|0:n-1}^{(k)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)^T\right]_+$

$\widehat{\boldsymbol{R}}_{v,n}^{(i,k)} = \beta \widehat{\boldsymbol{R}}_{v,n-1}^{(k)} + (1-\beta) \boldsymbol{R}_{v,n}$

$\boldsymbol{R}_{e,n}^{(k)} = H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) \boldsymbol{C}_{n|0:n-1}^{(k)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) + \widehat{\boldsymbol{R}}_{v,n}^{(k)}$

$\boldsymbol{K}_n^{(k)} = \boldsymbol{C}_{n|0:n-1}^{(k)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) \left(\boldsymbol{R}_{e,n}^{(k)}\right)^{-1}$ & $\widehat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(i)} = \widehat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(i)} + \boldsymbol{K}_n^{(k)} \left[\boldsymbol{y}_n - \widehat{\boldsymbol{y}}_{n|0:n-1}^{(k)}\right]$

$\boldsymbol{C}_{n|0:n}^{(k)} = \left[I - \boldsymbol{K}_n^{(k)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)\right] \boldsymbol{C}_{n|0:n-1}^{(k)}$

$\widehat{\boldsymbol{x}}_{n+1|0:n}^{\mathrm{L},(i,k)} = F_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) \widehat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(k)}$ & $\boldsymbol{C}_{n+1|0:n}^{(i,k)} = F_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) \boldsymbol{C}_{n|0:n}^{(k)} F_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right) + \boldsymbol{Q}_u^{\mathrm{L}}$

end for

MAP Estimation

$k^* = \mathrm{argmax}_k \alpha_{n-1}^{(k)} p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{\mathrm{N},(i)}, \widehat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(k)}\right) p\left(\boldsymbol{x}_n^{\mathrm{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\mathrm{N},(k)}\right)$

$\alpha_n^{(i)} = \alpha_{n-1}^{(k^*)} p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{\mathrm{N},(i)}, \widehat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(k^*)}\right) p\left(\boldsymbol{x}_n^{\mathrm{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\mathrm{N},(k^*)}\right)$

$\boldsymbol{C}_{n+1|0:n}^{(i)} = \boldsymbol{C}_{n+1|0:n}^{(i,k^*)}$ & $\widehat{\boldsymbol{R}}_{v,n}^{(i)} = \widehat{\boldsymbol{R}}_{v,n}^{(i,k^*)}$

$\widehat{\boldsymbol{x}}_{n+1|0:n}^{\mathrm{L},(i)} = \widehat{\boldsymbol{x}}_{n+1|0:n}^{\mathrm{L},(i,k^*)}$ & $\widehat{\boldsymbol{x}}_n^{(i)} = \left[\widehat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(i,k^*)}, \boldsymbol{x}_n^{\mathrm{N},(i)}\right]^{\mathrm{T}}$ &

ISSN 2057-0058 (Online)

$$\boldsymbol{z}_{0:n}^{(i)} = \left[\boldsymbol{z}_{0:n-1}^{(k^*)}, \widehat{\boldsymbol{x}}_n^{(i)}\right]$$

end for

$i^* = \text{argmax}_i \alpha_n^{(i)}$ & $\widehat{\boldsymbol{x}}_{0:n} = \boldsymbol{z}_{0:n}^{i^*}$

end for

2.4 Fast MAP adaptive marginalized particle filter (Fast MAM-PF)

A key computational disadvantage of the Optimal MAMPF algorithm is that the linear Kalman filter recursions must be applied N_p times for each particle, which results in N_p^2 Kalman filter recursions for each time update of the state estimate. This ensures that the maximization over all possible previous trajectories correctly accounts for the effect of the linear state estimates on the likelihood function, which is given by $p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{\text{N},(i)}, \widehat{\boldsymbol{x}}_{n|0:n-1}^{\text{L},(k)}\right)$ for the i^{th} particle. In most cases the likelihood function does not strongly affect the selection of the previous trajectory and this term can be eliminated from the MAP estimation step. Specifically,

$$k^* = \text{argmax}_k \alpha_{n-1}^{(k)} p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{\text{N},(i)}, \widehat{\boldsymbol{x}}_{n|0:n-1}^{\text{L},(k)}\right) p\left(\boldsymbol{x}_n^{\text{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\text{N},(k)}\right) \quad (8)$$

$$\approx \text{argmax}_k \alpha_{n-1}^{(k)} p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{\text{N},(i)}, \widehat{\boldsymbol{x}}_{n|0:n-1}^{\text{L},(i)}\right) p\left(\boldsymbol{x}_n^{\text{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\text{N},(k)}\right) \quad (9)$$

$$= \text{argmax}_k \alpha_{n-1}^{(k)} p\left(\boldsymbol{x}_n^{\text{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\text{N},(k)}\right). \quad (10)$$

This approximation sacrifices the asymptotic optimality of the Optimal MAM-PF, but substantially reduces the computational burden because the selection of the best past trajectory for a particle no longer requires the linear state estimates or Kalman filter recursions for each possible past trajectory. Rather, the best past trajectory can be determined before the Kalman filter recursions. These recursions can then be calculated once per particle instead of N_p times per particle. A complete description of the Fast MAM-PF algorithm is as follows:

Initialization

for $i = 1, \ldots, N_p$ do

Sample $\boldsymbol{x}_0^{\text{N},(i)} \sim \pi_0\left(\boldsymbol{x}_0^{\text{N}}\right)$ & $\boldsymbol{x}_{0:-1}^{\text{L},(i)} = \text{E}\left[\boldsymbol{x}_0^{\text{L},(i)} | \boldsymbol{x}_0^{\text{N},(i)}\right]$

end for

for $i = 1, \ldots, N_p$ do

$\alpha_0^{(i)} = \pi_0\left(\boldsymbol{x}_0^{\text{N},(i)}\right) p\left(\boldsymbol{y}_0 \middle| \boldsymbol{x}_0^{\text{N},(i)}, \boldsymbol{x}_0^{\text{L},(i)}\right)$ & $\boldsymbol{z}_0^{(i)} = \boldsymbol{x}_0^{(i)}$

end for

ISSN 2057-0058 (Online)

$i^* = \mathrm{argmax}_i \alpha_0^{(i)}$ & $\hat{\boldsymbol{x}}_0 = \boldsymbol{x}_0^{i^*}$

for $i = 1, \ldots, N_\mathrm{T}$ do

for $i = 1, \ldots, N_\mathrm{p}$ do

Particle Propagation

$$\boldsymbol{x}_n^{\mathrm{N},(i)} \sim q_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\mathrm{N},(i)}, \boldsymbol{y}_n\right)$$

MAP Estimation

$$\boldsymbol{k}^* = \mathrm{argmax}_k \alpha_{n-1}^{(k)} p\left(\boldsymbol{x}_n^{\mathrm{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\mathrm{N},(k)}\right)$$

Marginalized Sequential Estimation

$$\hat{\boldsymbol{y}}_{n|n-1} = H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)\hat{\boldsymbol{x}}_{n|0:n-1}^{\mathrm{L},(k^*)} \;\&\; \boldsymbol{e}_n = \boldsymbol{y}_n - \hat{\boldsymbol{y}}_{n|0:n-1}$$

$$\boldsymbol{R}_{v,n} = \left[\boldsymbol{e}_n \boldsymbol{e}_n^{\mathrm{T}} - H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)\boldsymbol{C}_{n|0:n-1}^{(k^*)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)^T\right]_+$$

$$\hat{\boldsymbol{R}}_{v,n} = \beta \hat{\boldsymbol{R}}_{v,n-1}^{(k^*)} + (1-\beta)\boldsymbol{R}_{v,n}$$

$$\boldsymbol{R}_{e,n} = H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)\boldsymbol{C}_{n|0:n-1}^{(k^*)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)^T + \hat{\boldsymbol{R}}_{v,n}$$

$$\boldsymbol{K}_n^{(k)} = \boldsymbol{C}_{n|0:n-1}^{(k^*)} H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)^T \left(\boldsymbol{R}_{e,n}^{(k)}\right)^{-1}$$

$$\hat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L},(i)} = \hat{\boldsymbol{x}}_{n|0:n-1}^{\mathrm{L},(k^*)} + \boldsymbol{K}_n \boldsymbol{e}_n$$

$$\boldsymbol{C}_{n|0:n}^{(k)} = \left[I - \boldsymbol{K}_n H_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)\right]\boldsymbol{C}_{n|0:n-1}^{(k^*)}$$

$$\hat{\boldsymbol{R}}_{v,n}^{(i)} = \hat{\boldsymbol{R}}_{v,n} \;\&\; \hat{\boldsymbol{x}}_{n+1|0:n}^{\mathrm{L},(i)} = F_n\left(\boldsymbol{x}_n^{\mathrm{N},(i)}\right)^T + \boldsymbol{Q}_u^{\mathrm{L}}$$

$$\alpha_n^{(i)} = \alpha_{n-1}^{(k^*)} p\left(\boldsymbol{y}_n \middle| \boldsymbol{x}_n^{\mathrm{N},(i)}, \hat{\boldsymbol{x}}_{n|0:n-1}^{\mathrm{L},(k^*)}\right) p\left(\boldsymbol{x}_n^{\mathrm{N},(i)} \middle| \boldsymbol{x}_{n-1}^{\mathrm{N},(k^*)}\right)$$

$$\hat{\boldsymbol{x}}_n^{(i)} = \left[\hat{\boldsymbol{x}}_{n|0:n}^{\mathrm{L}}, \boldsymbol{x}_n^{\mathrm{N},(i)}\right]^{\mathrm{T}} \;\&\; \boldsymbol{z}_{0:n}^{(i)} = \left[\boldsymbol{z}_{0:n-1}^{(k^*)}, \hat{\boldsymbol{x}}_n^{(i)}\right]$$

end for

$i^* = \mathrm{argmax}_i \alpha_n^{(i)}$ & $\hat{\boldsymbol{x}}_{0:n} = \boldsymbol{z}_{0:n}^{i^*}$

end for

2.5 State-space model for multi-harmonic tracking

2.5.1 State space model

Many natural signals contain nearly periodic rhythms with slowly varying morphologies. In order to model such type of signals, we assumed that those signals are generated by a sum of harmonically-related time-varying components with independent amplitudes and phase relationships. Then, we used the so-called rectangular model proposed in [24],

ISSN 2057-0058 (Online)

$$\boldsymbol{y}_n = h(\boldsymbol{x}_n) + \boldsymbol{v}_n = \left[\sum_{k=1}^{N_h} \boldsymbol{a}_{k,n}\cos(k\boldsymbol{\theta}_n) + \boldsymbol{b}_{k,n}\sin(k\boldsymbol{\theta}_n)\right] + \bar{\boldsymbol{y}}_n + \boldsymbol{v}_n \quad (11)$$

where N_h is the number of harmonics (assumed known), $\boldsymbol{\theta}_n$ is the instantaneous phase of the fundamental frequency, $\bar{\boldsymbol{y}}_n$ is the signal mean, $\boldsymbol{a}_{k,n}$ and $\boldsymbol{b}_{k,n}$ are the sinusoidal coefficients, and $\boldsymbol{v}_n$ is white Gaussian noise with variance $r_{\boldsymbol{v}}$. The variations in the state variables over time are modelled as follows:

$$\boldsymbol{\theta}_{n+1} = \boldsymbol{\theta}_n + 2\pi T_s \boldsymbol{f}_n, \quad (12)$$

$$\bar{\boldsymbol{f}}_{n+1} = g\left[\bar{\boldsymbol{f}}_n + \boldsymbol{u}_{\bar{\boldsymbol{f}},n}\right], \quad (13)$$

$$\boldsymbol{f}_{n+1} = \bar{\boldsymbol{f}}_n + \alpha\left(\boldsymbol{f}_n - \bar{\boldsymbol{f}}_n\right) + \boldsymbol{u}_{\boldsymbol{f},n}, \quad (14)$$

$$\boldsymbol{a}_{k,n+1} = \boldsymbol{a}_{k,n} + \boldsymbol{u}_{\boldsymbol{a},n}, \quad (15)$$

$$\boldsymbol{b}_{k,n+1} = \boldsymbol{b}_{k,n} + \boldsymbol{u}_{\boldsymbol{b},n}, \quad (16)$$

$$\bar{\boldsymbol{y}}_{n+1} = \bar{\boldsymbol{y}}_n + \boldsymbol{u}_{\bar{\boldsymbol{y}},n} \quad (17)$$

where $\boldsymbol{f}_n$ is the fundamental frequency, $\bar{\boldsymbol{f}}_n$ is the mean fundamental frequency, $T_s = \frac{1}{f_s}$ is the sampling interval, α is an autoregressive (AR) process coefficient for the instantaneous frequency model, and $\boldsymbol{u}_{\cdot,n}$ are mutually uncorrelated white noise processes. A value of $\alpha = 1$ results in a random walk model and $\alpha = 0$ results in a white noise model. The variance of $\boldsymbol{u}_n$ determines how quickly the parameters are expected to change over time.

Unlike previous frequency tracking models, in this model the mean frequency $\bar{\boldsymbol{f}}_n$ is modelled as a state variable that is distinct from the instantaneous frequency $\boldsymbol{f}_n$. In the particle filter framework this permits the model to account for many possible mean frequencies with relatively small amplitude fluctuations about their means, though at an expense of increasing the dimension of the nonlinear portion of the state space.

In many applications the range of the possible mean frequencies is known from domain knowledge. For example, in an application to track the heart rate of an adult, the range of typical adult heart rates is known. We model this by designing the state model such that $\bar{\boldsymbol{f}}_n$ has a uniform distribution $\bar{f} \sim \mathcal{U}(f_{\min}, f_{\max})$. This is achieved by selecting a uniform distribution for the random step $\boldsymbol{q}_{\bar{\boldsymbol{f}},n} \sim \mathcal{U}\left(-\delta_{\bar{f}}, \delta_{\bar{f}}\right)$ and using a nonlinear reflecting function to account for boundary effects

$$g[f] = \begin{cases} f_{\max} - (f - f_{\max}) & f_{\max} < f \\ f & f_{\min} < f \le f_{\max} \\ f_{\min} + (f_{\min} - f) & f \le f_{\min} \end{cases}. \quad (18)$$

ISSN 2057-0058 (Online)

This essentially causes the mean frequency $\bar{\boldsymbol{f}}_n$ to bounce elastically from the boundaries at $f_{\max}$ and $f_{\min}$, which in turn ensures that at any given time n the mean frequency is uniformly distributed within this range.

2.5.2 Multi-modal posterior

To demonstrate that the posterior distribution of instantaneous fundamental frequencies is multi-modal, it is useful to first consider a simplified case in which the posterior distribution can be analytically computed. In general it is difficult to solve for the posterior distribution exactly, even if a state space model of the process is known. However, if we use the simplifying assumptions that the sinusoidal coefficients and fundamental frequency are not changing over time, have uniform prior distributions, and the measurement noise is Gaussian, then we can solve for the posterior distribution with a least squares approach. In this case the observation model is

$$\boldsymbol{y}_n = \left[\sum_{k=1}^{N_h} \boldsymbol{a}_k \cos(k\boldsymbol{\omega} n) + \boldsymbol{b}_k \sin(k\boldsymbol{\omega} n)\right] + \bar{\boldsymbol{y}} + \boldsymbol{v}_n \tag{19}$$

where $\boldsymbol{\omega}$ is the fixed fundamental frequency in radians, $\bar{\boldsymbol{y}}$ is the fixed signal mean, $\boldsymbol{a}_k$ and $\boldsymbol{b}_k$ are the fixed sinusoidal coefficients. If the fundamental frequency and measurement noise are known, this model is a linear function of the signal mean and sinusoidal coefficients.

If we collect the unknown parameters into a single state vector

$$\boldsymbol{x} = [\bar{\boldsymbol{y}} \quad \boldsymbol{a}_1 \quad \ldots \quad \boldsymbol{a}_k \quad \boldsymbol{b}_1 \quad \ldots \quad \boldsymbol{b}_k \quad \boldsymbol{\omega}] \tag{20}$$

then the posterior distribution is given by

$$p(\boldsymbol{x}|\boldsymbol{y}_{0:n}; r_v) = \frac{p(\boldsymbol{y}_{0:n}|\boldsymbol{x}; r_v)p(\boldsymbol{x})}{p(\boldsymbol{y}_{0:n}; r_v)}. \tag{21}$$

Since we have assumed a uniform prior, $p(\boldsymbol{x}) = c$ for some constant c, the posterior distribution is proportional to the likelihood function. For a specified value of the fundamental frequency, the remaining parameters can be estimated by linear least squares, or, equivalently, maximum likelihood. An unbiased estimate of the measurement noise variance r_v is then given by

$$\hat{r}_v = \frac{1}{N_{\mathrm{T}} - 1 - 2N_h} \sum_{n=1}^{N_{\mathrm{T}}} (\boldsymbol{y}_n - \hat{\boldsymbol{y}}_n)^2 \tag{22}$$

and the marginal posterior is then

ISSN 2057-0058 (Online)

$$p(\boldsymbol{\omega}|\boldsymbol{y}_{0:n}) = \max_{\bar{y},\boldsymbol{a}_k,\boldsymbol{b}_k} p(\boldsymbol{y}_{0:n}|\boldsymbol{x};\hat{r}_v). \quad (23)$$

This estimated posterior distribution evaluated as a function of frequency for a signal with $N_{\mathrm{T}} = 500$ samples, $N_{\mathrm{h}} = 5$ harmonics, and a signal-to-noise ratio (SNR) of 3 dB is shown in Figure 1. This demonstrates that the marginal posterior distribution contains many modes, which limits the accuracy of state space methods such as the extended Kalman filter (EKF) and unscented Kalman filter (UKF) that characterize the posterior by the mean and covariance alone.

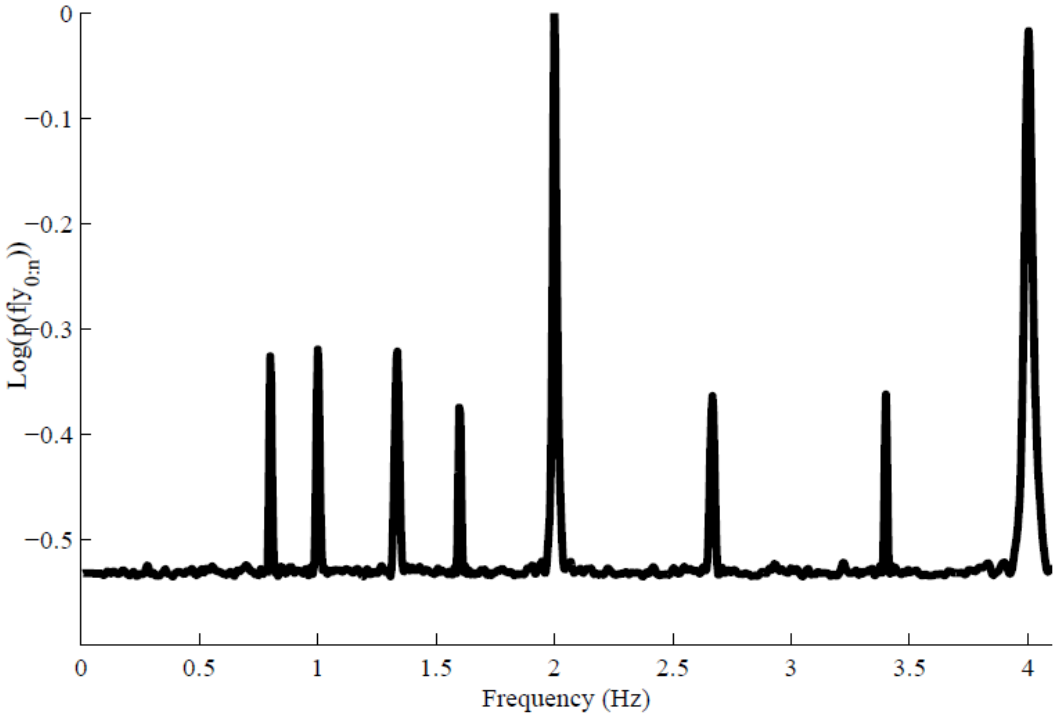

Figure 1 Logarithmic posterior PDF of fundamental frequencies for a multi-harmonic signal with 10 harmonics.

2.6 Performance measure

There are several performance measures for the harmonic tracking problem such as normalized mean-squared-error (NMSE) of the predicted signal, frequency mean-squared-error (FMSE), and mean-time-to-lock (MTL) for the frequency estimate [8]. Among them, we chose two measures: the NMSE of the predicted signal and the FMSE of the fundamental frequency. FMSE measures the accuracy of the fundamental frequency estimation alone while NMSE represents how accurately the state estimates describe the original signal. For real signals only NMSE can be computed since other measures require knowledge of the true state, which is not known.

ISSN 2057-0058 (Online)

FMSE is an important measure when the main goal of tracking is accurate estimation of the fundamental frequency. It can be written as follows,

$$\mathrm{FMSE} = \frac{\sum_{n=1}^{N_\mathrm{T}} \left(\boldsymbol{f}_n - \hat{\boldsymbol{f}}_n\right)^2}{N_\mathrm{T}} \tag{24}$$

whose unit is Hz^2. We will report $\sqrt{\mathrm{FMSE}}$ in Hz.

NMSE of signal estimation ranges from 0 to ∞. When its value is below 1, the tracker does a better job than a simple mean estimator. If its value is greater than 1, the tracker performs worse than estimating the signal to be equal to the signal mean. The normalized mean-square-error (NMSE) can be calculated as follows,

$$\mathrm{NMSE} = \frac{\sum_{n=1}^{N_\mathrm{T}} (\boldsymbol{y}_n - \hat{\boldsymbol{y}}_n)^2}{\sum_{n=1}^{N_\mathrm{T}} (\boldsymbol{y}_n - \bar{\boldsymbol{y}}_n)^2}. \tag{25}$$

3 RESULTS

In subsequent sections we apply three particle filters (MPF, Optimal MAM-PF, and Fast MAM-PF) to synthetic and real signals, where the fundamental frequency's marginal posterior distribution is non-Gaussian and multi-modal.

3.1 Multi-harmonic tracker comparison

Given the same number of particles, each of the trackers requires different computational load and yields different performance results. In order to compare the computational load and performance of each tracker, we generated synthetic signals mimicking realistic multi-harmonic signals. The sampling frequency was 40 Hz and the signal duration was 5 min. The mean of the fundamental frequency ($\bar{\boldsymbol{f}}$) was 1.5 Hz. The synthetic signal was generated with Formula 11 with a signal mean of $\bar{\boldsymbol{y}}_n = 0$ and constant coefficients $\boldsymbol{a}_k$ and $\boldsymbol{b}_k$. An example of a synthetic signal's spectrogram is illustrated in Figure 2. The higher harmonic components of the signal have more power than the lower harmonic components, which is common in real multi-harmonic signals such as electrocardiograms. The parameters used to generate the synthetic signals are listed in Table 1.

ISSN 2057-0058 (Online)

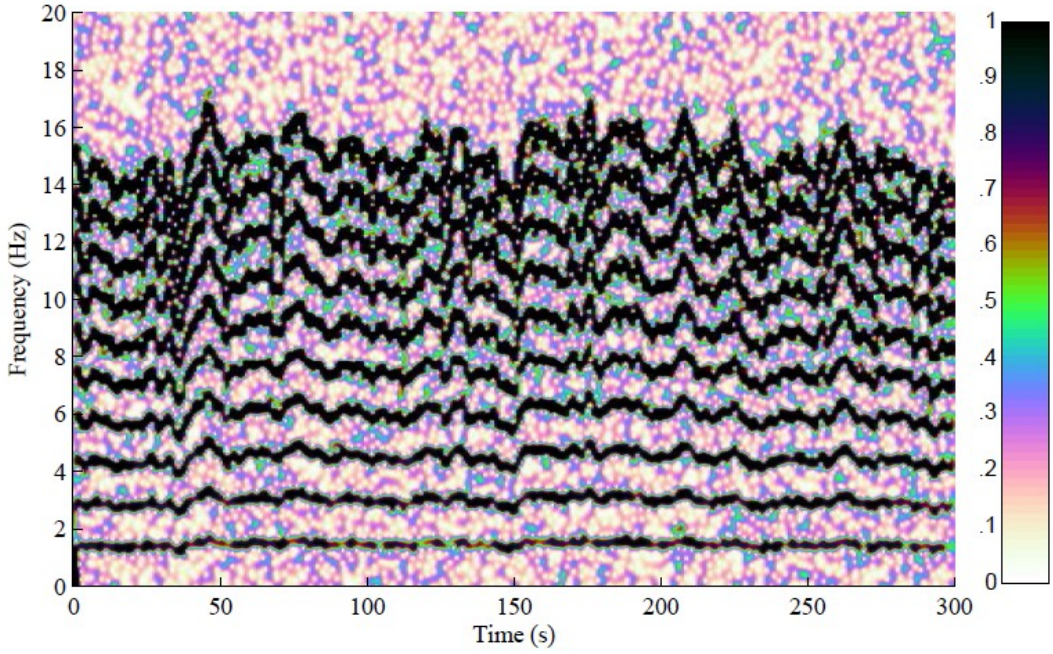

Figure 2 Spectrogram of a synthetic multi-harmonic signal whose fundamental frequency is centred at 1.5 Hz**.**

Table 1. Realistic synthetic signal generation parameters.

Name	**Symbol**	**Value**
Number of particles	N_{p}	100
Number of harmonics	N_{h}	10
Sampling frequency	f_{s}	40 Hz
Signal duration	l	5 min
Number of samples	N_{T}	12,000
Signal-to-noise ratio	SNR	10 dB
Fundamental mean frequency	$\bar{\omega}$	3π
Frequency coefficient	α	0.99

The two plots in Figure 3 depict the $\sqrt{\mathrm{FMSE}}$ and NMSE versus relative simulation times of the three multi-harmonic trackers. The Optimal MAM-PF substantially outperformed the MPF in terms of both $\sqrt{\mathrm{FMSE}}$ and NMSE. The main reason why the MPF has a larger $\sqrt{\mathrm{FMSE}}$ than the Optimal MAM-PF is that the MPF loses track of the true fundamental frequency occasionally and tracks one of the subharmonics, which are local maxima in the posterior distribution. The Optimal MAM-PF has less chance to erroneously track the subharmonics than the MPF since the particles of the Optimal MAM-PF are distributed throughout the entire possible fundamental frequency range of the signal and do not suffer from sample degeneracy. The erroneous frequency

ISSN 2057-0058 (Online)

tracking of the MPF causes it to have a very large NMSE, which is slightly greater than 1.

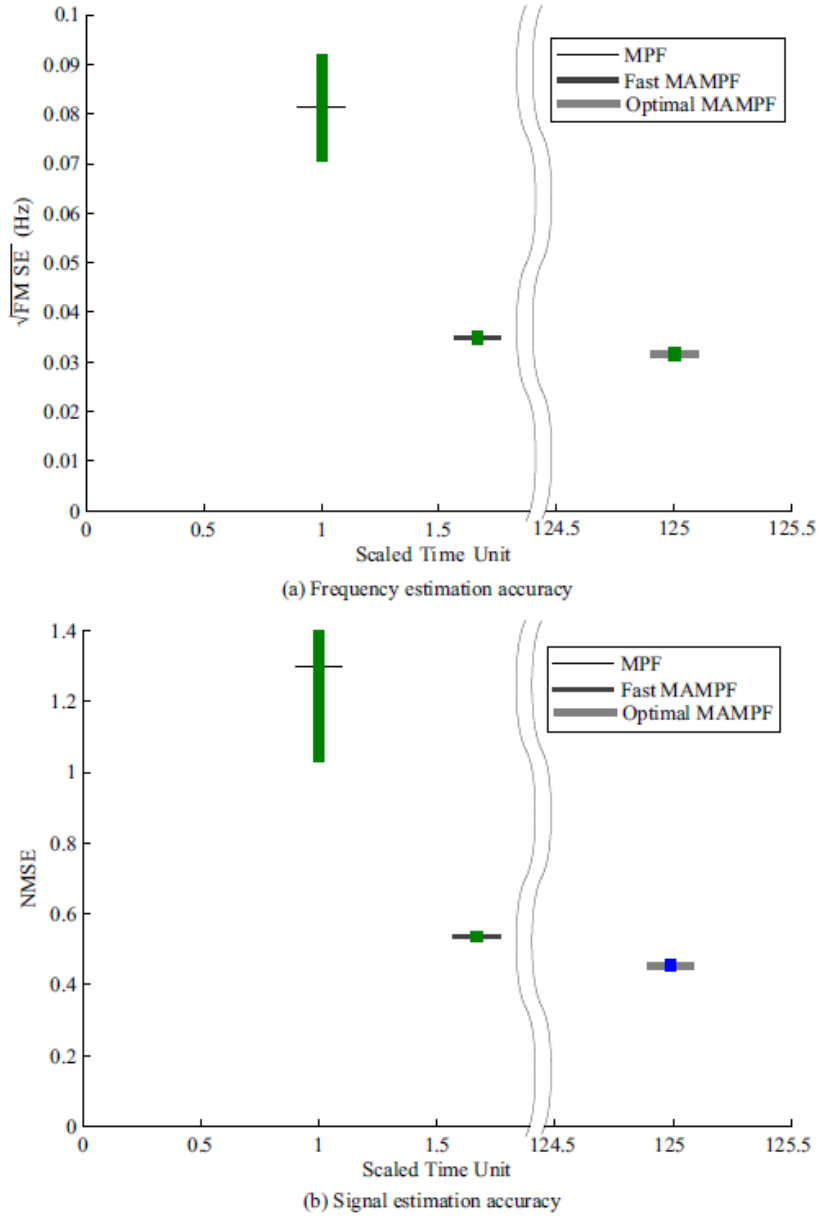

Figure 3 (a) $\sqrt{\mathrm{FMSE}}$ versus simulation time. (b) NMSE versus simulation time. The horizontal lines represent the mean values while the vertical bars the one-standard-deviation ranges around the means.

As shown in Figure 3, the simulation time required to obtain the Optimal MAM-PF state estimates was approximately 125 times greater than the MPF. However, the performance of the Fast MAM-PF was comparable to that of the Optimal MAM-PF, while the computation duration was only 1:6 times greater than the MPF's. This result is in line with the fact that the computational

ISSN 2057-0058 (Online)

burdens of the MPF and MAM-PF are proportional to $\mathcal{O}(N_\mathrm{p})$ and $\mathcal{O}(N_\mathrm{p}^2)$, respectively, where the total number of particles N_p was 100 in this simulation study. Since the Fast MAM-PF is nearly as accurate as the Optimal MAM-PF, but requires substantially less computation, we limit the remaining performance analysis to the MAM-PF and Fast MAM-PF trackers.

3.2 Equalizing computational load

The Fast MAM-PF requires more computation than the MPF due to the Viterbi search for the most probable past path of each particle. In order to make a fair comparison between these two filters, we first ran the Fast MAM-PF with 150 particles and measured the simulation duration. Then, we ran the MPF with various numbers of particles to find the number of particles with which the MPF uses the same simulation duration. With 180 particles the simulation time of the MPF was equal to that of the Fast MAM-PF. This process approximately equalizes the processing time required for each filter for the frequency tracking applications discussed in the following sections.

3.3 Adaptation ability of MPF and Fast MAM-PF

The spectrogram of a synthetic multiharmonic signal with an abrupt jump in the fundamental frequency from 2 Hz to 4 Hz after 60 s is shown in Figure 4(a). We applied both the Fast MAM-PF and the MPF harmonic trackers to this signal to compare their tracking ability. We hypothesized that the Fast MAM-PF harmonic tracker would be able to detect abrupt changes to the posterior distribution of fundamental frequencies and regain track, while the MPF would fail to regain track. Table 2 summarizes the further details of the synthetic signal generator. The user-specified parameters of the filters are listed in Table 3.

Table 2. Frequency shifting synthetic signal generation parameters.

Name	Symbol	Value
Number of harmonics	N_h	5
Sampling frequency	f_s	50 Hz
Signal duration	l	5 min
Number of samples	N_T	1.5e4
Signal-to-noise ratio	SNR	10 dB
Fundamental frequency	ω	4π & 8π

The estimated fundamental frequencies for the trackers in dark and light grey lines are shown in Figure 4(b). The advantage of the Fast MAM-PF approach is demonstrated by the response to the abrupt shift in frequency. Although it takes time for the Fast MAM-PF tracker to regain track after the sudden shift, it eventually locks on to the true frequency. It is able to make this transition, even though it is inconsistent with the underlying statistical model, because shortly after the frequency jump those particles at high frequencies

ISSN 2057-0058 (Online)

become more probable than the particles at low frequencies. The MPF, however, continues to estimate the fundamental frequency around 2 Hz because all of the particles are clustered about this frequency. Thus, the poor performance of the MPF is due to sample impoverishment.

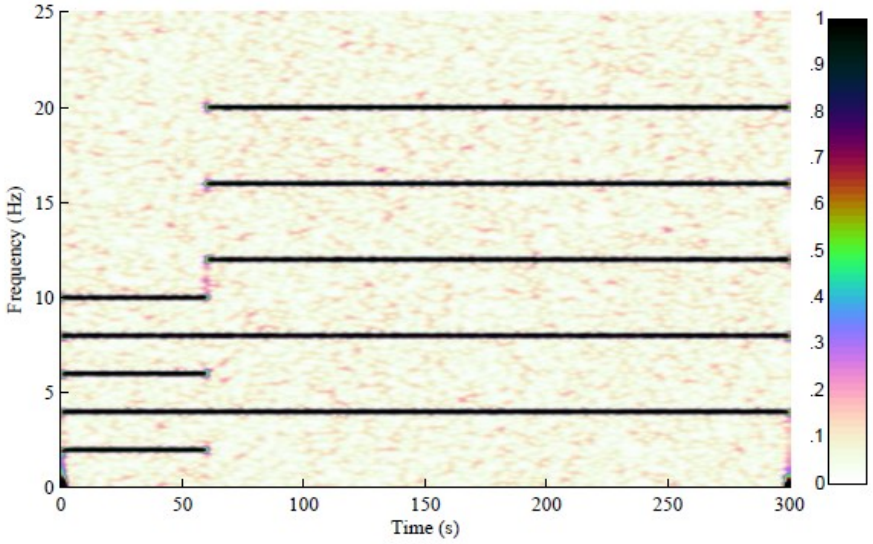

(a) Spectrogram

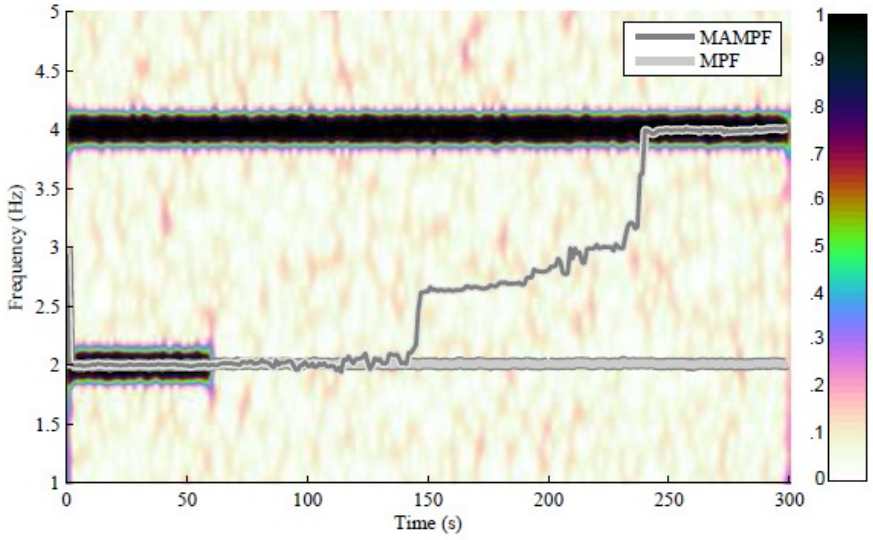

(b) Estimated fundamental frequencies

Figure 4 (a) Spectrogram of a synthetic multi-harmonic signal. (b) Estimated fundamental frequencies (Fast MAM-PF: black line, MPF: grey line).

ISSN 2057-0058 (Online)

Table 3. User-specified parameters for the filters (Fast MAM-PF/MPF).

Name	Symbol	Value
Number of particles	N_{p}	150/180
Resampling threshold	N_{t}	N_{p}/100
Number of harmonics	N_{h}	5
Frequency coefficient	α	0.999
Covariance coefficient	β	0.98
Frequency slew rate	$\delta_{\bar{f}}$	$1/f_s$Hz
Max. mean frequency range	$f_{\max}$	5 Hz
Min. mean frequency range	$f_{\min}$	1 Hz
Phase initial	θ_0	0
Frequency initial	f_0	2 Hz
Frequency mean initial	$\bar{f}_0$	2 Hz
Amplitudes initial	a_0 & b_0	0.01
Signal mean initial	$\bar{y}_0$	0
Frequency variance	R_{q_f}	1e-4
Mean Frequency variance	$R_{q_{\bar{f}}}$	1e-6
Amplitudes variance	R_{q_a} & R_{q_b}	var(y)/10
Signal mean variance	$R_{q_{\bar{f}}}$	1e-2
Measurement variance	R_v	var(y)/10

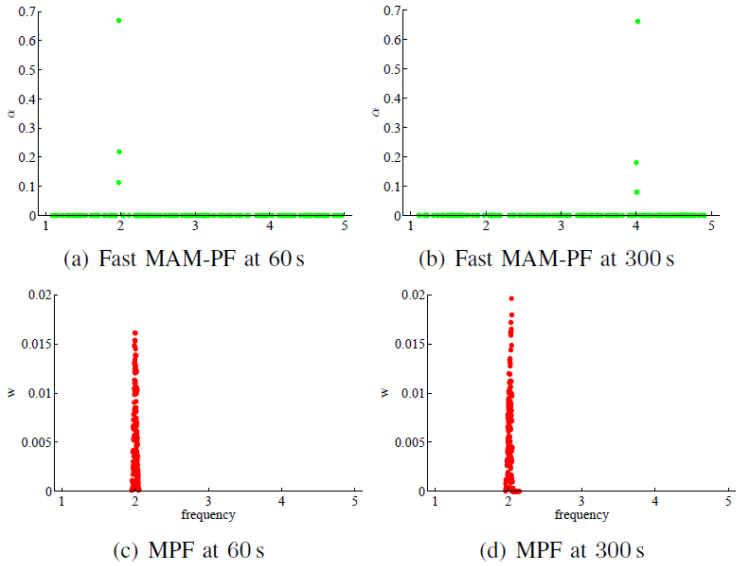

Figure 5 Distributions of particles (α) of the Fast MAM-PF and those (w) of the MPF. (a) α versus frequency at 60 s. (b) α versus frequency at 300 s. (c) w versus frequency at 60 s. (d) w versus frequency at 300 s.

ISSN 2057-0058 (Online)

It is illustrated in Figure 5 how the particles of the Fast MAM-PF and MPF are distributed at 60 s and 300 s. The true fundamental frequency value was 2 Hz up to 60 s and 4 Hz thereafter. Plot (d) in Figure 5 shows that the particles of the MPF were clustered about 2 Hz although the true fundamental frequency value was 4 Hz. This is a good example of the particles of the MPF locking onto one of the subharmonics of the fundamental frequency.

3.4 Real signal example

We applied both harmonic trackers to a real electrocardiogram (ECG) signal sampled at 500 Hz containing a high level of noise. The goal of this experiment is to compare the ability of the trackers to retain track on a noisy signal amid realistic signal artifacts such as signal dropouts and medical interventions. The ECG signal chosen here is one of the noisiest signals in the MIMIC database [22] on PhysioNet [10]. The estimated frequencies are the results of averaging the results of 100 simulations.

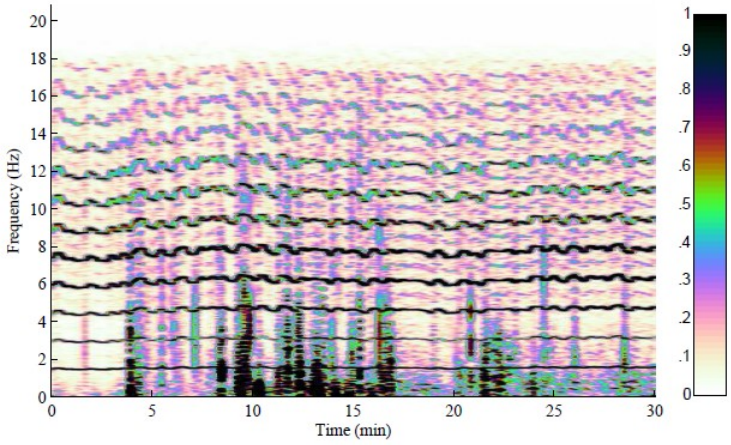

(a) Spectrogram

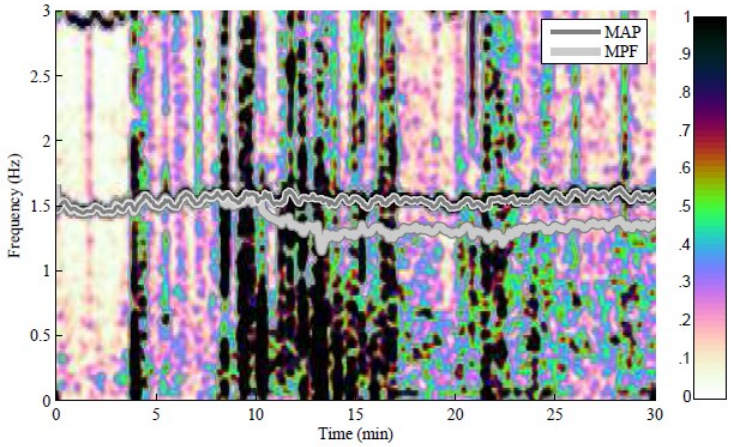

(b) Estimated fundamental frequencies

Figure 6 (a) Spectrogram of a real ECG signal. (b) Estimated fundamental frequencies (MAM-PF: dark grey line, MPF: light grey line).

ISSN 2057-0058 (Online)

The spectrogram of a real ECG signal (a) and the estimated fundamental frequencies using the two trackers (b) are depicted in Figure 6. The Fast MAM-PF does not lose track of the fundamental frequency during the entire duration of the signal, including the noisy segment from 9 to 16 min. In contrast, on average the MPF loses track of the fundamental frequency starting from around 9 min and never locks on to the true fundamental frequency again.

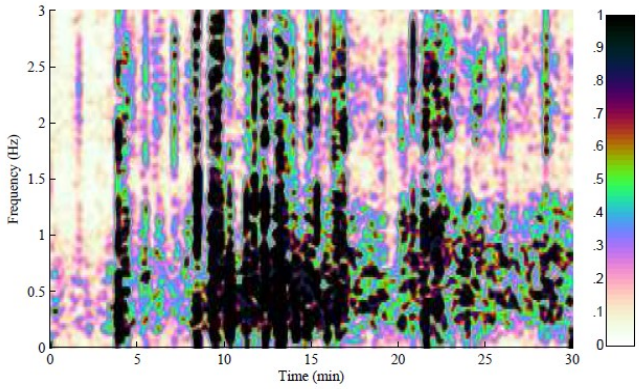

(a) MAM-PF: Averaged NMSE of 0.76

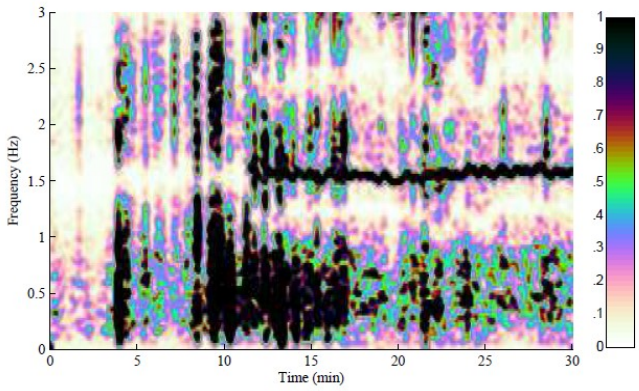

(b) MPF: Averaged NMSE of 0.85

Figure 7 (a) Spectrogram of residuals after applying the MAM-PF. (b) Spectrogram of residuals after applying the MPF.

The spectrograms of the signal estimation errors of the Fast MAM-PF (a) and MPF (b) trackers are illustrated in Figure 7. The Fast MAM-PF removes the fundamental frequency component and its harmonic components successfully. The MPF, however, fails to remove the harmonic components in the ECG signal approximately from 11 min. The residual spectrogram of the MPF shows very low power (white area) around 1.2 Hz after 11 min because the MPF removed one of the subharmonics of the fundamental frequency, which was at approximately 1.5 Hz.

4 DISCUSSION AND SUMMARY

We introduced two new algorithms that integrate marginalization and MAP estimation for Sequential Monte Carlo methods. Both algorithms

ISSN 2057-0058 (Online)

overcome two critical issues that arise when marginalization and MAP estimation are integrated, and that do not occur when these methods are used separately. We used a multi-harmonic frequency tracking problem as an example application to demonstrate the benefits of the new algorithm (MAM-PF) over a traditional marginalized particle filter (MPF). The new algorithms can find the correct mode more often and regain track after sudden shifts in the state more quickly than the MPF. They are also more immune to signal drop-outs and severe noise than the MPF.

The MAM-PF tracker can be applied to various quasiperiodic biomedical signals such as ECG, ABP, MER, and EEG signals. For example, the proposed algorithm (Fast MAMPF) has led to the development of a novel pulse pressure variation tracking method for ABP signals [15, 16].

In the previous work, we conducted a thorough simulation study to compare the performance of two tracking algorithms, which were based on the extended and sigma-point Kalman filters [19]. From the simulation study we learned that losing the track of the instantaneous frequency of a given signal is the main factor causing the performance difference between two tracking algorithms. That is the reason why the present work is focused on demonstrating the superior tracking ability of the MAM-PFs in extreme cases such as a sudden frequency shift and low signal-to-noise ratio. A thorough simulation study may be of value to compare the tracking performance of the MAM-PFs and MPF. However, the results from such study will be in line with what is shown in Figure 3 and Figure 6.

COMPETING INTERESTS STATEMENT

No competing interests.

REFERENCES

[1] Cadini F, Zio E, Peloni G (2012) Particle filtering for the detection of fault onset time in hybrid dynamics systems with autonomous transitions. IEEE Transactions on Reliability 61(1): 130-139

[2] Cappe O, Godsill S, Moulines E (2007) An overview of existing methods and recent advances in sequential Monte Carlo. Proceedings of the IEEE 95(5): 899-924

[3] Casella G, Robert CP (1996) Rao-blackwellisation of sampling schemes. Biometrika 83(1): 81-94

[4] Cemgil AT, Kappen B (2003) Monte Carlo methos for tempo tracking and rhythm quantization. Journal of Artificial Intelligence Research 18: 45-81

[5] Doucet A, Andrieu C (2001) Iterative algorithms for state estimation of jump markov linear systems. IEEE Transactions on Signal Processing 49(6): 1216-1227

[6] Doucet A, Godsill S, Andrieu C (2000) On sequential Monte Carlo sampling methods for Bayesian filtering. Statistics and Computing 10(3):197-208

ISSN 2057-0058 (Online)

[7] Dubois C, Davy M (2005) Harmonic tracking using sequential monte carlo. In: IEEE Workshop on Statistical Signal Processing, pp. 1292-1297
[8] Fischler E, Bobrovsky BZ (2006) Mean time to loose lock of phase tracking by particle filtering. Signal Processing 86: 3481-3485
[9] Godsill S, Doucet A, West M (2001) Maximum a posteriori sequence estimation using Monte Carlo particle filters. Annals of the Institute of Statistical Mathematics 53(1): 82-96
[10] Goldberger AL, Amaral AN, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng C-K, Stanley HE (2000) PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals 101(23): e215-e220
[11] Gordon N, Salmond D, Smith A (1993) Novel approach to nonlinear/non-Gaussian Bayesian state estimation. IEE Proceedings-F 140(2): 107-113
[12] Hutter F, Dearden R (2003) Te gaussian particle filter for diagnosis of non-linear systems. Proceedings of the 14th International Conference on Principles of Diagnosis, pp. 65-70
[13] Kitagawa G (1996) Monte-carlo filter and smoother for non-gaussian non-linear state sapace models. Journal of Computational and Graphical Statistics 1: 1-25
[14] Klaas M, de Freitas N, Doucet A (2005) Toward practical n2 monte carlo: The marginal particle filter. In: Proceedings of the Twenty-First Conference Annual Conference on Uncertainty in Artiticial Intelligence, pp. 308-315
[15] Kim S, Aboy M, McNames J (2009) Pulse pressure variation estimation using a sequential Monte Carlo method. In: Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 5713-5716
[16] Kim S, Aboy M, McNames J (2013) Pulse pressure variation tracking using sequential Monte Carlo methods. Biomedical Signal Processing and Control 8(4): 333-340
[17] Kim S, Homstrom L, McNames J (2008) Multiharmonic tracking using marginalized particle filters. In: Proceedings of the 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 29-33
[18] Kim S, McNames J (2010) Detecting and tracking tremor in spike trains using the rectangular model based extended Kalman smoother. Journal of Neuroscience Methods 188(1):97-104
[19] Kim S, Paul AS, Wan EA, McNames J (2010) Multiharmonic frequency tracking method using the signal-point Kalman smoother. EURASIP Journal on Advances in Signal Processing 2010: article no. 36
[20] Liu JS, Chen R (1998) Sequential monte carlo methods for dynamic systems. Journal of the American Statistical Association 93(443): 1032-1033

ISSN 2057-0058 (Online)

[21] Metropolis N, Rosenbluth A, Rosenbluth M (1953) Equation of state calculations by fast computing machines. Journal of Chemical Physics 21: 1087
[22] Moody GB, Mark RG (1996) A database to support development and evaluation of intelligent intensive care monitoring. Computers in Cardiology 23: 657-660
[23] Myers KA, Tapley BD (1976) Adaptive sequential estimation with unknown noise statistics. IEEE Transactions on Automatic Control 21(4): 520-523
[24] Parker PJ, Anderson BD (1990) Frequency tracking of nonsinusoidal periodic signals in noise. Signal Processing 20(2): 127-152
[25] Shivappa S, Rao B, Trivedi M (2010) Audio-visual fusion and tracking with multilevel iterative decoding: Framework and experimental evaluation. IEEE Journal of Selected Topics in Signal Processing 4(5): 882-894
[26] Sivia DS (1996) Data Analysis: A Bayesian Tutorial. New York, USA: Oxford University Press
[27] Viterbi A (1967) Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. IEEE Transactions on Information Theory 13(2): 260-269
[28] Wolfel M (2009) Enhanced speech features by single-channel joint compensation of noise and reverberation. IEEE Transactions on Audio, Speech, and Language Processing 17(2): 312-323
[29] Yee D, Reilly J, Kirubarajan T (2007) A blind sequential monte carlo detector for ofdm systems in the presence of phase noise, multipath fading, and channel order uncertainty. IEEE Transactions on Signal Processing 55(9): 4581-4598
[30] Zhou H, Taj M, Cavallaro A (2008) Target detection and tracking with heterogeneous sensors. IEEE Journal of Selected Topics in Signal Processing 2(4): 503-513

ISSN 2057-0058 (Online)