

Sigma Point Belief Propagation

Florian Meyer, *Student Member, IEEE*, Ondrej Hlinka, *Member, IEEE*, and Franz Hlawatsch, *Fellow, IEEE*

Abstract—The sigma point (SP) filter, also known as unscented Kalman filter, is an attractive alternative to the extended Kalman filter and the particle filter. Here, we extend the SP filter to nonsequential Bayesian inference corresponding to loopy factor graphs. We propose *sigma point belief propagation* (SPBP) as a low-complexity approximation of the belief propagation (BP) message passing scheme. SPBP achieves approximate marginalizations of posterior distributions corresponding to (generally) loopy factor graphs. It is well suited for decentralized inference because of its low communication requirements. For a decentralized, dynamic sensor localization problem, we demonstrate that SPBP can outperform nonparametric (particle-based) BP while requiring significantly less computations and communications.

Index Terms—Sigma points, belief propagation, factor graph, unscented transformation, cooperative localization.

I. INTRODUCTION

The sigma point (SP) filter, also known as unscented Kalman filter, is a sequential Bayesian estimator for nonlinear systems that outperforms the extended Kalman filter while being typically less complex than the particle filter [1], [2]. Sequential Bayesian estimation corresponds to a “sequential” factor structure of the joint posterior probability density function (pdf). For more general—possibly loopy—factor structures, the belief propagation (BP) message passing scheme can be used to (approximately) perform the marginalizations required for Bayesian inference [3]. Gaussian BP (GBP) [4] and nonparametric BP (NBP) [5] are reduced-complexity approximations of BP that extend the Kalman filter [6] and the particle filter [7], respectively to general factor structures. GBP assumes a linear, Gaussian system and uses Gaussian message representations, whereas NBP is suited to nonlinear, non-Gaussian systems due to its use of particle representations.

Here, we propose the *sigma point BP* (SPBP) message passing scheme as a new low-complexity approximation of BP for general nonlinear systems. SPBP extends the SP filter to general factor structures. We demonstrate that the performance of SPBP can be similar to or even better than that of NBP, at a far lower complexity. SPBP is well suited to certain distributed (decentralized) inference problems in wireless sensor networks because of its low communication requirements—only a mean vector and a covariance matrix are communicated between neighboring sensors. We note that SPBP is different from the algorithm proposed in [8], which uses SPs to approximate BP in Bayesian networks and an information fusion technique to multiply messages. In contrast, SPBP is based on a factor graph and a reformulation of BP in higher-dimensional spaces. In addition to the advantages of factor graphs discussed in [3],

F. Meyer, O. Hlinka, and F. Hlawatsch are with the Institute of Telecommunications, Vienna University of Technology, A-1040 Vienna, Austria (email: {florian.meyer, ondrej.hlinka, franz.hlawatsch}@tuwien.ac.at). This work was supported by FWF grant S10603 and by the Newcom# Network of Excellence in Wireless Communications of the European Commission.

this simplifies a decentralized implementation.

This letter is organized as follows. Some basics of SPs are reviewed in Section II. In Section III, the system model is described and BP is reviewed. The SPBP scheme is developed in Section IV, and its computation and communication requirements are discussed in Section V. Section VI presents simulation results for a decentralized, cooperative, dynamic self-localization problem.

II. SIGMA POINT BASICS

Consider a general (non-Gaussian) random vector $\mathbf{x} \in \mathbb{R}^J$ whose mean $\boldsymbol{\mu}_x = E\{\mathbf{x}\}$ and covariance matrix $\mathbf{C}_x = E\{(\mathbf{x} - \boldsymbol{\mu}_x)(\mathbf{x} - \boldsymbol{\mu}_x)^T\}$ are known, and a transformed random vector $\mathbf{y} = H(\mathbf{x})$, where $H(\cdot)$ is a generally nonlinear function. SPs $\{\mathbf{x}^{(j)}\}_{j=0}^{2J}$ and corresponding weights $\{w_m^{(j)}\}_{j=0}^{2J}$ and $\{w_c^{(j)}\}_{j=0}^{2J}$ are chosen such that the weighted sample mean $\tilde{\boldsymbol{\mu}}_x = \sum_{j=0}^{2J} w_m^{(j)} \mathbf{x}^{(j)}$ and weighted sample covariance matrix $\tilde{\mathbf{C}}_x = \sum_{j=0}^{2J} w_c^{(j)} (\mathbf{x}^{(j)} - \tilde{\boldsymbol{\mu}}_x)(\mathbf{x}^{(j)} - \tilde{\boldsymbol{\mu}}_x)^T$ are exactly equal to $\boldsymbol{\mu}_x$ and \mathbf{C}_x , respectively. Closed-form expressions of the SPs and weights are provided in [2]. The spread of the SPs around the mean $\boldsymbol{\mu}_x$ can be adjusted via tuning parameters, whose choice depends on the dimension J of \mathbf{x} [1], [2]. Next, each SP is propagated through $H(\cdot)$, resulting in $\mathbf{y}^{(j)} = H(\mathbf{x}^{(j)})$, $j \in \{0, \dots, 2J\}$ (“unscented transformation”). The set $\{(\mathbf{x}^{(j)}, \mathbf{y}^{(j)}, w_m^{(j)}, w_c^{(j)})\}_{j=0}^{2J}$ then represents the joint second-order statistics of \mathbf{x} and \mathbf{y} in an approximate manner. In particular, $\boldsymbol{\mu}_y$, \mathbf{C}_y , and $\mathbf{C}_{xy} = E\{(\mathbf{x} - \boldsymbol{\mu}_x)(\mathbf{y} - \boldsymbol{\mu}_y)^T\}$ are approximated by

$$\tilde{\boldsymbol{\mu}}_y = \sum_{j=0}^{2J} w_m^{(j)} \mathbf{y}^{(j)} \quad (1)$$

$$\tilde{\mathbf{C}}_y = \sum_{j=0}^{2J} w_c^{(j)} (\mathbf{y}^{(j)} - \tilde{\boldsymbol{\mu}}_y)(\mathbf{y}^{(j)} - \tilde{\boldsymbol{\mu}}_y)^T \quad (2)$$

$$\tilde{\mathbf{C}}_{xy} = \sum_{j=0}^{2J} w_c^{(j)} (\mathbf{x}^{(j)} - \tilde{\boldsymbol{\mu}}_x)(\mathbf{y}^{(j)} - \tilde{\boldsymbol{\mu}}_y)^T. \quad (3)$$

It has been shown in [1] and [2] that these approximations are at least as accurate as those resulting from a linearization (first-order Taylor series approximation) of $H(\cdot)$. Note also that the number $2J + 1$ of SPs grows linearly with the dimension of \mathbf{x} and is typically much smaller than the number of random samples in a particle representation.

Next, we consider the use of SPs for Bayesian estimation of a random vector \mathbf{x} from an observed vector

$$\mathbf{z} = \mathbf{y} + \mathbf{n}, \quad \text{with } \mathbf{y} = H(\mathbf{x}).$$

Here, the noise \mathbf{n} is statistically independent of \mathbf{x} and generally non-Gaussian, with zero mean and known covariance matrix

\mathbf{C}_n . Bayesian estimation relies on the posterior pdf

$$f(\mathbf{x}|\mathbf{z}) \propto f(\mathbf{z}|\mathbf{x})f(\mathbf{x}), \quad (4)$$

where $f(\mathbf{z}|\mathbf{x})$ is the likelihood function and $f(\mathbf{x})$ is the prior pdf. Direct calculation of (4) is usually infeasible. An important exception is the case where \mathbf{x} and \mathbf{n} are Gaussian random vectors and $H(\mathbf{x}) = \mathbf{H}\mathbf{x}$ with some known matrix \mathbf{H} . Then $f(\mathbf{x}|\mathbf{z})$ is also Gaussian, and the posterior mean $\boldsymbol{\mu}_{\mathbf{x}|\mathbf{z}}$ and posterior covariance matrix $\mathbf{C}_{\mathbf{x}|\mathbf{z}}$ can be calculated as

$$\boldsymbol{\mu}_{\mathbf{x}|\mathbf{z}} = \boldsymbol{\mu}_{\mathbf{x}} + \mathbf{K}(\mathbf{z} - \boldsymbol{\mu}_{\mathbf{y}}), \quad \mathbf{C}_{\mathbf{x}|\mathbf{z}} = \mathbf{C}_{\mathbf{x}} - \mathbf{K}(\mathbf{C}_{\mathbf{y}} + \mathbf{C}_{\mathbf{n}})\mathbf{K}^T, \quad (5)$$

where

$$\boldsymbol{\mu}_{\mathbf{y}} = \mathbf{H}\boldsymbol{\mu}_{\mathbf{x}}, \quad \mathbf{C}_{\mathbf{y}} = \mathbf{H}\mathbf{C}_{\mathbf{x}}\mathbf{H}^T \quad (6)$$

and

$$\mathbf{K} = \mathbf{C}_{\mathbf{x}\mathbf{y}}(\mathbf{C}_{\mathbf{y}} + \mathbf{C}_{\mathbf{n}})^{-1}, \quad \text{with } \mathbf{C}_{\mathbf{x}\mathbf{y}} = \mathbf{C}_{\mathbf{x}}\mathbf{H}^T. \quad (7)$$

These expressions are used in the measurement update step of the Kalman filter [6]. The minimum mean-square error estimate of \mathbf{x} is given by $\boldsymbol{\mu}_{\mathbf{x}|\mathbf{z}}$, and a characterization of the accuracy of estimation by $\mathbf{C}_{\mathbf{x}|\mathbf{z}}$.

In the general case of nonlinear $H(\cdot)$, the basic approximation underlying the extended Kalman filter [6] is obtained by using (essentially) (5)–(7) with \mathbf{H} being the Jacobian matrix resulting from a linearization of $H(\cdot)$. A more accurate alternative is to approximate $\boldsymbol{\mu}_{\mathbf{x}|\mathbf{z}}$ and $\mathbf{C}_{\mathbf{x}|\mathbf{z}}$ by means of SPs. For this, we use (5) and the first equation in (7), with $\boldsymbol{\mu}_{\mathbf{y}}$, $\mathbf{C}_{\mathbf{y}}$, and $\mathbf{C}_{\mathbf{x}\mathbf{y}}$ replaced by the SP approximations $\tilde{\boldsymbol{\mu}}_{\mathbf{y}}$, $\tilde{\mathbf{C}}_{\mathbf{y}}$, and $\tilde{\mathbf{C}}_{\mathbf{x}\mathbf{y}}$ in (1)–(3). This gives

$$\tilde{\boldsymbol{\mu}}_{\mathbf{x}|\mathbf{z}} = \boldsymbol{\mu}_{\mathbf{x}} + \tilde{\mathbf{K}}(\mathbf{z} - \tilde{\boldsymbol{\mu}}_{\mathbf{y}}), \quad \tilde{\mathbf{C}}_{\mathbf{x}|\mathbf{z}} = \mathbf{C}_{\mathbf{x}} - \tilde{\mathbf{K}}(\tilde{\mathbf{C}}_{\mathbf{y}} + \mathbf{C}_{\mathbf{n}})\tilde{\mathbf{K}}^T, \quad (8)$$

with $\tilde{\mathbf{K}} = \tilde{\mathbf{C}}_{\mathbf{x}\mathbf{y}}(\tilde{\mathbf{C}}_{\mathbf{y}} + \mathbf{C}_{\mathbf{n}})^{-1}$. We thus obtain the following approximate SP implementation of (4).

Step 1: SPs and weights $\{(\mathbf{x}^{(j)}, w_m^{(j)}, w_c^{(j)})\}_{j=0}^{2J}$ are calculated from $\boldsymbol{\mu}_{\mathbf{x}}$ and $\mathbf{C}_{\mathbf{x}}$ [2].

Step 2: The transformed SPs $\mathbf{y}^{(j)} = H(\mathbf{x}^{(j)})$, $j \in \{0, \dots, 2J\}$ are calculated.

Step 3: From $\{(\mathbf{x}^{(j)}, \mathbf{y}^{(j)}, w_m^{(j)}, w_c^{(j)})\}_{j=0}^{2J}$, the means and covariances $\tilde{\boldsymbol{\mu}}_{\mathbf{y}}$, $\tilde{\mathbf{C}}_{\mathbf{y}}$, and $\tilde{\mathbf{C}}_{\mathbf{x}\mathbf{y}}$ in (1)–(3) and, in turn, $\tilde{\boldsymbol{\mu}}_{\mathbf{x}|\mathbf{z}}$ and $\tilde{\mathbf{C}}_{\mathbf{x}|\mathbf{z}}$ in (8) are calculated.

This algorithm can be extended to nonadditive noise [2].

III. BELIEF PROPAGATION

We will next describe our system model and review BP. We consider K states \mathbf{x}_k , $k \in \{1, \dots, K\}$ and observations $\mathbf{z}_{k,l}$ that involve pairs of states $\mathbf{x}_k, \mathbf{x}_l$ according to

$$\mathbf{z}_{k,l} = G(\mathbf{x}_k, \mathbf{x}_l) + \mathbf{n}_{k,l}, \quad (k, l) \in \mathcal{E}. \quad (9)$$

Here, the set $\mathcal{E} \subseteq \{1, \dots, K\}^2$ is symmetric, i.e., $(k, l) \in \mathcal{E}$ implies $(l, k) \in \mathcal{E}$; $G(\cdot, \cdot)$ is a generally nonlinear symmetric function, i.e., $G(\mathbf{x}_k, \mathbf{x}_l) = G(\mathbf{x}_l, \mathbf{x}_k)$; and $\mathbf{n}_{k,l} = \mathbf{n}_{l,k}$ is zero-mean with known covariance matrix $\mathbf{C}_{\mathbf{n}_{k,l}}$. Note that $\mathbf{z}_{k,l} = \mathbf{z}_{l,k}$. (Methods for enforcing symmetry and a modified BP scheme that does not require it will be discussed presently.) We assume that $\mathbf{n}_{k,l}$ is independent of all \mathbf{x}_k , that $\mathbf{n}_{k,l}$ and $\mathbf{n}_{k',l'}$ are independent unless $(k, l) = (k', l')$ or $(k, l) = (l', k')$,

and that all \mathbf{x}_k are *a priori* independent. The BP algorithm—as well as the SPBP algorithm presented in Section IV—can be easily extended to more general system models, i.e., it is not limited to additive noise and “pairwise” observations.

In what follows, $\mathbf{x} \triangleq (\mathbf{x}_1^T \cdots \mathbf{x}_K^T)^T$; similarly, \mathbf{z} is defined by stacking all $\mathbf{z}_{k,l}$ in arbitrary order. Because of (4) and our assumptions, the joint posterior pdf $f(\mathbf{x}|\mathbf{z})$ factorizes as

$$f(\mathbf{x}|\mathbf{z}) \propto \prod_{k=1}^K f(\mathbf{x}_k) \prod_{\substack{(k',l) \in \mathcal{E} \\ k' > l}} f(\mathbf{z}_{k',l}|\mathbf{x}_{k'}, \mathbf{x}_l). \quad (10)$$

Bayesian estimation of the states \mathbf{x}_k relies on the marginal posterior pdfs $f(\mathbf{x}_k|\mathbf{z})$. Whereas direct marginalization of $f(\mathbf{x}|\mathbf{z})$ in (10) is usually infeasible, approximate marginal posterior pdfs (“beliefs”) $b(\mathbf{x}_k) \approx f(\mathbf{x}_k|\mathbf{z})$ can be computed by executing iterative BP message passing on the factor graph corresponding to (10) [3]. Let the “neighbor” set $\mathcal{N}_k = \{l_1, l_2, \dots, l_{|\mathcal{N}_k|}\}$ of variable node $k \in \{1, \dots, K\}$ comprise all $l \in \{1, \dots, K\} \setminus \{k\}$ such that $(k, l) \in \mathcal{E}$. Then, at iteration $p \geq 1$, the belief of variable node k is obtained as [3]

$$b^{(p)}(\mathbf{x}_k) \propto f(\mathbf{x}_k) \prod_{l \in \mathcal{N}_k} m_{l \rightarrow k}^{(p)}(\mathbf{x}_k), \quad (11)$$

where

$$m_{l \rightarrow k}^{(p)}(\mathbf{x}_k) = \int f(\mathbf{z}_{k,l}|\mathbf{x}_k, \mathbf{x}_l) n_{l \rightarrow k}^{(p-1)}(\mathbf{x}_l) d\mathbf{x}_l, \quad l \in \mathcal{N}_k, \quad (12)$$

with

$$n_{l \rightarrow k}^{(p)}(\mathbf{x}_l) = f(\mathbf{x}_l) \prod_{k' \in \mathcal{N}_l \setminus \{k\}} m_{k' \rightarrow l}^{(p)}(\mathbf{x}_l), \quad l \in \mathcal{N}_k. \quad (13)$$

This recursion is initialized by setting $n_{l \rightarrow k}^{(0)}(\mathbf{x}_l)$ equal to the prior pdf of \mathbf{x}_l .

In a decentralized scenario where the variable nodes in the factor graph are simultaneously sensor nodes, the symmetry condition $\mathbf{z}_{k,l} = \mathbf{z}_{l,k}$ is usually not satisfied. It can be enforced by considering the averaged measurement $(\mathbf{z}_{k,l} + \mathbf{z}_{l,k})/2$ or the stacked measurement $(\mathbf{z}_{k,l}^T \ \mathbf{z}_{l,k}^T)^T$; this requires communication between the sensor nodes k and l . Alternatively, the condition $\mathbf{z}_{k,l} = \mathbf{z}_{l,k}$ can be avoided by using the SPAWN message passing scheme [9]. SPAWN differs from the standard BP operations (11)–(13) in that $n_{l \rightarrow k}^{(p)}(\mathbf{x}_l) = b^{(p)}(\mathbf{x}_l)$ for all $l \in \mathcal{N}_k$. Consequently, (12) is replaced by

$$m_{l \rightarrow k}^{(p)}(\mathbf{x}_k) = \int f(\mathbf{z}_{k,l}|\mathbf{x}_k, \mathbf{x}_l) b^{(p-1)}(\mathbf{x}_l) d\mathbf{x}_l, \quad l \in \mathcal{N}_k,$$

and (13) need not be calculated.

IV. SIGMA POINT BELIEF PROPAGATION

The proposed SPBP algorithm performs a low-complexity, approximate calculation of $b^{(p)}(\mathbf{x}_k)$ in (11) based on SPs. The key idea is to reformulate the BP operations in higher-dimensional spaces that are defined by the “composite” vectors $\bar{\mathbf{x}}_k \triangleq (\mathbf{x}_k^T \ \mathbf{x}_{l_1}^T \ \mathbf{x}_{l_2}^T \cdots \mathbf{x}_{l_{|\mathcal{N}_k|}}^T)^T$ (\mathbf{x}_k and its neighbor states) and $\bar{\mathbf{z}}_k \triangleq (\mathbf{z}_{k,l_1}^T \ \mathbf{z}_{k,l_2}^T \cdots \mathbf{z}_{k,l_{|\mathcal{N}_k|}}^T)^T$ (all observations involving \mathbf{x}_k). The dimension of $\bar{\mathbf{x}}_k$ is $\bar{J}_k \triangleq J_k + \sum_{l=1}^{|\mathcal{N}_k|} J_l$, where J_k denotes the dimension of \mathbf{x}_k . Let $\bar{\mathbf{x}}_k^{\sim k}$ (resp. $\bar{\mathbf{x}}_k^{\sim k,l}$) denote $\bar{\mathbf{x}}_k$ with the subvector \mathbf{x}_k (resp. the subvectors \mathbf{x}_k and \mathbf{x}_l) removed,

and let $\bar{\mathbf{z}}_k^{\sim l}$ denote $\bar{\mathbf{z}}_k$ with the subvector $\mathbf{z}_{k,l}$ removed. By inserting (12) into (11) and (13), we obtain

$$b^{(p)}(\mathbf{x}_k) \propto \int f(\bar{\mathbf{z}}_k|\bar{\mathbf{x}}_k) f^{(p-1)}(\bar{\mathbf{x}}_k) d\bar{\mathbf{x}}_k^{\sim k} \quad (14)$$

$$n_{l \rightarrow k}^{(p)}(\mathbf{x}_l) = \int f(\bar{\mathbf{z}}_l^{\sim k}|\bar{\mathbf{x}}_l^{\sim k}) f^{(p-1)}(\bar{\mathbf{x}}_l^{\sim k}) d\bar{\mathbf{x}}_l^{\sim l,k}, \quad l \in \mathcal{N}_k, \quad (15)$$

where

$$\begin{aligned} f(\bar{\mathbf{z}}_k|\bar{\mathbf{x}}_k) &= \prod_{l \in \mathcal{N}_k} f(\mathbf{z}_{k,l}|\mathbf{x}_k, \mathbf{x}_l) \\ f^{(p-1)}(\bar{\mathbf{x}}_k) &\propto f(\mathbf{x}_k) \prod_{l \in \mathcal{N}_k} n_{l \rightarrow k}^{(p-1)}(\mathbf{x}_l) \\ f(\bar{\mathbf{z}}_l^{\sim k}|\bar{\mathbf{x}}_l^{\sim k}) &= \prod_{k' \in \mathcal{N}_l \setminus \{k\}} f(\mathbf{z}_{l,k'}|\mathbf{x}_l, \mathbf{x}_{k'}) \\ f^{(p-1)}(\bar{\mathbf{x}}_l^{\sim k}) &\propto f(\mathbf{x}_l) \prod_{k' \in \mathcal{N}_l \setminus \{k\}} n_{k' \rightarrow l}^{(p-1)}(\mathbf{x}_{k'}). \end{aligned}$$

(If SPAWN is used, $f^{(p-1)}(\bar{\mathbf{x}}_k) \propto f(\mathbf{x}_k) \prod_{l \in \mathcal{N}_k} b^{(p-1)}(\mathbf{x}_l)$, and thus $n_{l \rightarrow k}^{(p)}(\mathbf{x}_l)$ need not be calculated.) Note that $f(\bar{\mathbf{z}}_k|\bar{\mathbf{x}}_k)$ and $f^{(p-1)}(\bar{\mathbf{x}}_k)$ are, respectively, the likelihood function and ‘‘iterated prior pdf’’ for the composite observation model

$$\bar{\mathbf{z}}_k = \bar{\mathbf{y}}_k + \bar{\mathbf{n}}_k, \quad \text{with } \bar{\mathbf{y}}_k = H(\bar{\mathbf{x}}_k),$$

where $H(\bar{\mathbf{x}}_k) \triangleq ((G(\mathbf{x}_k, \mathbf{x}_{l_1}))^T \cdots (G(\mathbf{x}_k, \mathbf{x}_{l_{|\mathcal{N}_k|}}))^T)^T$ and $\bar{\mathbf{n}}_k \triangleq (\mathbf{n}_{k,l_1}^T \cdots \mathbf{n}_{k,l_{|\mathcal{N}_k|}}^T)^T$.

To develop an SP-based approximate calculation of $b^{(p)}(\mathbf{x}_k)$, we first note that (14) is a marginalization, i.e.,

$$b^{(p)}(\mathbf{x}_k) = \int b^{(p)}(\bar{\mathbf{x}}_k) d\bar{\mathbf{x}}_k^{\sim k}, \quad (16)$$

with

$$b^{(p)}(\bar{\mathbf{x}}_k) \propto f(\bar{\mathbf{z}}_k|\bar{\mathbf{x}}_k) f^{(p-1)}(\bar{\mathbf{x}}_k). \quad (17)$$

Because the expression (17) of the ‘‘composite belief’’ $b^{(p)}(\bar{\mathbf{x}}_k)$ is analogous to (4), we can obtain an approximate SP representation of $b^{(p)}(\bar{\mathbf{x}}_k)$ in a similar way as we obtained an approximate SP representation of $f(\mathbf{x}|\mathbf{z})$ in Section II. We first define a mean vector and a covariance matrix corresponding to the ‘‘composite prior’’ $f^{(p-1)}(\bar{\mathbf{x}}_k) \propto f(\mathbf{x}_k) \prod_{l \in \mathcal{N}_k} n_{l \rightarrow k}^{(p-1)}(\mathbf{x}_l)$:

$$\boldsymbol{\mu}_{\bar{\mathbf{x}}_k}^{(p-1)} \triangleq \left(\boldsymbol{\mu}_{\mathbf{x}_k}^T \boldsymbol{\mu}_{l_1 \rightarrow k}^{(p-1)T} \boldsymbol{\mu}_{l_2 \rightarrow k}^{(p-1)T} \cdots \boldsymbol{\mu}_{l_{|\mathcal{N}_k|} \rightarrow k}^{(p-1)T} \right)^T \quad (18)$$

$$\mathbf{C}_{\bar{\mathbf{x}}_k}^{(p-1)} \triangleq \text{diag} \left\{ \mathbf{C}_{\mathbf{x}_k}, \mathbf{C}_{l_1 \rightarrow k}^{(p-1)}, \mathbf{C}_{l_2 \rightarrow k}^{(p-1)}, \dots, \mathbf{C}_{l_{|\mathcal{N}_k|} \rightarrow k}^{(p-1)} \right\}. \quad (19)$$

Here, we interpreted $\prod_{l \in \mathcal{N}_k} n_{l \rightarrow k}^{(p-1)}(\mathbf{x}_l)$ as the product of the pdfs of statistically independent random variables (up to a normalization); furthermore, $\boldsymbol{\mu}_{\mathbf{x}_k}$ and $\mathbf{C}_{\mathbf{x}_k}$ are the mean and covariance matrix of the prior $f(\mathbf{x}_k)$; $\boldsymbol{\mu}_{l_i \rightarrow k}^{(p-1)}$ and $\mathbf{C}_{l_i \rightarrow k}^{(p-1)}$ are the mean and covariance matrix of $n_{l_i \rightarrow k}^{(p-1)}(\mathbf{x}_{l_i})$; and $\text{diag}\{\cdot\}$ denotes the block diagonal matrix whose diagonal blocks are the listed matrices. (For SPAWN, $\boldsymbol{\mu}_{l_i \rightarrow k}^{(p-1)}$ and $\mathbf{C}_{l_i \rightarrow k}^{(p-1)}$ are the mean and covariance matrix of $b^{(p-1)}(\mathbf{x}_{l_i})$.) The following steps are now performed for each $k \in \{1, \dots, K\}$ (note that the first three steps are analogous to those in Section II).

Step 1: SPs and weights $\{(\bar{\mathbf{x}}_k^{(j)}, w_m^{(j)}, w_c^{(j)})\}_{j=0}^{2\bar{J}_k}$ corresponding to $f^{(p-1)}(\bar{\mathbf{x}}_k)$ are calculated from $\boldsymbol{\mu}_{\bar{\mathbf{x}}_k}^{(p-1)}$ and $\mathbf{C}_{\bar{\mathbf{x}}_k}^{(p-1)}$ [2].

(Note that the dimension and number of the SPs depend on the number of neighbors $|\mathcal{N}_k|$, and thus the tuning parameters that adjust the spread of the SPs should be adapted to $|\mathcal{N}_k|$.)

Step 2: The transformed SPs $\bar{\mathbf{y}}_k^{(j)} = H(\bar{\mathbf{x}}_k^{(j)})$, $j \in \{0, \dots, 2\bar{J}_k\}$ are calculated.

Step 3: From $\{(\bar{\mathbf{x}}_k^{(j)}, \bar{\mathbf{y}}_k^{(j)}, w_m^{(j)}, w_c^{(j)})\}_{j=0}^{2\bar{J}_k}$, the means and covariances $\tilde{\boldsymbol{\mu}}_{\bar{\mathbf{y}}_k}^{(p)}$, $\tilde{\mathbf{C}}_{\bar{\mathbf{y}}_k}^{(p)}$, and $\tilde{\mathbf{C}}_{\bar{\mathbf{x}}_k \bar{\mathbf{y}}_k}^{(p)}$ are calculated as in (1)–(3). Subsequently, $\tilde{\boldsymbol{\mu}}_{b(\bar{\mathbf{x}}_k)}^{(p)}$ and $\tilde{\mathbf{C}}_{b(\bar{\mathbf{x}}_k)}^{(p)}$ (the SP approximations of the mean and covariance matrix of $b^{(p)}(\bar{\mathbf{x}}_k)$) are calculated as in (8), using $\tilde{\boldsymbol{\mu}}_{\bar{\mathbf{y}}_k}^{(p)}$, $\tilde{\mathbf{C}}_{\bar{\mathbf{y}}_k}^{(p)}$, and $\tilde{\mathbf{C}}_{\bar{\mathbf{x}}_k \bar{\mathbf{y}}_k}^{(p)}$ instead of $\tilde{\boldsymbol{\mu}}_{\mathbf{y}}$, $\tilde{\mathbf{C}}_{\mathbf{y}}$, and $\tilde{\mathbf{C}}_{\mathbf{x}\mathbf{y}}$, respectively.

Step 4: From $\tilde{\boldsymbol{\mu}}_{b(\bar{\mathbf{x}}_k)}^{(p)}$ and $\tilde{\mathbf{C}}_{b(\bar{\mathbf{x}}_k)}^{(p)}$, the elements related to \mathbf{x}_k are extracted (this corresponds to the marginalization (16)). More specifically, the approximate mean $\tilde{\boldsymbol{\mu}}_{b(\mathbf{x}_k)}^{(p)}$ and covariance matrix $\tilde{\mathbf{C}}_{b(\mathbf{x}_k)}^{(p)}$ of the ‘‘marginal belief’’ $b^{(p)}(\mathbf{x}_k)$ are given by the first J_k elements of $\tilde{\boldsymbol{\mu}}_{b(\bar{\mathbf{x}}_k)}^{(p)}$ and the upper-left $J_k \times J_k$ submatrix of $\tilde{\mathbf{C}}_{b(\bar{\mathbf{x}}_k)}^{(p)}$, respectively (cf. the stacked structure of $\boldsymbol{\mu}_{\bar{\mathbf{x}}_k}^{(p-1)}$ in (18) and the block structure of $\mathbf{C}_{\bar{\mathbf{x}}_k}^{(p-1)}$ in (19)).

An SP-based approximate calculation of the messages $n_{l \rightarrow k}^{(p)}(\mathbf{x}_l)$, $l \in \mathcal{N}_k$ in (15) can be performed in a similar manner, due to the structural analogy of (15) to (14). (For SPAWN, $n_{l \rightarrow k}^{(p)}(\mathbf{x}_l)$ is not needed.) We note that, as loopy BP in general [10], SPBP typically exhibits convergence of the mean but suffers from overconfident covariance matrices.

V. COMPUTATION AND COMMUNICATION REQUIREMENTS

Similar to the SP filter [2], SPBP requires the computation of the square root of the $\bar{J}_k \times \bar{J}_k$ matrices $\mathbf{C}_{\bar{\mathbf{x}}_k}^{(p-1)}$ to calculate the SPs $\bar{\mathbf{x}}_k^{(j)}$ in Step 1. This is the most complex part of the SPBP algorithm. An efficient computation of the matrix square root uses the Cholesky decomposition [11], whose complexity is cubic in $\bar{J}_k = J_k + \sum_{l=1}^{|\mathcal{N}_k|} J_l$. Thus, the complexity of SPBP is cubic in $|\mathcal{N}_k|$ and, also, in the number of SPs (which is $2\bar{J}_k + 1$). The complexity of NBP is linear in $|\mathcal{N}_k|$ and quadratic in the number of particles [12]. However, the number of particles in NBP is usually much higher than the number of SPs in SPBP. Moreover, the quadratic and cubic complexity terms of the Cholesky decomposition are rather small (about $\bar{J}_k^3/6$ multiplications, $\bar{J}_k^2/2$ divisions, and \bar{J}_k square root operations are used [11]). Therefore, in many applications, SPBP is significantly less complex than NBP.

SPBP is especially advantageous in decentralized signal processing applications where each variable node in the factor graph corresponds to a sensor node in a wireless sensor network. Because $b^{(p)}(\mathbf{x}_k)$ and $n_{k \rightarrow l}^{(p)}(\mathbf{x}_k)$ are represented by a mean vector and a covariance matrix, at most $J_k + \frac{J_k(J_k+1)}{2} = \frac{J_k(J_k+3)}{2}$ real values per message passing iteration $p \in \{1, \dots, P\}$ have to be transmitted from sensor k to neighboring sensor nodes, rather than hundreds or thousands of particles in NBP. More specifically, at message passing iteration p , sensor k receives $\boldsymbol{\mu}_{l \rightarrow k}^{(p-1)}$ and $\mathbf{C}_{l \rightarrow k}^{(p-1)}$ from all $l \in \mathcal{N}_k$ (this is needed to calculate $\boldsymbol{\mu}_{\bar{\mathbf{x}}_k}^{(p-1)}$ and $\mathbf{C}_{\bar{\mathbf{x}}_k}^{(p-1)}$, see (18) and (19)), and it broadcasts $\boldsymbol{\mu}_{k \rightarrow l}^{(p-1)}$ and $\mathbf{C}_{k \rightarrow l}^{(p-1)}$ to all $l \in \mathcal{N}_k$. These communications are a precondition for Step 1 of the SPBP

algorithm. If the measurement model in (9) involves only substates λ_k of the states \mathbf{x}_k , only the mean and covariance matrix corresponding to λ_k have to be transmitted. (In NBP, similarly, only subparticles corresponding to λ_k have to be transmitted.)

VI. SIMULATION RESULTS

We simulated¹ a decentralized, cooperative, dynamic self-localization scenario [9] using a network of $K = 5$ sensors, of which three are mobile and two are anchors, i.e., static sensors with perfect location information. The state of mobile sensor $k \in \{1, 2, 3\}$ at time $i \in \{0, 1, \dots, 50\}$ consists of the location $\lambda_{k,i} \triangleq (x_{1,k,i} \ x_{2,k,i})^T$ and the velocity, i.e., $\mathbf{x}_{k,i} \triangleq (x_{1,k,i} \ x_{2,k,i} \ v_{1,k,i} \ v_{2,k,i})^T$. Each mobile sensor moves within a field of size 50×50 , performs distance measurements relative to all other sensors, communicates the mean and covariance matrix of its current location to all other sensors, and estimates its own state. We assume that each mobile sensor is able to associate its measurements with the individual sensors. Each anchor sensor $k \in \{4, 5\}$ communicates its own (true) location $\bar{\lambda}_k$. The distance measurement of mobile sensor $k \in \{1, 2, 3\}$ relative to sensor l at time i is (cf. (9))

$$z_{k,l,i} = \begin{cases} \|\lambda_{k,i} - \lambda_{l,i}\| + n_{k,l,i}, & l \in \{1, 2, 3\} \setminus \{k\} \\ \|\lambda_{k,i} - \bar{\lambda}_l\| + n_{k,l,i}, & l \in \{4, 5\}, \end{cases}$$

where $n_{k,l,i}$ is zero-mean Gaussian measurement noise with variance $\sigma_n^2 = 1$.

The states of the mobile sensors evolve independently according to $\mathbf{x}_{k,i} = \mathbf{G}\mathbf{x}_{k,i-1} + \mathbf{W}\mathbf{u}_{k,i}$ [13]. Here, the matrices $\mathbf{G} \in \mathbb{R}^{4 \times 4}$ and $\mathbf{W} \in \mathbb{R}^{4 \times 2}$ are chosen as in [14] and the driving noise vectors $\mathbf{u}_{k,i} \in \mathbb{R}^2$ are Gaussian, i.e., $\mathbf{u}_{k,i} \sim \mathcal{N}(\mathbf{0}, \sigma_u^2 \mathbf{I})$, with variance $\sigma_u^2 = 10^{-4}$; furthermore, $\mathbf{u}_{k,i}$ and $\mathbf{u}_{k',i'}$ are independent unless $(k, i) = (k', i')$. In the generation of the state sequences, this recursive evolution of the $\mathbf{x}_{k,i}$ was initialized with $\mathbf{x}_{1,0} = (0 \ 0 \ 0.2 \ 1)^T$, $\mathbf{x}_{2,0} = (25 \ 50 \ 0.5 \ -0.8)^T$, and $\mathbf{x}_{3,0} = (50 \ 0 \ -1 \ 0.4)^T$. The anchor sensors are located at $\bar{\lambda}_4 = (0 \ 25)^T$ and $\bar{\lambda}_5 = (50 \ 25)^T$ for all i . In the simulation of the various self-localization algorithms, for the mobile sensors, we used the initial prior pdf $f(\mathbf{x}_{k,0}) = \mathcal{N}(\boldsymbol{\mu}_{k,0}, \mathbf{C}_{k,0})$. Here, $\mathbf{C}_{k,0} = \text{diag}\{1, 1, 0.01, 0.01\}$ represents the uncertainty in knowing $\mathbf{x}_{k,0}$, and $\boldsymbol{\mu}_{k,0}$ is a random hyperparameter that was randomly sampled (for each simulation run) from $\mathcal{N}(\mathbf{x}_{k,0}, \mathbf{C}_{k,0})$. For the anchor sensors, the true locations were used. The number of message passing iterations p at each time i was set to $P = 2$.

We compare the proposed SPBP algorithm (using 25 SPs) with two NBP methods for cooperative self-localization, referred to as NBP-1 and NBP-2. NBP-1 [12] is an extension of the method in [5] to moving sensors. NBP-2 differs from NBP-1 in that it performs the message multiplication (11) using Monte Carlo integration instead of Gaussian kernels [15]. All three methods are based on SPAWN. The NBP methods use 250, 500, or 1000 particles. In NBP-1, the bandwidth of the Gaussian kernels was equal to the measurement noise

¹The simulation source files and further information about the simulation setting are available online at <http://www.nt.tuwien.ac.at/about-us/staff/florian-meyer/>.

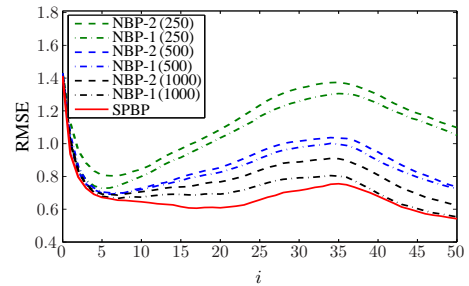


Fig. 1. Root-mean-square location and velocity error (RMSE) of the simulated self-localization algorithms versus time i .

variance $\sigma_n^2 = 1$ [5]. Fig. 1 shows the simulated root-mean-square location and velocity error of the various methods for $i = 1, \dots, 50$. This error was determined by averaging over the three mobile sensors and 1000 simulation runs. It is seen that, for the considered simulation parameters, SPBP outperforms the two NBP methods. We note, however, that NBP would outperform SPBP if the number of particles in NBP was further increased or in the case of an uninformative prior pdf $f(\mathbf{x}_{k,n})$. Also, we expect performance advantages of NBP over SPBP in problems with stronger nonlinearities.

The average runtime of our SPBP implementation on an Intel Xeon X5650 CPU, for all 50 time steps of one simulation run, was 0.61s. The average runtime of NBP-1 was 1.53s, 5.16s, and 19.57s (for 250, 500, and 1000 particles, respectively), that of NBP-2 was 2.01s, 7.27s, and 28.10s. Thus, in this scenario, SPBP is less complex than the two NBP methods.

With SPBP, since our measurement model involves only the two-dimensional location $\lambda_{k,i}$, each mobile sensor broadcasts the mean vector and covariance matrix of $b^{(p)}(\lambda_{k,i}) = \int \int b^{(p)}(\mathbf{x}_{k,i}) dv_{1,k,i} dv_{2,k,i}$ at each message passing iteration p , corresponding to $2 + 3 = 5$ real values. By contrast, for the NBP methods with 250, 500, and 1000 particles, the number of real values broadcast by each mobile sensor at each message passing iteration is 500, 1000, and 2000, respectively. Thus, SPBP requires significantly less communications than the NBP methods.

VII. CONCLUSION

We proposed SPBP as a low-complexity approximation of the BP message passing scheme. SPBP extends the SP filter, also known as unscented Kalman filter, to nonsequential Bayesian inference for general (loopy) factor graphs. Messages and marginal posteriors are represented by mean vectors and covariance matrices, which are calculated using SPs and the unscented transformation. Thereby, SPBP avoids both the linearity assumption of Gaussian BP and the typically high complexity of nonparametric (particle-based) BP. SPBP is especially well suited to certain decentralized inference problems in wireless sensor networks because of its low communication requirements. In particular, we simulated a decentralized, cooperative, dynamic sensor self-localization scenario and demonstrated significant advantages of SPBP over nonparametric BP regarding performance, complexity, and communication requirements.

REFERENCES

- [1] S. J. Julier and J. K. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *Proc. AeroSense-97*, Orlando, FL, pp. 182–193, Apr. 1997.
- [2] E. A. Wan and R. van der Merwe, "The unscented Kalman filter," in *Kalman Filtering and Neural Networks* (S. Haykin, ed.), ch. 7, pp. 221–280, New York, NY: Wiley, 2001.
- [3] H.-A. Loeliger, "An introduction to factor graphs," *IEEE Trans. Signal Process.*, vol. 21, pp. 28–41, Jan. 2004.
- [4] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," *Neural Computation*, vol. 13, pp. 2173–2200, Oct. 2001.
- [5] A. T. Ihler, J. W. Fisher, R. L. Moses, and A. S. Willsky, "Nonparametric belief propagation for self-localization of sensor networks," *IEEE J. Sel. Areas Comm.*, vol. 23, pp. 809–819, Apr. 2005.
- [6] S. Haykin, "Kalman filters," in *Kalman Filtering and Neural Networks* (S. Haykin, ed.), ch. 1, pp. 1–21, New York, NY: Wiley, 2001.
- [7] A. Doucet, N. De Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer, 2001.
- [8] W. Sun and K.-C. Chang, "Unscented message passing for arbitrary continuous variables in Bayesian networks," in *Proc. AAAI '07*, Vancouver, Canada, pp. 1902–1903, 2007.
- [9] H. Wymeersch, J. Lien, and M. Z. Win, "Cooperative localization in wireless networks," *Proc. IEEE*, vol. 97, pp. 427–450, Feb. 2009.
- [10] H. Wymeersch, F. Penna, and V. Savic, "Uniformly reweighted belief propagation for estimation and detection in wireless networks," *IEEE Trans. Wireless Comm.*, vol. 11, no. 4, pp. 1587–1595, 2012.
- [11] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*. New York, NY: Cambridge University Press, 1992.
- [12] J. Lien, J. Ferner, W. Srichavengsup, H. Wymeersch, and M. Z. Win, "A comparison of parametric and sample-based message representation in cooperative localization," *Int. J. Navig. Observ.*, 2012.
- [13] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. Part I: Dynamic models," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, pp. 1333–1364, Oct. 2003.
- [14] J. H. Kotecha and P. M. Djuric, "Gaussian particle filtering," *IEEE Trans. Signal Process.*, vol. 51, pp. 2592–2601, Oct. 2003.
- [15] V. Savic and S. Zazo, "Cooperative localization in mobile networks using nonparametric variants of belief propagation," *Ad Hoc Networks*, vol. 11, no. 1, pp. 138–150, 2013.