JOINT EMITTER DETECTION AND TRACKING USING DISTRIBUTED RANDOM EXCHANGE DIFFUSION PARTICLE FILTERING

Stiven S. Dias

Embraer Defense & Security
R&D Department
stiven.dias@embraer.com.br

Marcelo G. S. Bruno

Instituto Tecnológico de Aeronáutica
Eletrical Engineering Division
bruno@ita.br

ABSTRACT

We introduce in this paper a new fully distributed particle filter (PF) algorithm based on random information diffusion that is capable of performing joint multi-frame detection and tracking of a single moving emitter using a cooperative network of multiple received-signal-strength (RSS) sensors. Unlike previous consensus-based distributed PF schemes, the proposed Random Exchange Diffusion Particle Filter (ReDif-PF) does not require multiple iterative inter-node communication in the time interval between the arrival of two consecutive sensor measurements. Inter-node communication cost is further reduced by suitable parametric approximations.

Index Terms— Distributed Particle Filters, RSS Emitter Tracking, Diffusion, Joint Detection and Tracking

1. INTRODUCTION

In modern engineering systems, multiple processors with sensing and communication capabilities of their own cooperate to execute a global task without forwarding their local measurements to a data fusion center. Most existing distributed estimation algorithms, e.g. distributed Kalman filters [1], [2], [3] rely, however, on the assumption of linear, Gaussian state-space models. Distributed particle filtering [4] is an emerging technique that seeks to overcome the limitations of linear distributed Kalman filters in nonlinear/non-Gaussian scenarios. In particular, consensus-based distributed PF algorithms can reasonably approximate [5], [6], [7] or exactly reproduce [8] the global state estimate generated by a centralized PF running on a data fusion center. However, they require multiple iterative inter-node communication between consecutive sensor measurements, which is undesirable in real-world applications.

To circumvent the communication-cost limitations of consensus-based PF, we introduced in [9] the suboptimal Random Exchange Diffusion Particle Filter (ReDif-PF), which does not require iterative message passes between nodes. ReDif-PF diffuses information over the network by building over time, at each network node, different Monte Carlo representations of the posterior distribution of the target state conditioned on different sets of measurements coming from random locations in the entire network.

The ReDif-PF algorithm in [9] was applied specifically to a problem of cooperative emitter tracking using a network of received-signal-strength (RSS) sensors assuming that the target is always present. In this paper, we modify the algorithm in [9] to perform joint detection and tracking with a probability of detection less than one and a probability of false alarm greater than zero. Unlike the conventional contact/association approach [10] that decouples detection and tracking, we follow the methodology in [11] that enables integrated, multi-frame detection and tracking from raw sensor measurements only, but now adapted to a scenario with multiple observers and fully distributed processing over a partially-connected network where each node is allowed to communicate only with its immediate neighbors.

The paper is divided into 6 sections. Sec. 1 is this Introduction. Sec. 2 describes the emitter's state model and the RSS sensor model. In Sec. 3, we review the optimal centralized PF detector/tracker. Sec. 4 introduces the new ReDif-PF detector/tracker and discusses suitable parametric approximations that further reduce the inter-node communication cost. The performance of the algorithm is investigated in Sec. 5 using simulated data. Finally, we present our conclusions in Sec. 6.

2. TARGET AND OBSERVATION MODEL

The unknown target state at instant n is represented by a pair $(s_n, \mathbf{x}_n)$[1] where s_n is a *discrete* random variable that takes the value zero if the target is absent from the surveillance space at instant n, and the value one if the target is present. The real-valued random vector $\mathbf{x}_n \triangleq \begin{bmatrix} x_n & \dot{x}_n & y_n & \dot{y}_n \end{bmatrix}^T$ is in turn the hidden kinematic state vector of the target at time step n consisting of the positions and velocities of the target's centroid respectively in dimensions x and y.

Following the model in [11], an absent target enters the surveillance space $\mathbf{\Omega} \subset \mathbb{R}^2$ with probability p_a. Once a target becomes present, it can only become absent, i.e. change its

[1] For simplicity of notation, we use lowercase letters to denote both random variables/vectors and samples of random variables/vectors.

state s_n from one to zero, if it physically leaves the space $\mathbf{\Omega}$. For simplicity, we also assume in this paper that no more than one target may be present in the surveillance space at any given time instant. Denoting by $Pr(A)$ the probability of an event A, the aforementioned assumptions imply then that the state sequence $\{(s_n,\, \mathbf{x}_n)\}$ evolves in time according to a coupled probabilistic model such that $Pr(\{s_n = i\} \mid \{s_{n-1} = j\}\,,\, \mathbf{x}_{n-1})$ is equal to

$$\begin{cases} Pr(\{\mathbf{x}_n \in \mathbf{\Omega}\} \mid \mathbf{x}_{n-1}, \{s_{n-1} = j\}) & i = 1; j = 1 \\ 1 - Pr(\{\mathbf{x}_n \in \mathbf{\Omega}\} \mid \mathbf{x}_{n-1}, \{s_{n-1} = j\}) & i = 0; j = 1 \\ p_a & i = 1; j = 0 \\ 1 - p_a & i = 0; j = 0 \end{cases} \tag{1}$$

We assume a present target moves inside the surveillance grid according to a specified kinematic transition probability density function (p.d.f.) $f(\mathbf{x}_n|\mathbf{x}_{n-1})$ and that, following a target birth, the initial state of the new target is randomly sampled from a memoryless initial state p.d.f. $f_a(\mathbf{x}_n)$ such that the coupled transition p.d.f. $p(\mathbf{x}_n|\mathbf{x}_{n-1}, \{s_n = i\}, \{s_{n-1} = j\})$ is given by $f(\mathbf{x}_n|\mathbf{x}_{n-1})$ if i and j are equal to one and $f_a(\mathbf{x}_n)$ if $i = 1$ and $j = 0$.

Without loss of generality, we make $f_a(\mathbf{x}_n)$ a non-informative uniform p.d.f. and $f(\mathbf{x}_n|\mathbf{x}_{n-1}) = \mathcal{N}(\mathbf{x}_n|\mathbf{F}\mathbf{x}_{n-1}, \mathbf{Q})$, where $\mathcal{N}(\cdot|\boldsymbol{\mu}, \mathbf{\Sigma})$ denotes a multivariate normal p.d.f. with mean vector $\boldsymbol{\mu}$ and covariance matrix $\mathbf{\Sigma}$. The matrices $\mathbf{F}$ and $\mathbf{Q}$, parameterized by the sampling period Δ and the acceleration noise standard deviation σ_{accel}, are specified by the white noise acceleration model [10].

Finally, when a target disappears from the surveillance grid or remains absent, we arbitrarily set its kinematic state to an identically zero vector.

Observation Model The measurements $z_{r,0:n} = \{z_{r,0}, \ldots z_{r,n}\}$ in dBm at the r-th node of a network of R RSS sensors are modeled as

$$z_{r,n} = g_r(\mathbf{x}_n, s_n) + v_{r,n} \tag{2}$$

where $\{v_{r,n}\}$ represents a zero-mean i.i.d. Gaussian noise process with known variance σ_r^2. If $s_n = 1$, $g_r(\cdot, 1)$ is a nonlinear function given by [12]

$$g_r(\mathbf{x}, 1) = P_0 - 10\,\zeta_r\, \log\left(\frac{\|\mathbf{H}\mathbf{x} - \mathbf{x}_r\|}{d_0}\right) \tag{3}$$

where $\mathbf{x}_r$ represents sensor position, $||.||$ is the Euclidean norm, (P_0, d_0, ζ_r) are known model parameters, see [12] for a full description, and $\mathbf{H}$ is a 2×4 matrix such that $H(1,1) = H(2,3) = 1$ and $H(i,j) = 0$ otherwise. We also denote by $\mathbf{N}_r$ the set of nodes in the neighborhood of node r.

Otherwise, if $s_n = 0$, we assume for simplicity, unlike in previous work, e.g. [13], [14], that there are no clutter measurements and the sensors only record background noise, i.e $g_r(\mathbf{x}, 0) = 0$.

3. JOINT DETECTION AND TRACKING USING A CENTRALIZED NETWORK PARTICLE FILTER

At each instant n, the optimal centralized particle filter represents the mixed posterior distribution of $(s_n, \mathbf{x}_n)$ conditioned on all present and past network observations $z_{1:R,0:n}$ by a weighted set of particles, $\{(s_n^{(q)}, \mathbf{x}_n^{(q)})\}$ with corresponding weights $\{w_n^{(q)}\}$, $q \in \mathcal{Q} \triangleq \{1, \ldots, Q\}$, such that [11], [15]

$$Pr(\{s_n = i\}|z_{1:R,0:n}) \;\approx \sum_{q|s_n^{(q)}=i} w_n^{(q)} \tag{4}$$

$$E\{\mathbf{x}_n|\{s_n = 1\}, z_{1:R,0:n}\} \;\approx \sum_{q|s_n^{(q)}=1} \frac{w_n^{(q)}}{\sum_{l|s_n^{(l)}=1} w_n^{(l)}} \mathbf{x}_n^{(q)} \tag{5}$$

where $E\{.\}$ stands for expected value and $i \in \{0, 1\}$ in (4).

Using a blind importance function [15], [16], the particles $(s_n^{(q)}, \mathbf{x}_n^{(q)})$ are sequentially sampled, given $\mathbf{x}_{n-1}^{(q)}$ and $s_{n-1}^{(q)}$ as

$$s_n^{(q)} \;\sim\; P(s_n|s_{n-1}^{(q)}, \mathbf{x}_{n-1}^{(q)}) \tag{6}$$

$$\mathbf{x}_n^{(q)} \;\sim\; p(\mathbf{x}_n|\mathbf{x}_{n-1}^{(q)}, s_n^{(q)}, s_{n-1}^{(q)}) \tag{7}$$

where the functions on the right-hand side of Eq. (6) and (7) are defined respectively as $P(i|j, \mathbf{x}_{n-1}) = Pr(\{s_n = i\} \mid \{s_{n-1} = j\}\,,\, \mathbf{x}_{n-1}^{(q)})$ and $p(\mathbf{x}_n|\mathbf{x}_{n-1}^{(q)}, s_n^{(q)}, s_{n-1}^{(q)}) = p(\mathbf{x}_n| \mathbf{x}_{n-1}^{(q)}, \{s_n = i\}\,,\, \{s_{n-1} = j\})$ for $i, j \in \{0, 1\}$. The proper importance weights can in turn be recursively propagated as [15], [16], [17]

$$w_n^{(q)} \;\propto\; w_{n-1}^{(q)}\, p(z_{1:R,n}|\mathbf{x}_{0:n}^{(q)}, s_{0:n}^{(q)}, z_{1:R,0:n-1}) \tag{8}$$

$$= \; w_{n-1}^{(q)} \prod_{r=1}^{R} p(z_{r,n}|\mathbf{x}_n^{(q)}, s_n^{(q)}) \tag{9}$$

where line (9) follows from the assumption that the observation noise vectors $\{v_{r,n}\}$ are mutually independent for all $n \geq 0$ and all $r \in \mathcal{R} \triangleq \{1, \ldots, R\}$ and also independent of the hidden state sequence $\{(s_n, \mathbf{x}_n)\}$. In a centralized algorithm, all observations $z_{r,n}$ are sent to a data fusion center that computes the likelihood functions $p(z_{r,n}|\mathbf{x}_n^{(q)}, s_n^{(q)}) = \mathcal{N}(z_{r,n}|g_r(\mathbf{x}_n^{(q)}, s_n^{(q)}),\, \sigma_r^2)$ for all $r \in \mathcal{R}$ and all $q \in \mathcal{Q}$ and then updates the weights $\{w_n^{(q)}\}$ using Eq. (8). Note that if the sampled value of $s_n^{(q)}$ in (6) is 0, the new particle $\mathbf{x}_n^{(q)}$ is supposed to be outside the surveillance space $\mathbf{\Omega}$ and, according to (2), is not observable, thus it does not affect the calculus of the corresponding importance weight $w_n^{(q)}$ in (8) and, as previously stated, can be arbitrarily set to an identically zero vector to avoid unnecessary computational processing. The filter then builds the Monte Carlo approximation $Pr(\{s_n = i\}\,|z_{1:R,0:n})$, $i \in \{0, 1\}$, on the right-hand side of Eq. (4) and applies the minimum probability of error detection test

$$Pr(\{s_n = 0\}\,|z_{1:R,0:n}) \underset{H_1}{\overset{H_0}{\gtrless}} 1 - Pr(\{s_n = 0\} \mid z_{1:R,0:n}). \tag{10}$$

If the target is declared present, i.e. if hypothesis H_1 is accepted, its kinematic state $\mathbf{x}_n$ is then estimated using the Monte Carlo approximation to $E\{\mathbf{x}_n | \{s_n = 1\}, z_{1:R,0:n}\}$, i.e. we make the estimate $\hat{\mathbf{x}}_{n|n}$ equal to the right-hand side of Eq. (5).

4. DISTRIBUTED DETECTION AND TRACKING USING REDIF-PF

Given the independence assumption in (9), the global importance weights in (8) can be computed *exactly* in a fully distributed fashion without a data fusion center using, as shown in [8], R times D minimum consensus iterations where D is the diameter of the network graph. Synchronized importance sampling and resampling [18] at each node according to the consensus weights ensures that all nodes keep an identical set of particles at all time instants. To reduce the communication burden associated with the consensus-based solution, we follow a different approach in this paper.

Specifically, assume that, at instant $n-1$, node t has a properly weighted set of particles $\{(s_{t,0:n-1}^{(q)}, \mathbf{x}_{t,0:n-1}^{(q)})\}$ with associated weights $\{w_{t,n-1}^{(q)}\}$, which form a Monte Carlo representation for the joint posterior distribution of $s_{0:n-1}$ and $\mathbf{x}_{0:n-1}$ given $\mathcal{Z}_{t,0:n-1}$ where $\mathcal{Z}_{t,0:n-1}$ is the set of all measurements (possibly coming from different network nodes) that have been processed by node t up to instant $n-1$. Similarly, node r with $r \neq t$ has at instant $n-1$ a set of samples $\{(s_{r,0:n-1}^{(q)}, \mathbf{x}_{r,0:n-1}^{(q)})\}$ with weights $\{w_{r,n-1}^{(q)}\}$ that represent the joint posterior distribution of $s_{0:n-1}$ and $\mathbf{x}_{0:n-1}$ given $\mathcal{Z}_{r,0:n-1}$.

In the sequel, assume now that nodes r and t *exchange* their sample sets and respective weights at instant $n-1$. At instant n, the new particle set at node r, $(s_{r,0:n}^{(q)}, \mathbf{x}_{r,0:n}^{(q)}) = (s_{t,0:n-1}^{(q)}, s_{r,n}^{(q)}, \mathbf{x}_{t,0:n-1}^{(q)}, \mathbf{x}_{r,n}^{(q)})$ with updated weights $w_{r,n}^{(q)}$ such that

$$
\begin{aligned}
s_{r,n}^{(q)} &\sim P(s_n | s_{t,n-1}^{(q)}, \mathbf{x}_{t,n-1}^{(q)}) && (11)\\
\mathbf{x}_{r,n}^{(q)} &\sim p(\mathbf{x}_n | \mathbf{x}_{t,n-1}^{(q)}, s_{r,n}^{(q)}, s_{t,n-1}^{(q)}) && (12)\\
w_{r,n}^{(q)} &\propto w_{t,n-1}^{(q)} \prod_{l \in \mathbf{N}_r \cup \{r\}} p(z_{l,n} | \mathbf{x}_{r,n}^{(q)}, s_{r,n}^{(q)}), && (13)
\end{aligned}
$$

can be shown, using a procedure analogous to that in the appendix of reference [9], to be, under the conditional independence assumptions in the model, a properly weighted set to represent the updated joint posterior distribution of $s_{0:n}$ and $\mathbf{x}_{0:n}$ conditioned on the new set of measurements $\mathcal{Z}_{r,0:n} = (\mathcal{Z}_{t,0:n-1}, \mathcal{Z}_{r,n})$, where $\mathcal{Z}_{r,n} = \{z_{l,n}\}_{l \in \mathbf{N}_r \cup \{r\}}$.

Random Exchange Protocol To build, at each instant n and at each node r, a Monte Carlo representation of a posterior distribution conditioned on a set of observations $\mathcal{Z}_{r,0:n}$ consisting of measurements coming, at each time step, from different random locations in the entire network and not only from the node's immediate neighborhood, it suffices to implement a protocol where each node r, starting from instant zero, performs the aforementioned particle/weight exchange with a randomly chosen neighboring node t, propagates the received particles using the blind importance function as in Eqs. (11) and (12), and then updates their weights as in Eq. (13).

Each node r, at instant n, can then build its own Monte Carlo approximations

$$
Pr(\{s_n = i\} | \mathcal{Z}_{r,0:n}) \approx \sum_{q | s_{r,n}^{(q)} = i} w_{r,n}^{(q)} \qquad i \in \{0, 1\}, \quad (14)
$$

perform the detection test

$$
Pr(\{s_n = 0\} | \mathcal{Z}_{r,0:n}) \underset{H_1}{\overset{H_0}{\gtrless}} 1 - Pr(\{s_n = 0\} \mid \mathcal{Z}_{r,0:n}), \quad (15)
$$

and, if the target is declared present, estimate its kinematic state as

$$
E\{\mathbf{x}_n | \{s_n = 1\}, \mathcal{Z}_{r,0:n}\} \approx \sum_{q | s_{r,n}^{(q)} = 1} \frac{w_{r,n}^{(q)}}{\sum_{l | s_{r,n}^{(l)} = 1} w_{r,n}^{(l)}} \mathbf{x}_{r,n}^{(q)}. \quad (16)
$$

4.1. Low-Communication-Cost ReDif-PF

The ReDif-PF algorithm in Sec. 4 eliminates the need for iterative inter-node communication, but still requires the transmission of Q particle pairs and weights per node at each time step. To circumvent that limitation, we propose to use the weighted particle set $\{(s_{t,0:n-1}^{(q)}, \mathbf{x}_{t,0:n-1}^{(q)})\}$ with associated weights $\{w_{t,n-1}^{(q)}\}$ at node t at instant $n-1$ to compute the approximate posterior probabilities

$$
\widetilde{Pr}(\{s_{n-1} = i\} | \mathcal{Z}_{t,0:n-1}) = \sum_{q | s_{t,n-1}^{(q)} = i} w_{t,n-1}^{(q)} \qquad i \in \{0, 1\} \quad (17)
$$

and the parametric Gaussian Mixture Model (GMM) approximation $\tilde{p}(\mathbf{x}_{n-1} | \{s_{t,n-1} = 1\}, \mathcal{Z}_{t,0:n-1})$ given by

$$
\sum_{k=1}^{K} \eta_{t,n-1}^{(k)} \mathcal{N}(\mathbf{x}_{n-1} | \boldsymbol{\mu}_{t,n-1}^{(k)}, \boldsymbol{\Sigma}_{t,n-1}^{(k)}), \quad (18)
$$

which is obtained using the Expectation-Maximization (EM) algorithm [19]. Node t then transmits to node r at instant n the approximate posterior probabilities in (17), i.e. only one real number (since the two posteriors add up to one), and the parameters that specify the approximate p.d.f. in (18), i.e. $15K$ real numbers, where $K << Q$.

Node r then locally resamples $s_{t,n-1}^{(q)} \sim \widetilde{P}(s_{n-1} | \mathcal{Z}_{t,0:n-1})$, $q \in \boldsymbol{\mathcal{Q}}$, and, if the resampled $s_{t,n-1}^{(q)} = 1$, draws $\mathbf{x}_{t,n-1}^{(q)} \sim \tilde{p}(\mathbf{x}_{n-1} | \{s_{t,n-1} = 1\}, \mathcal{Z}_{t,0:n-1})$. Next, node r draws new particles $\left\{(s_{r,n}^{(q)}, \mathbf{x}_{r,n}^{(q)})\right\}$ as in (11) and (12), and updates its local weights using (13). Note that, if the resampled $s_{t,n-1}^{(q)} = 0$, the particle $\mathbf{x}_{t,n-1}^{(q)}$ is not necessary to sample the new particles $(s_{r,n}^{(q)}, \mathbf{x}_{r,n}^{(q)})$.

5. SIMULATION RESULTS

The performance of the ReDif-PF tracker was assessed using 100 independent Monte Carlo runs in a simulated scenario consisting of $R = 25$ RSS sensors deployed on a jittered grid within a squared surveillance space $\mathbf{\Omega}$ of size $110\,\text{m} \times 110\,\text{m}$. Furthermore, each node communicates with other nodes within a range of $40\,\text{m}$. The sensors' parameters were kept fixed during all simulations and set to $P_0 = 1\,\text{dBm}$, $d_0 = 1\,\text{m}$, $\zeta_r = 3$, $\forall r \in \boldsymbol{\mathcal{R}}$. The variances σ_r^2's were independently sampled according to an Inverse Gamma p.d.f. with mean 16.

The initial state distribution $f_a(\mathbf{x}_n)$ for newly initialized tracks was assumed an uniform p.d.f. on a square $\mathbf{\Omega}_a \subset \mathbf{\Omega}$ of size $30\,\text{m} \times 30\,\text{m}$ for the emitter's initial position in Cartesian coordinates and a Gaussian p.d.f. with mean vector $\left[\sqrt{2}\,\text{m/s} \quad 45°\right]^T$ and covariance matrix $diag(0.3^2, 5^2)$ for the emitter's initial velocity in polar coordinates.

Fig. 1 shows the sensor positions and two consecutive realizations of the emitter trajectory generated within a simulation period of $200\,\text{s}$ for $\Delta = 1\,\text{s}$, $\sigma_{accel} = 0.05\,\text{m/s}^2$ and $p_a = 0.75$. It also depicts the available network connections. The diameter of the sensor network is $D = 5$ hops and the minimum number of neighbors for any possible node is three.

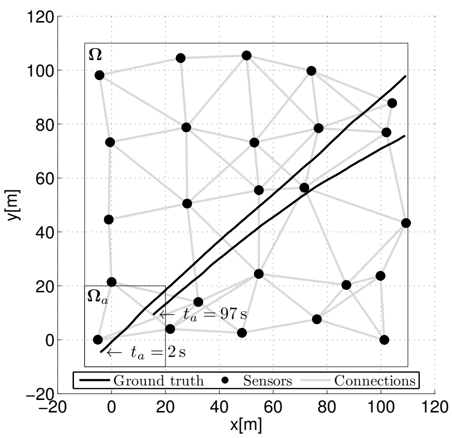

Fig. 1. Evaluated scenario.

All filters employed $Q = 500$ particles to estimate the emitter state. The particles at each node r were initialized at time step 0 with $s_{r,0}^{(q)} = 0$, $\forall q \in \boldsymbol{\mathcal{Q}}$. After a transition from $s_{r,n-1}^{(q)} = 0$ to $s_{r,n}^{(q)} = 1$ at a given time step n, a particle q had its state $\mathbf{x}_{r,n}^{(q)}$ initialized according to the distribution $f_a(\mathbf{x}_n)$.

Fig. 2 shows the evolution of the root-mean-square (RMS) error norm of the emitter position estimates for the proposed ReDif-PF using a single Gaussian, i.e. $K = 1$, and for the minimum-consensus CbPFa algorithm in [8], also adapted in this paper for joint detection and tracking. The bars shown in Fig. 2 along the ReDif-PF's curve represent the standard deviation of the error norm across all network nodes. The RMS error at time step 0 was calculated after the measurements $z_{1:R,0}$ were assimilated. Compared to CbPFa, which exactly mimics the optimal centralized PF [8], the ReDif-PF tracker has a greater RMS error since it assimilates less information at each network location than the former. Moreover, since the RMS error norm was computed at each time n considering just the first simulated emitter trajectory of each Monte Carlo run, it increases and then becomes more noisy for both algorithms after time instants 50 s and 100 s, respectively, as the first emitter tends to leave the surveillance space $\mathbf{\Omega}$ (see Fig. 1).

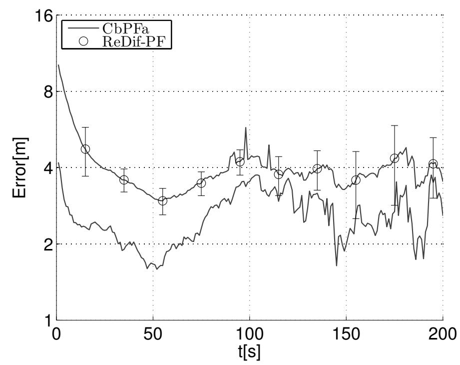

Fig. 2. Evolution of the estimated position RMS error norm.

Assuming a four-byte representation for real numbers, we recorded the total amount of bytes exchanged by all network nodes while running each filter. The average transmission (TX) and reception (RX) rates per node were then calculated by dividing respectively the average total amount of data transmitted and received considering all Monte Carlo runs by the simulation period (200 s) and, then, by the number of nodes R. Additionally, the average duty cycle per node was calculated by dividing the average total processing time at each node considering all Monte Carlo runs by the simulation period.

Table 1 summarizes the computed performance metrics for each algorithm. The ReDif-PF tracker has an average communication cost per node three orders of magnitude lower than that of the CbPFa with half the processing cost.

Table 1. Performance metrics of the evaluated algorithms.

	RX Rate	**TX Rate**	**Duty Cycle**
CbPFa	1.2 MB/s	244.1 KB/s	23.2 %
ReDif-PF	148 B/s	132 B/s	12.9 %

Detection Performance Finally, we estimated the average P_{fa} and P_d for the ReDif-PF detector dividing the total count of false alarms and accurate detections respectively by the total number of detection decisions made by all network nodes throughout the simulations, i.e. $201 \times 100 \times 25$. We obtained an average $P_d = 99.9\,\%$ with average $P_{fa} = 0.12\,\%$.

6. CONCLUSIONS

We introduced in this paper a new version of the fully distributed Random Exchange Diffusion Particle Filter which enables joint multi-frame detection and tracking of a single target in a partially connected network of RSS sensors using raw sensor data only. The proposed algorithm does not require iterative inter-node communication between sensor measurements and achieves high probability of detection with low false alarm rates and RMS tracking error close to that of the optimal centralized particle filter.

7. REFERENCES

[1] S. Kar and J. M. F. Moura, "Gossip and distributed Kalman filtering: weak consensus under weak detectability," *IEEE Transactions on Signal Processing*, v.59, n.4, pp. 1766-1784, April 2011.

[2] F. S. Cattivelli, and A. H. Sayed, "Diffusion Strategies for Distributed Kalman Filtering and Smoothing", *IEEE Transactions on Automatic Control*, vol. 55, n. 9, pp. 2069 - 2084, September 2010.

[3] A. Ribeiro, G. B. Giannakis, and S. I. Roumeliotis, "SOI-KF: Distributed Kalman Filtering With Low-Cost Communications Using the Sign of Innovations," *IEEE Transactions on Signal Processing*, vol. 54, n. 12, pp. 4782-4795, December 2006.

[4] O. Hlinka, F. Hlawatsch, and P. M. Djuric, "Distributed Particle Filtering in Agent Networks: A Survey, Classification, and Comparison," *IEEE Signal Processing Magazine*, vol. 30, no. 1, pp. 61-81, January 2013.

[5] O. Hlinka, O. Sluciak, F. Hlawatsch, P. Djurić, and M. Raupp "Likelihood Consensus and Its Application to Distributed Particle Filtering," *IEEE Transactions on Signal Processing*, vol. 60, n.8, pp. 4334-4349, August 2012.

[6] A. Mohammadi and A. Asif, "Consensus-based distributed unscented particle filter," *2011 IEEE Statistical Signal Processing Workshop*, Nice, France, June 2011.

[7] D. Uztebay, M. Coates, and M. Rabat, "Distributed Auxiliary Particle Filter using Selective Gossip," *IEEE Intl. Conf. on Acoustics Speech and Sig. Proc.*, Prague, Czech Republic, May 2011

[8] S. S. Dias and M. G. S. Bruno, "Cooperative target tracking using decentralized particle filtering and RSS sensors," *IEEE Transactions on Signal Processing*, vol. 61, no. 14, pp. 3632–3646, July 2013.

[9] S. S. Dias, and M. G. S. Bruno, "Distributed Emitter Tracking using Random Exchange Diffusion Particle Filters," *16th International Conference on Information Fusion*, Istanbul, Turkey, July 2013.

[10] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques.* YBS, Storrs, CT, 1995.

[11] M. G. S. Bruno, R. V. Araújo, and A. G. Pavlov, "Sequential Monte Carlo methods for joint detection and tracking of multiaspect targets in infrared radar images," *EURASIP Journal on Advances in Signal Processing*, 2008, doi:10.1155/2008/217373.

[12] N. Patwari, A. O. Hero III, M. Perkins, N. S. Correal, and R. J. O'Dea, "Relative location estimation in wireless sensor networks," *IEEE Transactions on Signal Processing*, vol. 51, n. 3, pp. 2137-2148, March 2003.

[13] B-T. Vo, D. Clark, B-N. Vo, and B. Ristic, "Bernoulli Forward-Backward Smoothing for Joint Target Detection and Tracking," *IEEE Transactions on Signal Processing*, vol. 59, n. 9, pp. 4473-4477, September 2011.

[14] R. P. Mahler, *Statistical Multisource-Multitarget Information Fusion.* Artech House Inc., Norwood, 2007.

[15] A. Doucet, S. J. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 3, pp. 197 - 208, 2000.

[16] M. G. S. Bruno, "*Sequential Monte Carlo methods for nonlinear discrete-time filtering*", Synthesis Lectures on Signal Processing, J. M. F. Moura, Editor, vol. 6, pp. 1-99, Morgan & Claypool Publishers, 2013.

[17] C. J. Bordin Jr. and M. G. S. Bruno, "Cooperative blind equalization of frequency-selective channels in sensor networks using decentralized particle filtering," in *42nd Asilomar Conf. on Sign., Syst. and Comp.*, pp. 1198–1201, Pacific Grove - CA, October 2008.

[18] M. J. Coates, "Distributed particle filtering for sensor networks," in *3rd Intl. Symp. on Inf. Proc. in Sensor Networks*, pp. 99–107, Berkeley - CA, April 2004.

[19] A. P. Dempster, N. M. Laird, D. B. Rubin (1977), "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society, Series B, 39 (1): 138*, 1977.