This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2010 proceedings

Bayesian Localization in Sensor Networks: Distributed Algorithm and Fundamental Limits

D. Fontanella[1], M. Nicoli[1] and L. Vandendorpe[2]
[1]Dip. di Elettronica e Informazione, Politecnico di Milano, Piazza L. da Vinci 32, I-20133 Milano, Italy
[2]Commun. & Remote Sensing Lab., Université Catholique de Louvain, Place du Levant 2, B-1348 Louvain la Neuve, Belgium

***Abstract*—Self-localization in ad-hoc sensor networks is becoming a crucial issue for several location-aware applications. This technology implies the combination of absolute anchor locations with relative inter-node information exchanged on a peer-to-peer basis. In this paper we investigate a distributed algorithm and fundamental performance bounds for Bayesian cooperative localization in stochastic networks. Nodes are assumed to be randomly deployed within a finite space according to a prior distribution. Bayesian inference is performed through an iterative local message passing procedure based on belief propagation and particle-filtering message representation. The algorithm performance is analyzed for a simplified scenario in which unknown node positions are randomly scattered along a line segment and anchors are fixed. Global Cramer-Rao bounds are derived and compared to the performance of the distributed algorithm.**

I. Introduction

Pervasive wireless sensor networks (WSN) are expected to revolutionize our ability to sense and control the surroundings by enabling a wide range of applications, such as environmental/habitat monitoring, disaster prevention and relief, assisted navigation, medical care, intelligent transportation systems. Networks for these applications are characterized by a topology that is dynamic and/or unknown prior to placement so that nodes have to self-organize and operate without the use of any existing infrastructure. Self-localization in particular is an extremely important feature [1]- [9], as sensor positions are often the first data to be achieved by the WSN (to match monitored quantities with spatial coordinates) and they are needed to enable efficient multi-hop communications.

In typical WSN deployments only a fraction of nodes (anchor nodes or ANs) have a measure of their own positions while the others (unknown nodes or UNs) have to be localized by the network itself through cooperative algorithms [1]- [2]. In this paper we are interested in a distributed approach in which each node estimates *locally* its own position using only noisy range measurements obtained from neighbors and some prior information, if available. We investigate distributed algorithms and fundamental performance limits for cooperative localization in a stochastic WSN where nodes are randomly deployed within a finite space according to a prior distribution.

Network localization is formulated as a distributed Bayesian inference problem. Starting from the results in [3]- [4], we consider a belief propagation (BP) algorithm that allows each node to compute the marginal a posteriori probability density function (pdf) of its own position by an iterative exchange of information (messages) with nearby nodes. Each node represents its beliefs through a random set of particles locally updated by importance re-sampling and weighting [10]. The main issue in particle-based BP localization is the fusion of messages obtained from the neighbors since they are defined over different sets of particles. In non-parametric BP (NBP) [3]- [4] data fusion is handled by modeling the messages as Gaussian mixtures and combining them either through Gibbs sampling or mixture importance sampling. Here we propose a modified particle filtering BP (PF-BP) approach which employs an efficient procedure for message fusion and a different weight computation for message update. Data fusion is obtained by smoothing the neighbors' messages with a kernel function (different functions are investigated) and resampling over a common set of particles. PF-BP is shown to provide a performance gain with respect to NBP.

To analyze the performance of distributed Bayesian localization in stochastic WSNs, we evaluate and compare the accuracy of the BP algorithm with fundamental performance limits. Without loss of generality, in this preliminary study we consider a one-dimensional (1D) network with nodes distributed along a line segment; the analysis can be extended to higher dimensions and more complex networks by adapting the Bayesian network model. While the deterministic (or local) Cramér-Rao bound (CRB) has been widely investigated for localization of a fixed configuration of nodes (see, e.g., [1], [6], [9] and references therein), here we propose a novel analysis based on Bayesian (or global) bounds that are valid for any set of unknown node positions drawn from the prior distribution. We derive the Bayesian CRB [12] and other types of global bounds for the considered localization scenario, comparing the results with the distributed PF-BP algorithm performance.

II. Network localization model

We consider a stochastic network composed of N nodes deployed over a finite space. As shown in Fig. 1, we assume that nodes are scattered along a straight line segment of length L. The k-th node location is denoted by the scalar coordinate $\theta_k \in [0, L]$, $k = 1, \ldots, N$. For convenience of notation, we assume that the first N_u nodes have unknown positions $\boldsymbol{\theta} = [\theta_1 \cdots \theta_{N_u}]^{\mathrm{T}}$, while the remaining $N_{\mathrm{a}} = N - N_{\mathrm{u}}$ nodes are anchors with known positions $\boldsymbol{\theta}_{\mathrm{a}} = [\theta_{N_u+1} \cdots \theta_N]^{\mathrm{T}}$. The UN positions $\boldsymbol{\theta}$ are modelled as independent random variables distributed within the 1D finite space with (known) probability

This work was supported by the European Commission in the framework of the FP7 Network of Excellence in Wireless COMmunications NEWCOM++ (contract n. 216715) and by the Ministero dell'Istruzione dell'Università e della Ricerca MIUR FIRB Integrated System for Emergency (InSyEme) project under the grant RBIP063BPH.

978-1-4244-6404-3/10/$26.00 ©2010 IEEE

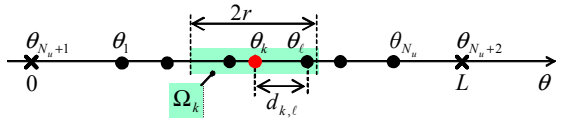

Fig. 1. Example of 1D WSN topology. Anchors are denoted as '×', while nodes of unknown positions as '•'.

density function $p(\boldsymbol{\theta}) = \prod_k p(\theta_k)$, while ANs are assumed to be placed in deterministic and fixed positions. An example of deployment, used later on for performance assessment, is shown in Fig. 1 with $N_\mathrm{a} = 2$ ANs placed at the ends of the line segment: $\theta_{N_\mathrm{u}+1} = 0$ and $\theta_{N_\mathrm{u}+2} = L$.

Cooperative estimation of the locations $\boldsymbol{\theta}$ is obtained from pairwise distance measurements $z_{k,\ell}$ made between any pair of (either AN or UN) nodes k and ℓ. To account for limited transmission coverage, we assume that each node can make measurements only to nodes within a distance r: $\Omega_k = \{\theta_\ell : |\theta_k - \theta_\ell| \leq r\}$. For $\ell \in \Omega_k$, the measurement from node k to node ℓ is modelled as unbiased Gaussian, according to:

$$z_{k,\ell} = |\theta_k - \theta_\ell| + e_{k,\ell} = d_{k,\ell} + e_{k,\ell}, \tag{1}$$

where $d_{k,\ell} = |\theta_k - \theta_\ell|$ is the distance between the two nodes and $e_{k,\ell} \sim \mathcal{N}(0, \sigma^2_{k,\ell})$ denotes the Gaussian measurement uncertainty with zero mean and variance $\sigma^2_{k,\ell}$. Errors associated with different links, i.e. $e_{m,n}$ and $e_{k,\ell}$ with $(m,n) \neq (k,\ell)$, are assumed to be independent to each other.

Distance measurements $\{z_{k,\ell}\}$ can be obtained by estimating either the power or delay of the radio signal [1]. Times of arrival (TOA) are known to provide more accurate localization especially in ultra-wide band (UWB) systems [6], [8]. The model in (1) is widely adopted to describe these type of observations; possible bias due to non-line-of-sight propagation can be included as discussed in [1], [7]. Regarding synchronization, we recall that nodes do not need to be synchronized relatively to one another when round trip TOA measurements are available (i.e., both $z_{k,\ell}$ and $z_{\ell,k}$ for each pair of nodes k, ℓ) as clock biases can be canceled.

III. Review of Belief Propagation

We are interested in a Bayesian estimation of the random vector $\boldsymbol{\theta}$ given the set of noisy observations $\mathbf{z} = \{z_{k,\ell}\}$ and the prior information $p(\boldsymbol{\theta})$. Here we restrict our attention to a distributed approach, in which the k-th node localizes itself by computing the marginal a posteriori pdf of its own location $p(\theta_k|\mathbf{z})$. An approximate solution is provided by the BP approach as recalled in the sequel (see [4] for more details).

The self-localization problem is formulated as an undirected graphical model (or Markov random field) in which the vertices are the nodes, with associated random positions θ_k, and the edges connect nodes that exchange measurements with each other. Vertices and edges are associated with potential functions, namely the single-node potential for the k-th vertex, $\psi_k(\theta_k) = p(\theta_k)$, and the pairwise potential for the edge (k,ℓ), $\psi_{k,\ell}(\theta_k, \theta_\ell) = P_0(\theta_k, \theta_\ell)\, p(z_{k,\ell}|\theta_k, \theta_\ell)$, with $P_0(\theta_k, \theta_\ell) = P(\ell \in \Omega_k)$. In the localization scenario herein considered, the measurement likelihood $p(z_{k,\ell}|\theta_k, \theta_\ell)$ is, from (1), Gaussian with mean $|\theta_k - \theta_\ell|$ and variance $\sigma^2_{k,\ell}$.

At the iteration t of the BP procedure, node k receives from each neighbor $\ell \in \Omega_k$ a message, $m^t_{\ell\to k}(\theta_k)$, which represents the belief about node-k position based on node-ℓ observations. The node k approximates its a posteriori pdf as:

$$\hat{p}^t(\theta_k|\mathbf{z}) \propto \psi_k(\theta_k) \prod\nolimits_{\ell\in\Omega_k} m^t_{\ell\to k}(\theta_k), \tag{2}$$

where $\propto$ stands for proportional. The new location estimate $\hat{\theta}^t_k$ can be drawn from (2) using either the maximum a posteriori (MAP) or the minimum mean square error (MMSE) criterion; here we select the former: $\hat{\theta}^t_k = \arg\max_{\theta_k} \hat{p}^t(\theta_k|\mathbf{z})$. The a posteriori pdf is then used to compute the message $m^{t+1}_{k\to\ell}(\theta_k)$ to be transmitted from node k to node ℓ for the next iteration:

$$m^{t+1}_{k\to\ell}(\theta_\ell) \propto \int \psi_{k,\ell}(\theta_k, \theta_\ell) \frac{\hat{p}^t(\theta_k|\mathbf{z})}{m^t_{\ell\to k}(\theta_k)} d\theta_k. \tag{3}$$

Steps (2)-(3) are repeated till convergence, i.e. till $|\hat{\theta}^t_k - \hat{\theta}^{t-1}_k| \leq \varepsilon$ where ε is a given threshold. The algorithm is initialized with $\hat{p}^0(\theta_k|\mathbf{z}) = \psi_k(\theta_k)$.

The main issue for practical implementation of the BP algorithm is the efficient representation of beliefs/messages to be calculated/exchanged among nodes. A grid-based approach relying on a regular sampling of the location space has an unacceptably high cost for both computation and transmission of beliefs. A more efficient method has been proposed in [3] based on Gaussian mixture representation of beliefs/messages. In the next section we propose a modified approach using particle filtering (PF). Fusion of messages from neighbors, that use different sets of particles, is obtained by an efficient re-sampling over an auxiliary common set of particles. The fusion method is less complex than techniques usually adopted for combining PF and BP, based on Gibbs sampling, as it is linear in the number of particles (while Gibbs sampling is quadratic). An alternative PF technique has been proposed recently in [11] for a generic Bayesian estimation problem, with messages from neighbors sharing the same set of particles. A comparison of all these methods will be carried out in Sec. VI.

IV. Particle Filter for Belief Representation

Consider the iteration t, assume that node k receives from neighbor $\ell \in \Omega_k$ a message in the particle-based form [10]:

$$m^t_{\ell\to k}(\theta_\ell) \approx \sum\nolimits_{n=1}^{N_p} \gamma^t_{\ell\to k,n}\, \delta\left(\theta_\ell - \theta^t_{\ell\to k,n}\right), \tag{4}$$

with weights normalized so that $\sum_{n=1}^{N_p} \gamma^t_{\ell\to k,n} = 1$. Based on this information, node k has to evaluate the particles $\{\theta^t_{k,n}\}$ and weights $\{\omega^t_{k,n}\}$ for its own a posteriori belief:

$$\hat{p}^t(\theta_k|\mathbf{z}) \approx \sum\nolimits_{n=1}^{N_p} \omega^t_{k,n} \delta\left(\theta_k - \theta^t_{k,n}\right). \tag{5}$$

According to importance sampling [10], particles $\{\theta^t_{k,n}\}$ can be obtained as i.i.d. samples drawn from the a posteriori pdf in (2). This belief, however, depends on the product of messages $m^t_{\ell\to k}(\theta_k)$ coming from different neighbors $\ell \in \Omega_k$

and thus sampled over different particle sets $\{\theta^t_{\ell\to k,n}\}_{n=1}^{N_p}$. A possibility to approximate the product in (2) is to smooth and resample the messages over a common set of particles. Here we propose to join all neighbors' particle in the common set $\{\tilde\theta^t_{k,i}\}_{i=1}^{\tilde N_p} = \bigcup_{\ell\in\Omega_k}\{\theta^t_{\ell\to k,n}\}_{n=1}^{N_p}$, with $\tilde N_p = |\Omega_k|\cdot N_p$ (a smaller set could be obtained by selecting a reduced number of significant samples from each message). To allow the computation of messages over the new particle set, we assume that each i-th particle $\tilde\theta^t_{k,i}$ represents a *finite* interval of positions of width L_w centered around the value $\tilde\theta^t_{k,i}$. The corresponding probability, $\tilde\gamma^t_{\ell\to k,i} = P(\theta_k\in[\tilde\theta^t_{k,i}-L_w/2, \tilde\theta^t_{k,i}+L_w/2])$, can be calculated by integrating the message (4) from node ℓ over the considered bin:

$$\tilde\gamma^t_{\ell\to k,i} = \int m^t_{\ell\to k}(\theta_k)w(\theta_k-\tilde\theta^t_{k,i})d\theta_k = \sum_{n=1}^{N_p}\gamma^t_{\ell\to k,n}w(\theta^t_{\ell\to k,n}-\tilde\theta^t_{k,i}) \tag{6}$$

where $w(\theta)$ is a normalized symmetric windowing function[1] with support L_w ($w(\theta)=0$ for $|\theta|\geq L_w/2$). The method can be seen as a smoothing of messages (4) with kernel $w(\cdot)$ and resampling over the new set of particles. Given the re-sampled messages, the weights associated with the new particles of the marginal a posteriori pdf can be obtained from (2) as:

$$\tilde\omega^t_{k,i} \propto \psi_k\left(\tilde\theta^t_{k,i}\right)\prod\nolimits_{\ell\in\Omega_k}\tilde\gamma^t_{\ell\to k,i}. \tag{7}$$

An example of message fusion is shown in Fig. 3 for the node $k=3$ of a network of five UNs and two ANs. Once the samples $\{\tilde\theta^t_{k,i},\tilde\omega^t_{k,i}\}$ have been evaluated, resampling is performed by drawing a new set of particles $\{\theta^t_{k,n}\}_{n=1}^{N_p}$ from the discrete belief $P(\theta^t_{k,n}=\tilde\theta^t_{k,i})=\tilde\omega^t_{k,i}$. The corresponding weights are set all equal to $\omega^t_{k,n}=1/N_p$.

Based on the updated marginal a posteriori belief, represented by the set $\{\theta^t_{k,i},\omega^t_{k,i}\}$, node k has to evaluate particles and weights for the message to be sent to each neighbor ℓ. Similarly to [3], particles $\{\theta^{t+1}_{k\to\ell,i}\}_{i=1}^{\tilde N_p}$ are generated using as importance density the belief for the node-ℓ position approximated at node k, here indicated as $\hat p^t_k(\theta_\ell|\mathbf{z})$, i.e. by propagating each i-th particle $\theta^t_{k,i}$ according to (1):

$$\theta^{t+1}_{k\to\ell,i} = \theta^t_{k,i} \pm z_{k,\ell} + e^{t+1}_{k\to\ell,i}, \tag{8}$$

with $e_{k,\ell,i}\sim\mathcal{N}(0,\sigma^2_{k,\ell})$. Weights are modified with respect to [3]; they are indeed calculated according to the selected importance density as $\gamma^{t+1}_{k\to\ell,i} = m^t_{k\to\ell}(\theta^{t+1}_{k\to\ell,i})/\hat p^t_k(\theta^{t+1}_{k\to\ell,i}|\mathbf{z})$ which, from (3) and (5), yields:

$$\gamma^{t+1}_{k\to\ell,i} \propto \frac{\sum_{n=1}^{\tilde N_p}\psi_{k,\ell}(\tilde\theta^t_{k,n},\theta^{t+1}_{k\to\ell,i})\tilde\omega^t_{k,n}/\tilde\gamma^t_{\ell\to k,n}}{\hat p^t_k(\theta^{t+1}_{k\to\ell,i}|\mathbf{z})}, \tag{9}$$

where $\hat p^t_k(\theta_\ell|\mathbf{z}) \propto \sum_{m=1}^{N_p} p(\theta_\ell|z_{k,\ell},\theta^t_{k,m})\omega^t_{k,m}$ and $p(\theta_\ell|z_{k,\ell},\theta^t_{k,m})$ is, from (1), Gaussian with mean $\theta^t_{k,m}\pm z_{k,\ell}$ and variance $\sigma^2_{k,\ell}$.

[1] A rectangular window or other smoothing functions can be adopted. The width and the shape of the function $w(\cdot)$ should be chosen so as to optimize the accuracy and the convergence rate of the localization algorithm.

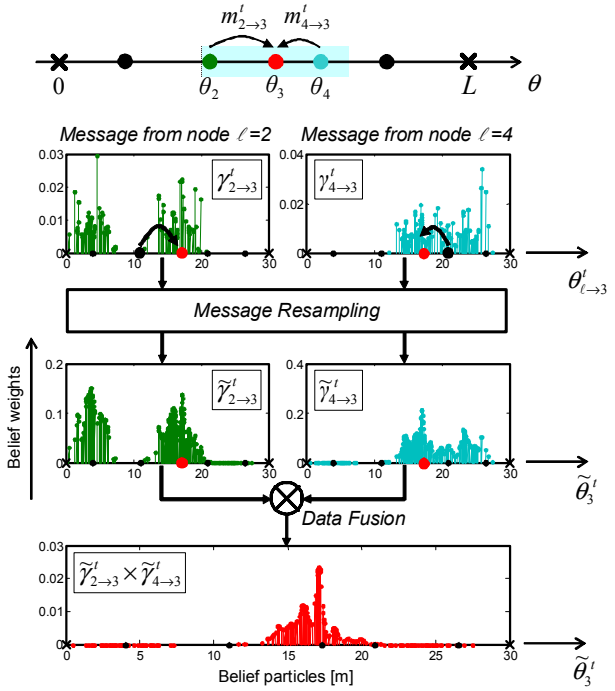

Fig. 2. Example of particle-based message fusion: at iteration t, node 3 resamples the messages from neighbors 2 and 4 over a common set of particles in order to compute the product of the two beliefs (data fusion).

V. Lower Bounds for Stochastic Networks

For a given configuration of nodes $\boldsymbol\theta$, the deterministic (or local) CRB, denoted by $\mathbf{C}_{\mathrm{CRB}}(\boldsymbol\theta)$, provides a lower bound on the covariance matrix $\mathbf{C}(\boldsymbol\theta)$ of any unbiased estimate $\hat{\boldsymbol\theta}(\mathbf{z})$ drawn from $\mathbf{z}$. It is defined as [12]:

$$\mathbf{C}(\boldsymbol\theta) \triangleq \mathrm{E}_{\mathbf{z}}[(\hat{\boldsymbol\theta}(\mathbf{z})-\boldsymbol\theta)(\hat{\boldsymbol\theta}(\mathbf{z})-\boldsymbol\theta)^T] \geq \mathbf{F}^{-1}(\boldsymbol\theta) \triangleq \mathbf{C}_{\mathrm{CRB}}(\boldsymbol\theta). \tag{10}$$

The term $\mathbf{F}(\boldsymbol\theta)$ denotes the Fisher information matrix (FIM) that, for the specific 1D localization problem, is given by [9]:

$$[\mathbf{F}(\boldsymbol\theta)]_{k,\ell} = \begin{cases} 2\sum_{i=1,i\neq k}^{N}\frac{I_{\Omega_k}(i)}{\sigma^2_{k,i}}, & \text{if } k=\ell \\ -2\frac{I_{\Omega_k}(\ell)}{\sigma^2_{k,\ell}}, & \text{if } k\neq\ell \end{cases} \tag{11}$$

where $I_{\Omega_k}(\ell)$ is the visibility function defined as $I_{\Omega_k}(\ell)=1$ if $\ell\in\Omega_k$, or $I_{\Omega_k}(\ell)=0$ if not.

The deterministic CRB depends on the actual value of $\boldsymbol\theta$, i.e. on a specific network deployment. When considering random deployments, on the other hand, it is of interest to evaluate a Bayesian (or global) bound that should be valid for any value of $\boldsymbol\theta$ drawn from $p(\boldsymbol\theta)$. The Bayesian CRB (BCRB) [12], denoted by $\mathbf{C}_{\mathrm{BCRB}}$, belongs to this class of bounds. It is a lower bound for the covariance matrix $\mathbf{C}$ of any Bayesian estimate $\hat{\boldsymbol\theta}(\mathbf{z})$, based on the measurements $\mathbf{z}$ *and* the prior information $p(\boldsymbol\theta)$:

$$\mathbf{C} \triangleq \mathrm{E}_{\boldsymbol\theta}[\mathbf{C}(\boldsymbol\theta)] \geq \mathbf{F}_{\mathrm{B}}^{-1} \triangleq \mathbf{C}_{\mathrm{BCRB}}, \tag{12}$$

with $\mathbf{F}_{\mathrm{B}}$ denoting the Bayesian Information Matrix (BIM):

$$\mathbf{F}_{\mathrm{B}} = \underbrace{\mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{F}(\boldsymbol{\theta})\right]}_{\mathbf{F}_{\mathrm{D}}} + \underbrace{\mathrm{E}_{\boldsymbol{\theta}}\left[\frac{\partial \ln p\left(\boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}} \frac{\partial^{T} \ln p\left(\boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}}\right]}_{\mathbf{F}_{\mathrm{P}}}. \tag{13}$$

The matrices $\mathbf{F}_{\mathrm{D}}$ and $\mathbf{F}_{\mathrm{P}}$ represent the information obtained from, respectively, the data and the a-prior pdf. We recall that in case of uniform prior (with finite support) the BCRB is not defined as the weak unbiasedness condition [12] is not fulfilled (the prior is not twice differentiable); still, the BCRB holds for $N \rightarrow \infty$ and/or large signal-to-noise ratio (SNR) since the measurements become more informative and the prior looses importance. Under these asymptotic conditions (high SNR bound), it is $\mathbf{F}_{\mathrm{B}} \approx \mathbf{F}_{\mathrm{D}}$ and the BCRB reduces to the inverse of the average FIM, i.e. to the modified bound:

$$\mathbf{C}_{\mathrm{MCRB}} \triangleq \left(\mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{F}(\boldsymbol{\theta})\right]\right)^{-1} = \mathbf{F}_{\mathrm{D}}^{-1}. \tag{14}$$

If we consider a non-Bayesian estimate, which does not exploit any prior information and treats $\boldsymbol{\theta}$ as a deterministic parameter, performances are expected to be worse than for the Bayesian estimate. In this case, as shown in Sec. VI, a tighter bound for the error covariance $\mathbf{C}$ is given by the average of the local CRBs with respect to $\boldsymbol{\theta}$, here called the average CRB:

$$\mathbf{C}_{\mathrm{ACRB}} \triangleq \mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{C}_{\mathrm{CRB}}(\boldsymbol{\theta})\right] = \mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{F}(\boldsymbol{\theta})^{-1}\right]. \tag{15}$$

From the Jensen's inequality, it is:

$$\left(\mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{F}(\boldsymbol{\theta})\right] + \mathbf{F}_{P}\right)^{-1} \leq \left(\mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{F}(\boldsymbol{\theta})\right]\right)^{-1} \leq \mathrm{E}_{\boldsymbol{\theta}}\left[\mathbf{F}(\boldsymbol{\theta})^{-1}\right], \tag{16}$$

and thus the following relationship holds:

$$\mathbf{C}_{\mathrm{BCRB}} \leq \mathbf{C}_{\mathrm{MCRB}} \leq \mathbf{C}_{\mathrm{ACRB}}. \tag{17}$$

In the next section all the above CRB bounds will be analyzed for the considered localization scenario.

VI. Simulation Results

We consider the deployment in Fig. 1 with $N_r = 2$ ANs at 0m and $L = 30$m, and $N_u = 5$ UNs. The coverage radius for cooperation is $r = 8$m. In practical systems prior information about node positions can be obtained during the installation phase, based on some knowledge of the environment configuration or the wireless network structure (e.g., UNs are known to be clustered around a number of ANs); the prior pdf $p(\boldsymbol{\theta})$ should be defined so as to account for this information. For simulation here we consider a simplified scenario where UN positions are independent random variables distributed over the interval $[0, L]$ according to the prior pdfs $\{p(\theta_k)\}_{k=1}^{5}$ shown in Fig. 3 (top figure): $p(\theta_k)$ is a raised-cosine of width 2Δ centered around the position $k\Delta$, where $\Delta = L/(N_u + 1) = 5$m is the mean inter-node distance. Measurements are generated according to (1) with $\sigma_{k,\ell} = \sigma \approx 0.2\mathrm{m}^2$.

[2]For TOA-based range measurements, this value of range accuracy could be reached theoretically using a UWB technology with effective bandwidth $B = 500$MHz, signal to noise ratio $\mathrm{SNR} = -10$dB: $\sigma = c/(2\sqrt{2}\pi\sqrt{\mathrm{SNR}B}) \approx 0.2$m, where c is the speed of light [6].

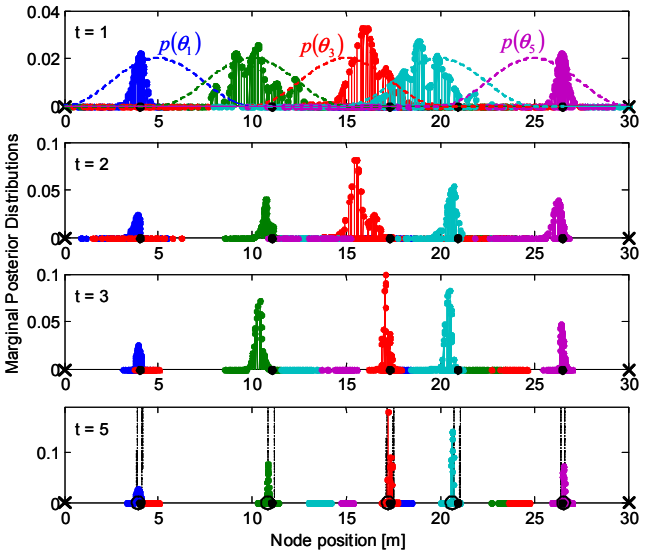

Fig. 3. Marginal posterior pdf estimated by nodes at iterations $t = 1, 2, 3, 5$. ANs are represented with '×' and UNs with '•'.

An example of node beliefs calculated by the BP procedure is shown in Fig 3. The algorithm is implemented as in Sec. IV using a triangular window $w(\theta)$ of support $L_w = 1$m and $N_p = 100$ particles. Fig. 3 shows the marginal posterior distributions approximated by nodes at iterations $t = 1, 2, 3, 5$ and the estimated positions (empty circles) at the last iteration. Looking at the evolution of the beliefs over the iterations, it can be seen that the accuracy and the convergence rate are worse for central nodes which are far away from ANs. The uncertainty interval evaluated according to the local CRB is indicated by dotted lines. After 5 iterations the root mean square error (RMSE) for the BP algorithm, averaged over the nodes, is RMSE ≈ 0.16m, which is approximately equal to the average local CRB calculated from (10): $\mathrm{tr}[\mathbf{C}_{\mathrm{CRB}}(\boldsymbol{\theta})]/N_u \approx 0.16$m.

We compare now the performance of the PF-BP algorithm with the lower bounds discussed in Sec. V. The network model is simulated according to the prior distributions in Fig. 3 (top figure) with $L = 55$m, $N_u = 10$, and mean inter-node distance $\Delta = 5$m. Network configurations for which the localization problem is not solvable (i.e., the FIM is singular) are not considered. The plot on the left in Fig. 4 shows the RMSE of the location estimate vs the inverse range error $1/\bar{\sigma}$ computed over 100 realizations of the network layout, each of them characterized by 100 sets of measurements. The range error is normalized with respect to the coverage radius $r = 10$m: $\bar{\sigma} = \sigma/r$. The accuracy of the PF-BP algorithm is shown to approach the BCRB while it is upper bounded by the ACRB. As expected, the BCRB is a more realistic bound with respect to the ACRB for Bayesian estimation as it accounts for the information provided by the prior pdf $p(\boldsymbol{\theta})$. For small range error σ (high SNR bound), we observe that the BCRB tends to the MCRB as the information drawn from measurements ($\mathbf{F}_{\mathrm{D}}$) becomes dominant with respect to the prior information ($\mathbf{F}_{\mathrm{P}}$). For $\sigma = 1.8$m, the right plot in Fig. 4 shows the localization accuracy as a function of the node index

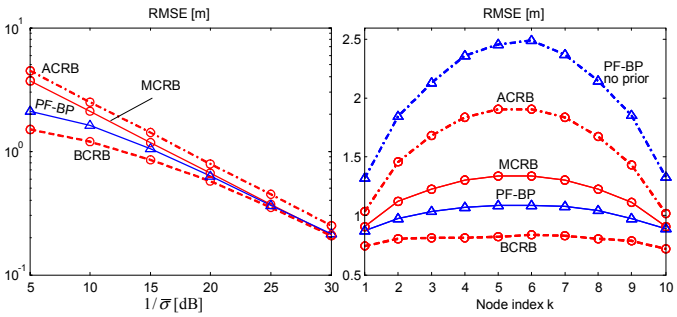

Fig. 4. Accuracy of PF-BP compared to the lower bounds vs. (left) the inverse normalized range accuracy and (right) the nodes along the 1D array.

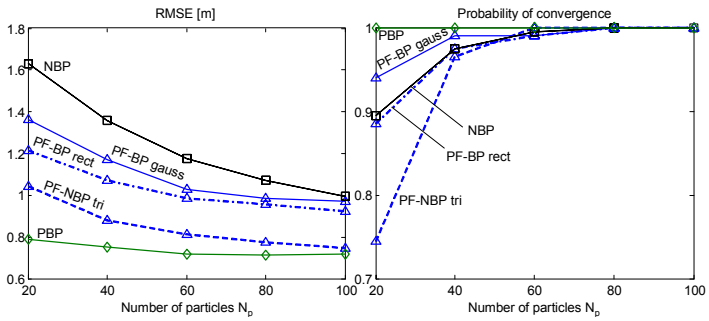

Fig. 5. Performance of BP algorithms vs the number of particles: (left) accuracy and (right) probability of convergence.

k. Since nodes are randomly scattered around a regular grid with inter-node interval Δ, the performances along the array are similar to those observed for regular networks (see [9]), with accuracy that decreases towards the centre of the array. In addition to PF-BP with known $p(\theta_k)$, the figure shows also the performance of the algorithm when prior information is not available, i.e. for $p(\theta_k)$ uniform in $[0, L]$. In this case, the BCRB is too low, as it assumes prior information at the estimator. On the other hand, the ACRB is a tighter bound as it is obtained as the average of the performance of localization applied separately on each network deployment without taking into account the statistics of the deployments.

Finally, Fig. 5 presents a comparison between PF-BP, NBP [4], and PBP [11], for different values of the number of particles N_p. The network deployment here is fixed and is the same as in Fig. 3, with range error $\sigma = 1.8$m. To perform a fair analysis, NBP is implemented choosing, for the product of incoming messages, a collection of $|\Omega_k| \cdot N_p$ particles for each node. The windowing function $w(\theta)$ for PF-BP has effective width $\sigma_w = \int \theta^2 w(\theta) d\theta$ equal to the standard deviation of the Gaussian smoothing kernel of NBP. We consider three types of function $w(\theta)$: triangular (tri), Gaussian (gauss), and rectangular (rect). To ease the interpretation of results we keep σ_w constant for the different values of N_p and we set $\sigma_w = \sigma$; a more efficient implementation can be obtained by optimizing σ_w for each value of N_p (i.e., decreasing σ_w for increasing N_p). PBP is implemented according to the procedure in [11, Sec. 4], using as sampling distributions the current belief estimates. Fig. 5-left shows the RMSE averaged over all nodes at the 5th BP iteration, while Fig. 5-right gives the probability of convergence of the BP algorithm. Results are obtained by averaging over 200 realizations of measurements. For $N_p \geq 40$ the triangular kernel is the best choice in terms of estimate error and convergence; the rectangual one slightly reduces the performance but it has the advantage of having lower complexity (no multiplications are needed in (6)). PF-BP slightly outperforms NBP in terms of RMSE. PBP is more accurate (as it does not perform message smoothing); however the transmission of updated particle sets to neighbors requires one additional broadcast at each iteration, and this operation is time-consuming especially for localization in dynamic scenarios. Notice that the performance gain of PF-BP over NBP is given at the price of a higher computational complexity: PF-BP has an efficient message fusion procedure, but the complexity required for the evaluation of weights in (9) is quadratic in N_p while for NBP is linear.

VII. Conclusions

We investigated Bayesian cooperative localization in wireless networks where nodes are randomly deployed over a finite space according to a (known) prior distribution. We proposed a modified version of the distributed NBP algorithm using particle filtering for belief representation. In order to assess the algorithm performance, we derived fundamental performance bounds based on the CRB for the considered stochastic networks. The BCRB turned out to be a realistic bound for the PF-BP method when prior information is available on node positions, otherwise the ACRB should be considered. A further development of the work will be the extension to 2D random networks and a more in-depth investigation of the performance of cooperative localization and its fundamental bounds for realistic stochastic network models.

References

[1] N. Patwari et al., "Locating the nodes," *IEEE Signal Proc. Mag.*, Vol. 22, No. 4, pp. 54-69, July 2005.
[2] H. Wymeersch, J. Lien, M.Z. Win, "Cooperative Localization in Wireless Networks," *Proceedings of the IEEE*, Vol. 97, pp. 427-450, Feb. 2009.
[3] E.B. Sudderth, A.T. Ihler, W.T. Freeman, A.S. Willsky, "Nonparametric Belief Propagation", in *Proc. IEEE CVPR 2003*, Vol. I, pp. 605-612.
[4] A.T. Ihler, J.W. Fisher, R.L. Moses, A.S. Willsky, "Nonparametric belief propagation for self-localization of sensor networks," *IEEE J. Sel. Areas Commun.*, Vol 23, No. 4, pp. 809-819, April 2005.
[5] M. Nicoli, C. Morelli, V. Rampa, "A jump Markov particle filter for localization of moving terminals in dense multipath indoor scenarios," *IEEE Trans. Signal Processing*, Vol. 56, N. 8, pp. 3801-3809, Aug. 2008.
[6] S. Gezici et al., "Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks," *IEEE Signal Processing Mag.*, vol. 22, No. 4, pp. 70-84, July 2005.
[7] F. Gustafsson and F. Gunnarsson, "Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements," *IEEE Signal Processing Mag.*, Vol. 22, No. 4, pp. 41-53, July 2005.
[8] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, M. Z. Win, "Ranging With Ultrawide Bandwidth Signals in Multipath Environments," *Proceedings of the IEEE*, Vol. 97, No. 2, pp. 404-426, Feb. 2009.
[9] M.Nicoli and D.Fontanella, "Fundamental Performance Limits of TOA-based Cooperative Localization," in *Proc. IEEE ICC 2009*.
[10] M.S. Arulampalam, S. Maskell, N. Gordon, T. Clapp, "A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking", *IEEE Trans. Signal Processing*, Vol. 50, No. 2, pp. 174-188, Feb. 2002.
[11] A. Ihler and D. McAllester, "Particle belief propagation", Proc. AISTATS 2009, Vol. 5, 2009, pp. 256-263.
[12] H. L. Van Trees, *Optimum Array Processing*, Wiley, 2002.