

A fast and unbiased minimalistic resampling approach for the particle filter

Ramakrishna Gurajala

Dept. of E.C.E., Andhra University
Visakhapatnam, India
ramagurajala@gmail.com

Praveen B. Choppala

Dept. of E.C.E., WISTM, Andhra Univ.
Visakhapatnam, India
praveen.b.choppala@gmail.com

James Stephen Meka

Dept. of C.S.E., WISTM, Andhra Univ.
Visakhapatnam, India
jamesstephenm@gmail.com

Paul D. Teal

School of E.C.S., Victoria Univ. of Wellington
Wellington, New Zealand
paul.teal@vuw.ac.nz

Abstract—The particle filter is an important approximation method for online state estimation in nonlinear nonGaussian scenarios. The resampling step in the particle filter is critical because it eliminates the wasteful use of particles that do not contribute to the posterior (degeneracy). The fully stochastic resamplers, despite being unbiased in approximating the posterior density, involve exhaustive and sequential communication within the particles and thus are computationally expensive. The alternate partial deterministic resamplers overcome this problem by reducing the communication within particles but this leads to approximation bias. This paper proposes a fast resampling procedure that gives an accurate approximation of the posterior and tracks as accurately as the conventional resamplers.

Index Terms—Particle filter, resampling, bias analysis, channel estimation, bearings-only tracking.

I. INTRODUCTION

THE Bayesian state estimation is an important solution to estimate hidden dynamic target states and is used in a wide range of applications including communications and tracking [1]. The particle filter (PF) [2], [3] is a popular Bayesian estimation method that provides a framework for inferring the hidden target state in nonlinear and nonGaussian state space models. The PF implements the Bayesian estimation by sequentially generating and updating a set of particles and their corresponding weights which together represent the state posterior probability density function (pdf). This operation is called sequential importance sampling (SIS). SIS by itself, leads to degeneracy, a situation where few particles take the full weight and the others become nonrepresentative of the posterior. This problem is overcome using resampling that duplicates the larger weight particles in place of the lower weight ones such that the PF represents the posterior more accurately [4], [5].

The stochastic resamplers are by far the most popularly used resamplers. These operate by first evaluating the cumulative sum of the normalised particle weights and then finding a value of the sum greater than a random sample drawn from $\mathcal{U}(0, 1)$, one sample per particle. The popular of these methods are the multinomial [2], stratified [6] and systematic [7] resamplers where the latter improve over the former by reducing the

associated Monte Carlo error variance. The second class are the partially deterministic resamplers. The residual resampler [8] resamples the low weight particles using the principle of proportional allocation of the high weight ones. The work in [9] used thresholding schemes to deterministically replicate large weight particles and stochastically sample the lower weight ones. A similar approach for treating wasteful duplication of large weight particles was proposed in [10], [11]. The soft systematic resampler [12] improves over the residual resampler and treats the tails of the posterior more accurately. Another class of resamplers developed to suit for parallel implementation include the recently proposed Metropolis resampler [13] and others [14], [15]. These aid in nearly-parallel PF implementation by reducing the communication within the particles; but this comes at the expense of reduced filtering accuracy and bias.

Contribution of this paper: The stochastic resamplers, especially the systematic and residual resamplers, are by far the most extensively used PF resampling methods. Leveraging the resampling on random samples, one sample per particle, sequentially for all the particles, allows for exhaustive communication within all the particles in a way that the total information (weight) is preserved; this results in the much desired unbiased representation of the posterior. However this exhaustive and sequential resampling operation leads to high Monte Carlo error variance [4] and computational complexity [13]. The partial deterministic resamplers overcome this computational complexity but a comprehensive study of their unbiasedness is not available. This paper proposes a fast deterministic resampling approach with minimal communication within the particles and practically satisfies the unbiasedness property of the PF for good resampling quality.

The rest of the paper is organised as follows. Section II sets the notation and presents the PF. Section III presents the proposed deterministic resampling method. This is followed by a analysis on the bias properties of the proposed method in section IV, evaluation results in section V and conclusion in section VI.

II. PARTICLE FILTERING

Consider a state space model defined by a Markovian state transition and observation models as

$$\mathbf{x}_t | \mathbf{x}_{t-1} \sim p(\mathbf{x}_t | \mathbf{x}_{t-1}), \mathbf{y}_t | \mathbf{x}_t \sim p(\mathbf{y}_t | \mathbf{x}_t) \quad (1)$$

for $t = 1, \dots, T$, where (a) the target state $\mathbf{x}_t \in \mathbb{R}^{d_x}$ at time instant $t \in \mathbb{N}$ is a hidden Markov process with initial distribution $p(\mathbf{x}_0)$ and the Markov state transition pdf $p(\mathbf{x}_t | \mathbf{x}_{t-1})$ with d_x denoting the state dimensionality, (b) the sensor observation $\mathbf{y}_t \in \mathbb{R}^{d_y}$ is conditionally independent given the state and has the observation density is $p(\mathbf{y}_t | \mathbf{x}_t)$ with d_y denoting the observation dimensionality. The model includes a vector of known parameters which is omitted here for convenience.

The aim of Bayesian state filtering is to recursively estimate the posterior pdf of the hidden target state using all available observations. If the posterior pdf $p(\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1})$ at time $t-1$ is available, then the posterior pdf at time t is computed as

$$p(\mathbf{x}_t | \mathbf{y}_{1:t}) = \frac{p(\mathbf{y}_t | \mathbf{x}_t) p(\mathbf{x}_t | \mathbf{y}_{1:t-1})}{p(\mathbf{y}_t | \mathbf{y}_{1:t-1})}$$

where the *a priori* state prediction pdf is $p(\mathbf{x}_t | \mathbf{y}_{1:t-1}) = \int p(\mathbf{x}_t | \mathbf{x}_{t-1}) p(\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1}) d\mathbf{x}_{t-1}$ and $p(\mathbf{y}_t | \mathbf{y}_{1:t-1})$ is the normalisation constant. The PF approximates (or represents) the pdfs in the Bayesian recursion using a set of weighted particles. Consider a set of N weighted particles $\{\mathbf{x}_{t-1}^i, w_{t-1}^i\}_{i=1}^N$ representing the posterior pdf of the target state $p(\mathbf{x}_{t-1} | \mathbf{y}_{1:t-1})$ at time $t-1$. In the SIS step, the PF generates a new set of particles from the old ones using a proposal distribution $q(\cdot)$ and weighs them according to

$$\bar{\mathbf{x}}_t^i \sim q(\mathbf{x}_t | \mathbf{x}_{t-1}^i, \mathbf{y}_t) \quad (2)$$

$$\bar{w}_t^i = w_{t-1}^i \frac{p(\mathbf{y}_t | \bar{\mathbf{x}}_t^i) p(\bar{\mathbf{x}}_t^i | \mathbf{x}_{t-1}^i)}{q(\bar{\mathbf{x}}_t^i | \mathbf{x}_{t-1}^i, \mathbf{y}_t)}, i = 1, \dots, N \quad (3)$$

The weights are then normalised and this weighted particle set is representative of the posterior at t . However, after a few iterations, the discrepancy between the weights increases, leading to degeneracy. The solution to this is resampling, that eliminates particles that have negligible weights and replaces them by copies of others with larger weights. Thus the unbiased PF representation of the posterior is

$$p(\mathbf{x}_t | \mathbf{y}_{1:t}) \approx \sum_{i=1}^M \bar{w}_t^i \delta(\mathbf{x}_t - \bar{\mathbf{x}}_t^i) = \sum_{i=1}^N \frac{n_t^i}{N} \delta(\mathbf{x}_t - \bar{\mathbf{x}}_t^i) \quad (4)$$

where n_t^i denotes the number of replications of the i th particle determined in accordance to its weight. Conventional resampling methods that give an unbiased representation of the posterior [4], [5] obtains a new weighted set of N particles $\{\bar{\mathbf{x}}_t^i, \bar{w}_t^i\}_{i=1}^{M=N} \rightarrow \{\mathbf{x}_t^i, w_t^i\}_{i=1}^N$ as follows, for $i = 1, \dots, N$, an index is sampled $j(i)$ distributed according to the probability $P(j(i) = m) = \bar{w}_t^m, m = 1, \dots, N$ and assign $\mathbf{x}_t^i = \bar{\mathbf{x}}_t^{j(i)}$ and set $w_t^i = 1/N$. Then (4) changes to

$$\hat{p}(\mathbf{x}_t | \mathbf{y}_{1:t}) \approx \sum_{i=1}^N w_t^i \delta(\mathbf{x}_t - \mathbf{x}_t^i) = \sum_{i=1}^N \frac{1}{N} \delta(\mathbf{x}_t - \mathbf{x}_t^i) \quad (5)$$

This resampling operation, being sequential and also exhaustive in nature over all the particles, renders the PF computationally expensive. Subsequently this paper proposes a minimalistic resampling strategy that will substantially speed up the PF.

III. PROPOSED MINIMALISTIC RESAMPLING

Here a deterministic redistributive resampling scheme is proposed to obtain $\{\bar{\mathbf{x}}_t^i, \bar{w}_t^i\}_{i=1}^N \rightarrow \{\mathbf{x}_t^i, w_t^i\}_{i=1}^N$ to satisfy an unbiased representation of the posterior. Consider a set of normalised weighted particles at time t that are sorted in accordance to their weights as $\{\bar{\mathbf{x}}_t^i, \bar{w}_t^i\}_{i=1}^N : \bar{w}_t^1 \geq \bar{w}_t^2 \geq \dots \geq \bar{w}_t^N$. Determine a set of indices corresponding to very low weights as $\{i : \bar{w}_t^i \leq T_h/N\}$ and then from it determine the number of large weights particles that would replicate the low weight ones as

$$L = N - |\{i : \bar{w}_t^i \leq T_h/N\}| \quad (6)$$

The proposed minimalistic resampling method eliminates only those low particles with low weights and replaces them with copies of the first L large weight particles. The aim here is to obtain a fixed number of replications for the first L particles as $\{n_t^i\}_{i=1}^L$ such that the total would count for all the N particles as $\sum_{i=1}^L n_t^i = N$. The way to do this is as follows.

Out of the first L particles, the first M particles will be replicated as

$$M = N - \left(\left\lfloor \frac{N-L}{L} \right\rfloor \right) + 1 \quad (7)$$

$$n_t^i = \left\lfloor \frac{N-L}{L} \right\rfloor + 2, i = 1, \dots, M \quad (8)$$

and the remaining $L-M$ particles will be replicated as

$$n_t^i = \left\lfloor \frac{N-L}{L} \right\rfloor + 1, i = M+1, \dots, L \quad (9)$$

This implies that the first L particles are redistributed within all the N particles. The weight of each of the L replicated particles is preserved by redistributing it within the replications of the particle. Thus the resampling indices and resampled particles and weights are

$$\begin{aligned} \{j(i)_{i=1, \dots, N}\} &= \left\{ \underbrace{1, \dots, 1}_{n_t^1 \text{ times}}, \underbrace{2, \dots, 2}_{n_t^2 \text{ times}}, \dots, \underbrace{L, \dots, L}_{n_t^L \text{ times}} \right\} \\ \{\mathbf{x}_t^i\}_{i=1}^N &= \left\{ \underbrace{\bar{\mathbf{x}}_t^1, \dots, \bar{\mathbf{x}}_t^1}_{n_t^1 \text{ times}}, \underbrace{\bar{\mathbf{x}}_t^2, \dots, \bar{\mathbf{x}}_t^2}_{n_t^2 \text{ times}}, \dots, \underbrace{\bar{\mathbf{x}}_t^L, \dots, \bar{\mathbf{x}}_t^L}_{n_t^L \text{ times}} \right\} \\ \{w_t^i\}_{i=1}^N &= \left\{ \underbrace{\frac{\bar{w}_t^1}{n_t^1}, \dots, \frac{\bar{w}_t^1}{n_t^1}}_{n_t^1 \text{ times}}, \underbrace{\frac{\bar{w}_t^2}{n_t^2}, \dots, \frac{\bar{w}_t^2}{n_t^2}}_{n_t^2 \text{ times}}, \dots, \underbrace{\frac{\bar{w}_t^L}{n_t^L}, \dots, \frac{\bar{w}_t^L}{n_t^L}}_{n_t^L \text{ times}} \right\} \end{aligned}$$

The weight of the discarded particles $w_{\text{spare}} = 1 - \sum_{i=1}^N w_t^i = \sum_{i=L+1}^M \bar{w}_t^i$ is redistributed among all the weights as

$$w_t^i = \tilde{w}_t^i + w_{\text{spare}}/N \quad (10)$$

such that the final weights sum to one.

It can be observed that the proposed method eliminates only the lowest weight particles and replaces them with the higher weight ones in a way that preserves the total weight. The approach is simple, extremely fast and importantly, gives and unbiased estimate of the posterior. The method is termed *minimalistic* as only those particles that should be replicated are treated within the resampling process and no computational effort is wasted on those that will be discarded. It is noteworthy that the proposed minimalistic resampler will be sensitive to the choice of the threshold that determines the set of low weights that should be discarded. This sensitivity will be discussed along with the evaluation in subsequent sections.

IV. BIAS ANALYSIS

In this section, the proposed resampler is analysed on the basis of its unbiasedness in preserving the total weight. A resampler is said to be unbiased if the weight of a particle is completely preserved, i.e., the difference between the actual weight and the total weight in each of its copies should be zero. The error in the weight preservation is $e^i = \bar{w}^i - \mathbb{E}(\sum_{j(i)=i: \mathbf{x}^j = \bar{\mathbf{x}}^i} w^j)$ which should be ideally zero.

Lemma: *The error in the weight preservation condition for the proposed redistributive resampling is negative for $i = 1, \dots, L$ thereafter positive and decaying.*

Proof:

$$\begin{aligned}
e^i &= \bar{w}^i - \mathbb{E}\left(\sum_{j(i)=i: \mathbf{x}^j = \bar{\mathbf{x}}^i} w^j\right) \\
&= \bar{w}^i - \left(\sum_{j(i)=i} \left(\frac{\bar{w}^i}{n^i} + \frac{w_{\text{spare}}}{N}\right)\right) \\
&= \bar{w}^i - \frac{n^i \bar{w}^i}{n^i} - \frac{n^i w_{\text{spare}}}{N} \\
&= \begin{cases} -\frac{n^i w_{\text{spare}}}{N}, & i = 1, \dots, L, \text{ as } n^i > 0 \\ \bar{w}^i, & i = 1, \dots, L+1, \text{ as } n^i = 0 \end{cases} \quad (11)
\end{aligned}$$

The variance in the error is $\mathbb{V}(e^i) = 0$ because the approach is fully deterministic. The proposed method is compared with the




Fig. 1. The weights for $N = 100$ samples. The left panel corresponds to the case when many samples having moderate weight, and the right to the case when few samples having significant weight.

stochastic, the residual and partial deterministic and Metropolis resamplers. All results in this paper are averaged over 500 Monte Carlo runs. Two cases of weight sequences as shown

in Fig 1, (a) many samples having moderate weights (the left panel), and (b) fewer samples having significant weights (the right panel) are tested. Fig 2 shows the error $e^{i=1:N}$ for different resamplers for the weight sequence in the left panel of Fig 1. Fig 3 shows the error $e^{i=1:N}$ for different resamplers for the weight sequence in the right panel of Fig 1. It can

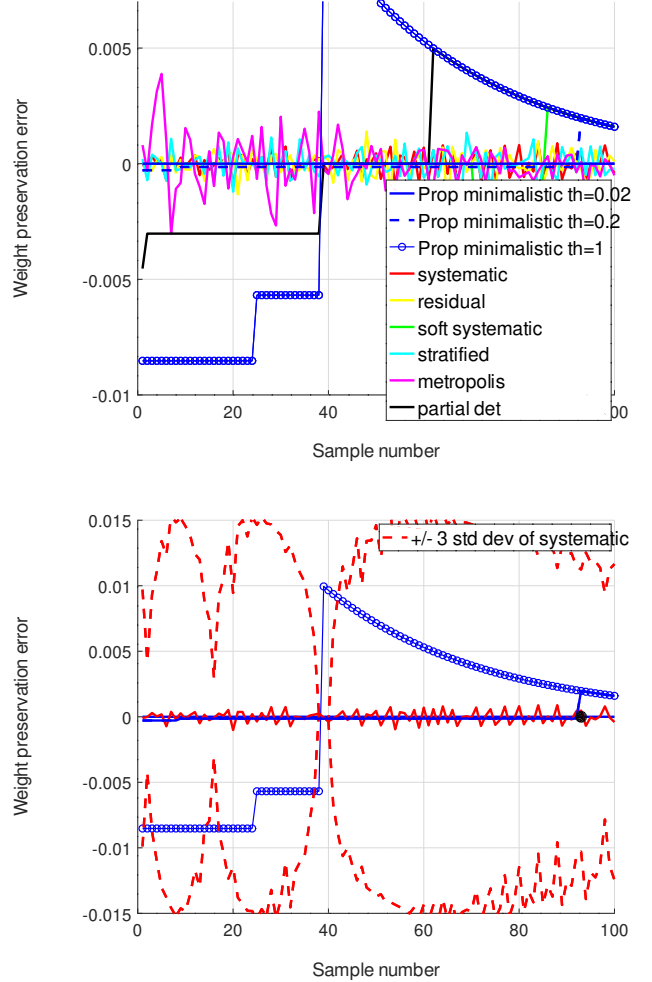


Fig. 2. The top panel shows the error e^i versus the sample index, and the bottom panel shows the error and error bounds of the systematic resampler and the error of the proposed resampler versus the sample index, for the weights in the left panel of Fig 1.

be observed from these figures that while all the resamplers show low error, the proposed minimalistic resampling, and systematic and residual resampling methods have the lowest error amongst them all. It can also be observed that the error in the proposed resampler lies within the bounds of the popularly used systematic resampler and the total weight is preserved. Moreover the error function $e^i, i = 1, \dots, L$ is negative for the first $L = 93$ samples in case (a) and for the first $L = 12$ samples in case (b), and thereafter positive and decaying (transition shown using black marker) for $T_h = 0.02, 0.2$ as derived in (11). It can also be seen that small $T_h \lll$ leads

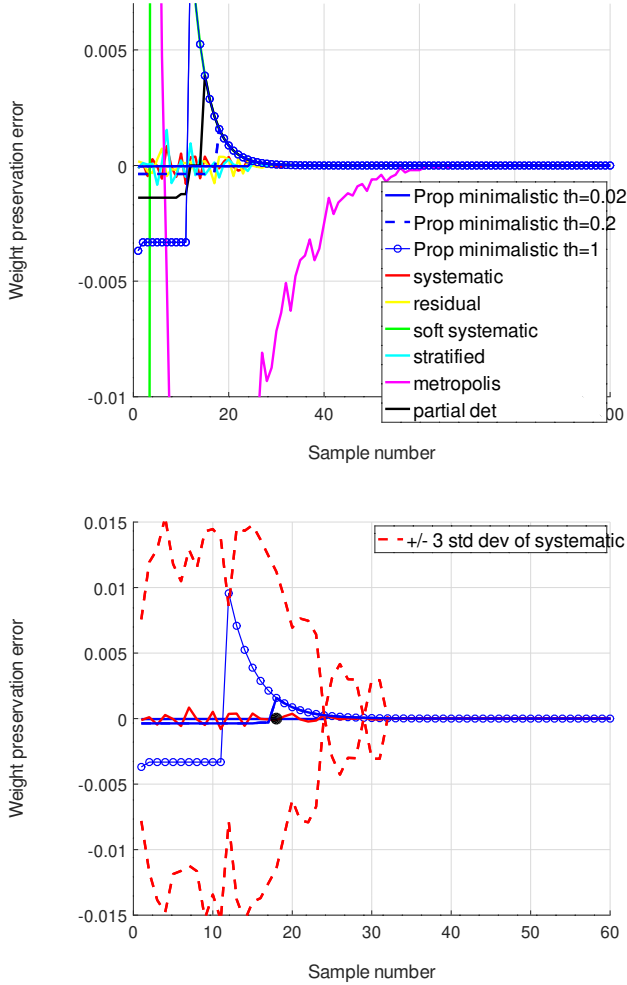


Fig. 3. The top panel shows the error e^i versus the sample index, and the bottom panel shows the error and error bounds of the systematic resampler and the error of the proposed resampler versus the sample index, for the weights in the right panel of Fig 1.

to no replacement and large $T_h \gg \gg$ leads to degeneracy; so a moderately low $T_h \in (0.1, 1)$, empirically tested, is a practical choice.

A. Posterior analysis for a simple channel estimation example

Consider a wireless communication system where the goal is to track a channel vector $\mathbf{x}_t \in \mathbb{R}^{d_x}$ where $d_x = 1$ is the number of dynamic coefficients to be estimated. The channel is estimated by sequentially transmitting pilots that are known to the receiver. For this, a simple linear Gaussian state space model is considered,

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{a}_t; \mathbf{y}_t = \mathbf{g}^\top \mathbf{x}_t + \mathbf{e}_t; t = 1, \dots, T = 100$$

where $\mathbf{F} = 0.9$, $\mathbf{g}^\top = 1$ and the noise densities are $\mathbf{a}_t \sim \mathcal{N}(0, \tau^2 = 1)$ and $\mathbf{e}_t \sim \mathcal{N}(0, \sigma^2 = 0.1)$. The initial actual target state is $\mathbf{x}_{t=0} = 10$ and the PFs are initialised with $p(\mathbf{x}_{t=0}) = \mathcal{N}(7, 1)$. For this model, the accuracy of

the PF with minimalistic resampling in representing the true posterior pdf (computed directly from the Kalman filter as a benchmark) according to the Kolmogorov-Smirnov (KS) statistic [16] is tested. For this we de-mean and de-correlate the final set of weighted particles using the Kalman mean and covariance. If indeed the particles are truly representative of the posterior, then the resultant will be a new set of particles with zero mean and unit covariance. This can be observed by comparing the cumulative distribution function (CDFs) or the weights with the new set and the CDF of the standard Gaussian pdf and then computing the maximum misfit between them, conventionally called the KS statistic (used in [15]). A low KS value is practically desirable. Fig 4 shows the CDFs for the minimalistic resampler with $T_h = 0.2$ and the popularly used systematic resampler. It can be observed that the proposal gives a unbiased posterior approximation as does the systematic resampler. It does not overshoot or undershoot at the tails implying that *every particle* is truly representative of the posterior. Fig 5 shows the KS statistic versus the number

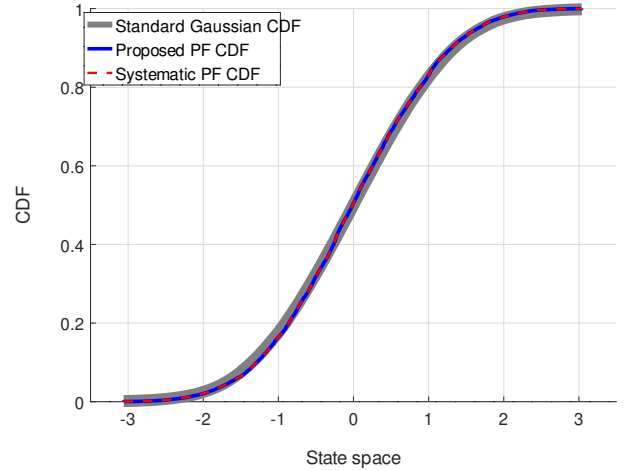


Fig. 4. The CDFs of the PFs plotted against the CDF of the standard Gaussian pdf (averaged over $T = 100$ time steps).

of particles. The proposed minimalistic resampler performs equivalently to the soft systematic resampler and closely to the computationally expensive systematic and residual resamplers. The higher KS values in the partial deterministic and Metropolis resamplers could be attributed to reduced communication within the particles. Low KS values indicate that the proposed resampler agrees closely to the Kalman filter posterior.

V. EVALUATION USING A TRACKING EXAMPLE

The proposed resampling method is now applied to a non-linear bearings-only tracking application [2], [5]. The target state is defined as $\mathbf{x}_t = (x_t, v_{x_t}, y_t, v_{y_t})^\top \in \mathbb{R}^4$ where the first and third entries are the $x - y$ target positions and the remaining are the corresponding velocities. The target moves




Fig. 5. The KS statistic versus the number of particles. in the $x - y$ plane via constant velocity (CV) motion model described as

$$\mathbf{x}_t = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{x}_{t-1} + \begin{bmatrix} 0.5 & 0 \\ 1 & 0 \\ 0 & 0.5 \\ 0 & 1 \end{bmatrix} \mathbf{a}_t$$

for $t = 1, \dots, T = 40$, where $\mathbf{a}_t \sim \mathcal{N}(\mathbf{0}, \tau^2 \mathbf{I}_2)$, $\tau^2 = 0.05^2$. The sensor located at the origin $(0,0)$ observes the target heading via the observation model given by

$$\mathbf{y}_t = \tan^{-1}(y_t/x_t) + \mathbf{e}_t$$

where $\mathbf{e}_t = \mathcal{N}(\mathbf{0}, \sigma^2 = 0.005^2)$. The initial target state is $\mathbf{x}_{t=0} = [-0.05, 0.001, 1.7, -0.055]^\top$. The PFs are initialised with mean $[-0.4, 0, 1.4, -0.5]^\top$ and covariance $\text{diag}(0.5, 0.005, 0.3, 0.01)$. Figs 6 and 7 respectively illustrate the particles using systematic resampling and the proposed minimalistic resampling at different time steps. The systematic resampler is set as a benchmark with $N = 10000$. The proposed method is run with $N = 2000$ and $T_h = 0.2$. It can be observed that the particle clusters in the proposed method conform to the particle clusters in the systematic resampler.

Fig 8 shows the RMSE versus the number of particles N and it can be seen that the proposed resampler compares favourably with the conventional resamplers. The time comparison is shown in Fig 9 and it can be observed that the proposed resampler exhibits orders of magnitude speed efficiency over the systematic resampler. At $N = 500, 5000$ particles by virtue of minimal particle interaction, the proposed method is nearly 3.8, 15.78 times faster than systematic resampling while maintaining equivalent tracking accuracy. In this analysis, it can be observed that the partial deterministic resampling is computationally very efficient. Unlike the proposed method, the partial deterministic resampler uses two thresholds to determine large weight and low weight particles. It can be observed in Fig 2 that the method has higher error in large weight particles than the proposed method due to the way the large weights are redistributed within. Moreover the performance




Fig. 6. The particles at times $t = 1, 3, 5, 10, 15, 20, 25, 30, 35, 40$ for the bearings-only example for one realisation.




Fig. 7. The particles at times $t = 1, 3, 5, 10, 15, 20, 25, 30, 35, 40$ for the bearings-only example for one realisation.

of the method is sensitive to the choice of the thresholds, that being said, this analysis helps to see that the method is nearly comparable to the minimalistic resampler in terms of speed.

For an unbiased approximation to the posterior, the number of particle replications should be $n_t^i = N \bar{w}_t^i, i = 1, \dots, N$. This is called the proper weighting condition [4]. A reliable measure of the resampling quality is the variance in the distance between the integral approximations (4) and (5). This variance is given by (see [5])

$$\mathbb{E} \left[\sum_{i=1}^N \frac{n_t^i - N \bar{w}_t^i}{N} \delta(\mathbf{x}_t) \right]. \quad (12)$$

By restricting n_t^i to lie close to $N \bar{w}_t^i$ for $i = 1, \dots, N$ reduces the variance and hence satisfies proper weighting criterion. Fig 10 shows this variance versus the number of particles

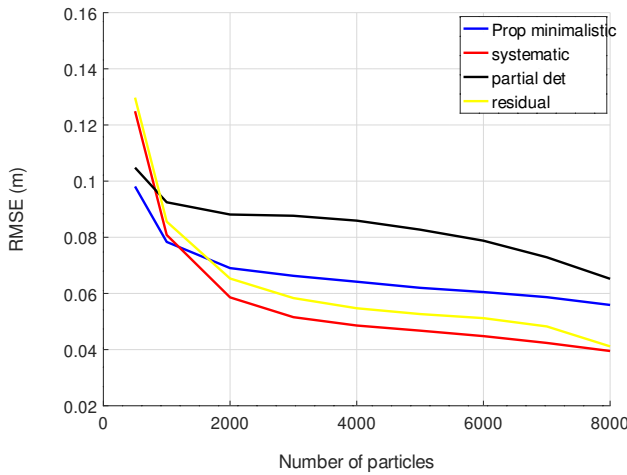


Fig. 8. The RMSE versus the number of particles N .

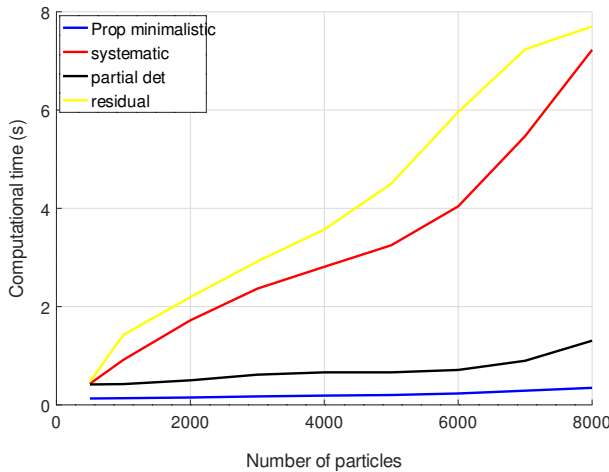


Fig. 9. The RMSE versus the number of particles N .

with $T_h = 0.2$ and it can be seen that the variance is close to zero indicating that the proposed resampler indeed replicates particles with a close-to-proper-weighting condition and hence gives an unbiased estimate of the posterior.

VI. CONCLUSION

This paper proposed a fast resampling scheme for the particle filter. The key innovation is to minimise the communication within the particles by deterministically replicating only the large weight ones. The proposal compares equivalently to the conventional systematic resampler in giving an unbiased representation of the true posterior and tracks accurately. This is illustrated using simulations on linear and nonlinear models.

REFERENCES

[1] P.M. Djuric, J. Zhang, T. Ghirmai, Y. Huang, and JH Kotecha, "Applications of particle filtering to communications: A review," in Proc. IEEE European Signal Processing Conference, pp. 1–4, 2002.

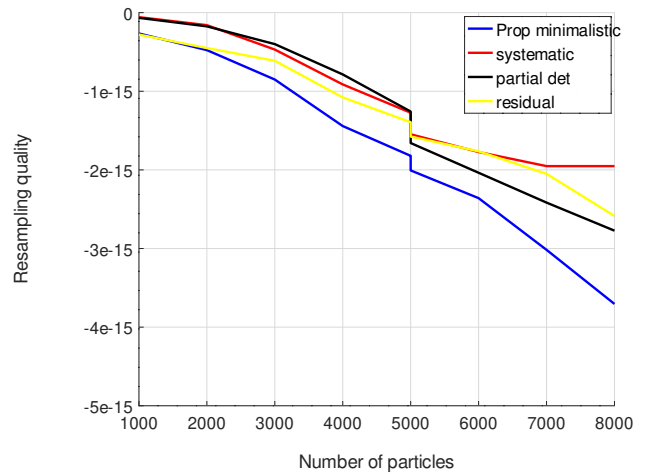


Fig. 10. The resampling quality versus N .

- [2] N. Gordon, David J. Salmond, and Adrian FM Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," in IEE proceedings F (radar and signal processing), vol. 140, no. 2, pp. 107–113, 1993.
- [3] P.M. Djuric, J.H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M.F. Bugallo, and J. Miguez, "Particle filtering," IEEE Signal Proc. Mag., vol. 20, no. 5, pp. 19–38, 2003.
- [4] R. Douc, and O. Cappe, "Comparison of resampling schemes for particle filtering," in Proc. IEEE Symp. on Image and Signal Processing and Analysis, pp. 64–69, 2005.
- [5] J. D. Hol, T. B. Schon, and F. Gustafsson, "On resampling algorithms for particle filters," Proc. IEEE Nonlinear Statistical Signal Processing Workshop, pp. 79–82, 2006.
- [6] G. Kitagawa, "Monte Carlo filter and smoother for non-Gaussian nonlinear state space models," J. of Computational and Graphical Statistics, pp. 1–25, 1996.
- [7] J. Carpenter, P. Clifford, and P. Fearnhead, "An improved particle filter for nonlinear problems," IEE Proc. Radar, Sonar and Navigation, vol. 146, pp. 2–7, 1999.
- [8] J. Liu, and R. Chen, "Sequential Monte Carlo Methods for Dynamic Systems," J. of the American Statistical Association, vol. 93, pp. 1032–1044, 1998.
- [9] M. Bolic, Petar M. Djuric, and Sangjin Hong, "Resampling algorithms for particle filters: A computational complexity perspective," Proc. EURASIP J. on Advances in Signal Processing, pp. 2267–2277, 2004.
- [10] P. Fearnhead, "Particle filters for mixture models with an unknown number of components," J. Statistics and Computing, vol. 14, no. 1, pp. 11–21, 2004.
- [11] S. Barenbruch, A. Garivier, and Eric Moulines, "On approximate maximum-likelihood methods for blind identification: How to cope with the curse of dimensionality," IEEE Trans. Signal Processing, vol. 57, no. 11, pp. 4247–4259, 2009.
- [12] P.B. Choppala, P.D. Teal, and M.R. Frean, "Soft systematic resampling for accurate posterior approximation and increased information retention in particle filtering," in Proc. IEEE Workshop on Statistical Signal Processing, pp. 260–263, 2014.
- [13] L. M. Murray, Anthony Lee, and Pierre E. Jacob, "Parallel resampling in the particle filter," J. of Computational and Graphical Statistics, vol. 25, no. 3, pp. 789–805, 2016.
- [14] Mehdi Chitchian, Andrea Simonetto, Alexander S. van Amesfoort, and T. Keviczky, "Distributed Computation Particle Filters on GPU Architectures for Real-Time Control Applications," IEEE Trans. Control Systems Technology, vol. 21, no. 6, pp. 2224–2238, 2013.
- [15] P.B. Choppala, P.D. Teal, and M.R. Frean, "Particle filter parallelisation using random network based resampling," in Proc. IEEE International Conference on Information Fusion (FUSION), pp. 1–8, 2014.
- [16] N. Smirnov, "Table for estimating the goodness of fit of empirical distributions," in J. Annals of Mathematical Statistics, pp. 279–281, 1948.