Constrained particle filtering using Gaussian sum approximations

Marc-André Beyer * **Gunter Reinig** **

* *Ruhr-University Bochum, Universitätsstraße 150, 44801, Bochum, Germany (Tel.: +49 - 234 - 32 24052; E-Mail: beyer@rus.rub.de)*
** *Ruhr-University Bochum, Universitätsstraße 150, 44801, Bochum, Germany (Tel.: +49 - 234 - 32 24060; E-Mail: reinig@rus.rub.de)*

Abstract: In many filtering problems, there are hard constraints in the state vector that can be a valuable source of information in the estimation process. In this contribution a method to incorporate hard state constraints in particle filters is proposed. The derived approach is based on Gaussian mixture model representation of probability distributions within the particle filter framework and a projection approach to generate constrained samples from these truncated distributions. The developed particle filters show significant improved state estimation performance and robustness against filter divergence compared to their unconstrained counterparts.

Keywords: Particle filters, Bayesian methods, Nonlinear systems

1. INTRODUCTION

In most problem domains it is impossible to create complete, consistent models of the "real" world which requires dealing with uncertainties. Probability theory serves as a methodology for representing and manipulating uncertain beliefs. In essence, probability theory allows to reason under uncertainty. This process is called probabilistic inference (PI). Thus, probabilistic inference is the problem of estimating the internal variables (states and/or parameters) of a system given only incomplete and/or noisy measurements. Application examples of PI are observers for (nonlinear) control, fault detection, inertial navigation, etc.. The complete solution of the PI problem is given by the posterior probability density (filtering density) of the state $x_k \in \mathbb{R}^n$ given the sequence of observations $Y_k \in \mathbb{R}^{m\times k}$, $p(x_k|Y_k = \{y_1, y_2, \ldots, y_k\})$, which allows an estimated state to be computed, e.g. by calculating the mean. The filtering density can be obtained recursively using Recursive Bayesian Estimation (RBE). Although the RBE is the optimal solution of the PI problem the multidimensional integrals are usually intractable for most real world systems. This demands practical solutions.

Particle filters (PF) or Sequential Monte Carlo (SMC) methods are a convenient and easy-to-implement way to solve the RBE numerically for arbitrary probability densities within the RBE framework. Comprehensive surveys regarding PFs can be found in (Doucet et al. (01), Ristic et al. (04)). Over the last few years there has been a proliferation of scientific papers on PFs and their applications. Most popular PFs are the Generic Particle Filter (GPF), the Auxiliary Particle Filter (APF), the Regularized Particle Filter (RPF), and the Kernel Particle Filter (KPF). Based on PFs several hybrids that combine PFs with Gaussian and Gaussian Sum algorithms have been derived, i.e. the Sigma-Point Particle Filter (SPPF) (v.d. Merwe et al. (01)), the Gaussian Mixture Sigma-Point Particle Filter (GMSPPF) (v.d. Merwe et al. (03)), and the Gaussian Sum Particle Filter (GSPF) (Kotecha et al. (03)), to name the most important ones. All PFs and derived hybrids have in common that they operate upon unconstrained state spaces. For certain systems with hard state constraints, this may lead to state space representations without physical meaning and/or with mathematical incorrectness, which typically results in significant performance and robustness decrease. An approach to overcome this problem is proposed in this paper. A new method is derived, that allows PFs to deal with hard state constraints by truncated Gaussian mixture model (GMM) sampling. The sequel is organized as follows:

Section 2 comprises the GMM approximation theory of arbitrary probability distribution functions (*pdf*s), their truncation and truncated sampling. The new proposed projection sampling approach is derived and compared to the rejection sampling approach by means of two multimodal distributions.

Section 3 contains the straightforward application of the developed projection sampling to the particle filter framework.

Section 4 shows the benefit of the derived constrained PFs by means of a nonlinear state estimation problem with hard state constraints.

2. GAUSSIAN SUM APPROXIMATION, PDF TRUNCATION AND SAMPLING

2.1 GMM approximation of pdfs

Every probability density function $p(x)$, $x \in \mathbb{R}^n$ can be approximated as accurately as desired by a GMM (for a proof see Anderson et al. (1979)) given by

$$p(x) \approx p_{GMM}(x) = \sum_{g=1}^{G} \alpha^{(g)} \cdot \mathcal{N}\left(x; \bar{x}^{(g)}, \mathbf{P}^{x,(g)}\right), \quad (1)$$

with component mean $\bar{x}^{(g)}$, positive definite component covariance matrix $\mathbf{P}^{x,(g)}$ and component weight $\alpha^{(g)}$. Furthermore GMMs with Gaussian components that have diagonal covariance matrices only ($\mathbf{P}_{ij}^{x,(g)} = 0, i \neq j$) can also approximate *pdf*s very accurately.

Given a discrete set of samples from $p(x)$ with weights $\{x^{(i)}, w^{(i)}\}_{i=1,...,N}$, an approximating GMM can be gained using a weighted Expectation-Maximization (WEM) clustering algorithm (McLachlan et al. (1997)). The number G of required GMM components is a critical design issue. Either a fixed G can be chosen or a model-order-selection WEM approach can be used to determine the optimal number of Gaussian components adaptively. The latter approach may lead to more accurate GMM approximations of the *pdf*s but with increased computational load. For the experiments in this contribution which are presented in section 4, a fixed $G = 5$ was chosen.

2.2 Bhattacharyya coefficient as metric for goodness of fit

In order to determine the goodness of approximating distributions the Bhattacharyya coefficient (BC) ρ (Bhattacharyya (1943)) can be considered. The BC between a proposal distribution $\pi(x)$ and $p(x)$ is determined by

$$\rho(p, \pi) = \int_{\mathbb{R}^n} \sqrt{p(x) \cdot \pi(x)}\ dx \times 100\,\% \quad (2)$$

$$\rho(p, \pi) = \sum_{x \in \mathbb{R}^n} \sqrt{p(x) \cdot \pi(x)} \times 100\,\% \quad (3)$$

for continuous and discrete probability densities respectively, where the conditions

$$\int_{\mathbb{R}^n} p(x)\ dx = \int_{\mathbb{R}^n} \pi(x)\ dx \stackrel{!}{=} 1 \quad (4)$$

$$\sum_{x \in \mathbb{R}^n} p(x) = \sum_{x \in \mathbb{R}^n} \pi(x) \stackrel{!}{=} 1 \quad (5)$$

must be satisfied. The BC is bounded from above by $\rho \leq 100\,\%$, $\rho = 100\,\%$ denoting a perfect proposal, i.e. $\pi(x) = p(x)$.

2.3 GMM sampling, truncation and truncated sampling

To develop constrained PFs, samples need to be generated from GMM approximated distributions. To do so, the following algorithm can be used:

(1) Define the number of samples to be generated by each component via:
  (a) Draw N samples $u \sim \mathcal{U}(u; 0, 1)$.
  (b) Compute the cumulative sum of weights vector $CSW^{(g)} = \sum_{i=1}^{g} \alpha^{(i)}, g = 1, \ldots, G, CSW^{(0)} = 0$.
  (c) The number of samples q that satisfy $CSW^{(g-1)} \leq u < CSW^{(g)}$ represents the number of samples to be drawn from component g.

(2) Draw q ($\sum q = N$) samples from each component g $x \sim \mathcal{N}(x; \bar{x}^{(g)}, \mathbf{P}^{x,(g)})$.
(3) Concatenate the samples from each component to one sample cloud.

This procedure results in a sample set distributed according to $p_{GMM}(x) \approx p(x)$.

Now consider a set of state vector constraints $\Psi \in \mathbb{R}^{n \times 2}$. The truncated density function $p(x, \Psi)$ of the unconstrained distribution $p(x)$ is then given by

$$p(x, \Psi) = \begin{cases} \dfrac{p(x)}{\int_{x \in \Psi} p(x)\, dx}, & x \in \Psi; \\ 0, & \text{else.} \end{cases} \quad (6)$$

Rejection Sampling (RS) from a truncated pdf

It is easily possible to sample from the unconstrained GMM $p_{GMM}(x)$. By RS variates are generated from $p_{GMM}(x)$, i.e. $x \sim p_{GMM}(x)$ until the condition $x \in \Psi$ is satisfied. The resulting sample set is distributed according to the truncated *pdf* $p(x, \Psi)$. RS is a simple and easy-to-implement technique to incorporate state constraints in *pdf* sampling. The major disadvantage of the rather naive RS approach is that sampling from truncated *pdf*s with little support and/or low probability within the allowed state space results in an exponentially increasing number of rejected samples and equally increasing computational effort. The examples in section 2.4 demonstrate this fact.

Projection Sampling (PS) from a truncated pdf

Consider a probability density $p(x)$ with strictly increasing cummulative distribution function (*cdf*)

$$\mathcal{P}(x) = \int_{-\infty}^{x} p(\tilde{x}) d\tilde{x}. \quad (7)$$

Assume $p(x)$ is truncated with constraints $\Psi = [\Psi_L\ \Psi_U]$, $\Psi_U > \Psi_L$, then $\mathcal{P}(x)$ is also constrained with $\Phi = [\Phi_L\ \Phi_U] = [\mathcal{P}(\Psi_L)\ \mathcal{P}(\Psi_U)]$. If the inverse function of $\mathcal{P}(x)$ exists and the inverse transform (Fishman (1973)) can be applied to $\mathcal{P}(x)$, an algorithm to generate $x \sim p(x, \Psi)$ is:

(1) Generate $u \sim \mathcal{U}(u; 0, 1)$.
(2) Let $\Omega = \Phi_L + (\Phi_U - \Phi_L) \cdot u$.
(3) $x = \mathcal{P}^{-1}(\Omega)$.

PS guarantees that every resulting sample x satisfies the condition $x \in \Psi$. This allows to sample from truncated state spaces with little support and/or low probability without increasing computational effort. The restriction is, that the inverse of the distribution function must, i.e. $\mathcal{P}(x)$ must be bijective.

The remainder is to develop an inverse transform method for arbitrary probability densities $p(x)$, $x \in \mathbb{R}^n$ which are approximated by a GMM that has components with diagonal covariance matrices only. To do so, the following propositions are stated (no proof due to page limitation):

Proposition 1. *For each GMM-component g the $n-$dimensional pdf $p(x)^{(g)}$ is given by*

$$p(x)^{(g)} = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}\sigma_i^{(g)}} \cdot \exp\left(-\frac{1}{2}\left(\frac{x_i - \bar{x}_i^{(g)}}{\sigma_i^{(g)}}\right)^2\right)$$
$$= \prod_{i=1}^{n} p_i(x_i)^{(g)}, \tag{8}$$

where $\boldsymbol{P}^{x,(g)} = diag([(\sigma_1^{(g)})^2 \ (\sigma_2^{(g)})^2 \ \dots \ (\sigma_n^{(g)})^2])$. *Consequently, the cdf is*

$$\mathcal{P}(x)^{(g)} = \prod_{i=1}^{n} \int_{-\infty}^{x_i} \frac{1}{\sqrt{2\pi}\sigma_i^{(g)}} \cdot \exp\left(-\frac{1}{2}\left(\frac{\tilde{x}_i - \bar{x}_i^{(g)}}{\sigma_i^{(g)}}\right)^2\right) d\tilde{x}_i$$
$$= \prod_{i=1}^{n} \mathcal{P}_i(x_i)^{(g)}. \tag{9}$$

Proposition 2. *For an univariate normal distribution with mean* $\bar{x}$ *and covariance* σ^2 *the cdf is bijective and the inverse is given by*

$$x = \mathcal{P}^{-1}\left(\mathcal{P}(x)|\bar{x}, \sigma^2\right) \tag{10}$$

where

$$\mathcal{P}(x) = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}\sigma} \cdot \exp\left(-\frac{1}{2}\left(\frac{\tilde{x} - \bar{x}}{\sigma}\right)^2\right) d\tilde{x}. \tag{11}$$

Proposition 3. *Since Propositions 1 and 2 hold, a sample from component g can be drawn via:*

(1) FOR $i = 1, \dots, n$
  (a) Generate $u \sim \mathcal{U}(u; 0, 1)$.
  (b) Let $\Omega = \mathcal{P}_i^{(g)}(\Psi_{L,i}) + (\mathcal{P}_i^{(g)}(\Psi_{U,i}) - \mathcal{P}_i^{(g)}(\Psi_{L,i})) \cdot u$.
  (c) $x_i = (\mathcal{P}_i^{(g)})^{-1}(\Omega|\bar{x}_i^{(g)}, (\sigma_i^{(g)})^2)$
(2) END FOR
(3) Concatenate $x = [x_1 \ x_2 \ \dots \ x_n]^T$.

The resulting sample satisfies $x \in \Psi$. *The combination of the GMM sampling and the constrained component sampling allows to generate samples resulting in a sample cloud distributed according to* $p(x, \Psi)$.

2.4 Two examples of GMM approximation and truncated sampling

One-dimensional gamma pdf

Consider the one-dimensional gamma distribution with parameters $\mathcal{A}$ and $\mathcal{B}$. The heavily tailed *pdf*, supported on the semi-infinite interval $x \in [0, \infty)$ has density function

$$p(x|\mathcal{A} = 3, \mathcal{B} = 1.25) = \frac{1}{\mathcal{B}^{\mathcal{A}}\Gamma(\mathcal{A})} x^{\mathcal{A}-1} e^{-\frac{x}{\mathcal{B}}}, \tag{12}$$

where Γ denotes the gamma function. Fig. 1 depicts the *pdf* with three GMMs having different number of components G. The 1-component GMM approximates the *pdf* quite poor (Gaussian) and achieves lowest BC, whereas the 6-component GMM nearly fits the *pdf* exactly.

Two-dimensional multimodal pdf

Now consider a two-dimensional distribution with two modes located at $x = [-2 \ -3]^T$ and $x = [5 \ 0]^T$, which is

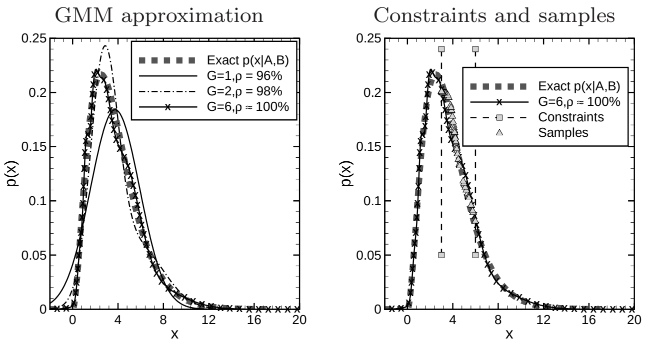

Fig. 1. 1D gamma *pdf* with GMM approximations, constraints and constrained samples.

depicted in Fig. 2. The multimodal *pdf* is constructed by addition of two unequally weighted normal distributions with the respective means and full (!) covariance matrices. The approximating GMM has diagonal component covariance matrices only but still achieves a BC of $\rho = 99.97\,\%$.

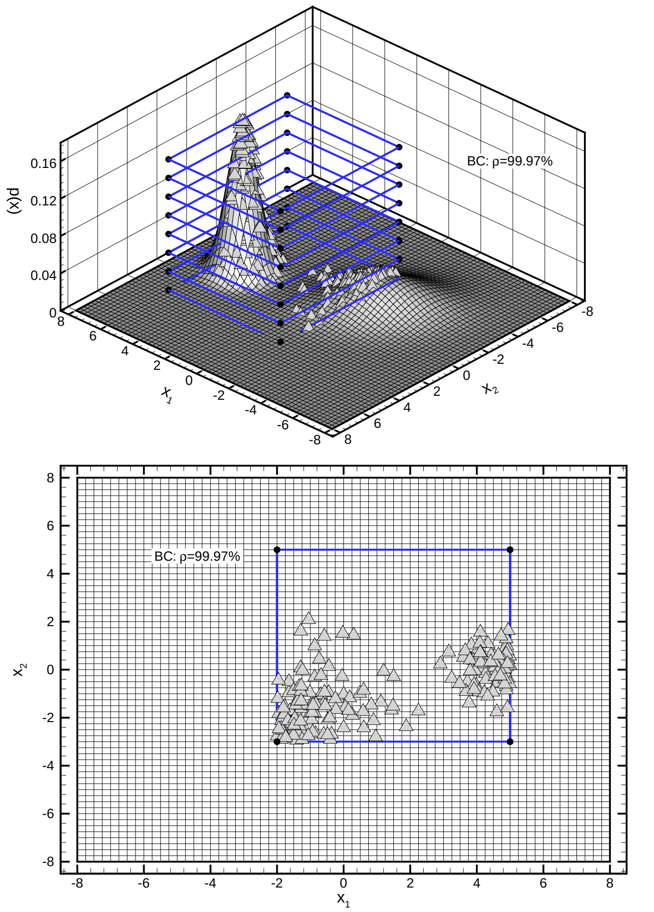

Fig. 2. 2D multimodal *pdf* with constraints and constrained samples.

Numerical experiments for the two examples

To experimentally prove the usability of the proposed sampling approach the computational costs of RS and PS for generating 100 samples from the two distributions and

Ψ	Rejection Sampling		Projection Sampling			True pdf	
	t	Rejections	t	$\mathbb{E}[x]$	$\mathbb{E}[(x-\mathbb{E}[x])^2]$	$\mathbb{E}[x]$	$\mathbb{E}[(x-\mathbb{E}[x])^2]$
1D gamma pdf example							
[3 6]	0.100 s	162	0.0021 s	4.07	0.61	4.25	0.69
[8 10]	0.837 s	2446	0.0021 s	8.78	0.31	8.81	0.31
[12 14]	8.944 s	27560	0.0021 s	12.70	0.28	12.79	0.31
[14 14.5]	92.34 s	286600	0.0021 s	14.19	0.02	14.24	0.02
2D multimodal pdf example							
$\begin{bmatrix} -2 & 5 \\ -3 & 5 \end{bmatrix}$	0.060 s	69	0.005 s	$\begin{bmatrix} 1.9277 \\ -0.6973 \end{bmatrix}$	$\begin{bmatrix} 7.1657 & 2.3000 \\ 2.3000 & 2.5881 \end{bmatrix}$	$\begin{bmatrix} 1.9308 \\ -0.6976 \end{bmatrix}$	$\begin{bmatrix} 7.1700 & 2.3026 \\ 2.3026 & 2.5685 \end{bmatrix}$
$\begin{bmatrix} -1 & 1 \\ -1 & 1 \end{bmatrix}$	0.69 s	1349	0.005 s	$\begin{bmatrix} -0.5653 \\ -0.3766 \end{bmatrix}$	$\begin{bmatrix} 0.2544 & 0.0476 \\ 0.0476 & 0.2390 \end{bmatrix}$	$\begin{bmatrix} -0.5603 \\ -0.3511 \end{bmatrix}$	$\begin{bmatrix} 0.2482 & 0.0388 \\ 0.0388 & 0.2490 \end{bmatrix}$
$\begin{bmatrix} -.1 & .1 \\ -.1 & .1 \end{bmatrix}$	56.95 s	115328	0.005 s	$\begin{bmatrix} -0.0088 \\ -0.0035 \end{bmatrix}$	$\begin{bmatrix} 0.0033 & \approx 0 \\ \approx 0 & 0.0033 \end{bmatrix}$	$\begin{bmatrix} -0.0088 \\ -0.0034 \end{bmatrix}$	$\begin{bmatrix} 0.0033 & \approx 0 \\ \approx 0 & 0.0033 \end{bmatrix}$

Table 1. Comparison of the two sampling approaches for different *pdf*s and constraints.

different Ψ are comprised in Table 1. Furthermore the first and second central moment computed from 1000 PS-generated samples are compared to the true moments of the *pdf*s. Not only consistently outperforms the PS the RS in computational load (especially with decreasing probability or support), but also matches the central moments of the true *pdf*s (also higher moments). Consequently, the proposed approach generates samples fast and is applicable to arbitrary multidimensional probability distributions.

3. CONSTRAINED PARTICLE FILTERING

Given the stochastic $n-$dimensional dynamic state space model

$$x_k = \mathcal{F}\left(x_{k-1}, u_{k-1}, \kappa_{k-1}\right) \tag{13}$$

$$y_k = \mathcal{H}\left(x_k, u_k, \gamma_k\right), \tag{14}$$

with state x_k, exogenous input $u \in \mathbb{R}^p$ and system and measurement noise $\kappa \in \mathbb{R}^v$ and $\gamma \in \mathbb{R}^z$, the PI problem is the problem of estimating the states of the system given only incomplete and/or noisy measurements y_k. As already stated in the introduction, the optimal solution of the PI problem is given by the posteriori probability density of the state x_k given the measurement sequence Y_k, denoted as $p(x_k|Y_k)$, with the RBE as closed recursive solution. The result of the RBE, i.e. the posterior *pdf* is computed in two steps, namely prediction and update (or filtering). The prediction step uses the stochastic state equation (13) to obtain the prior *pdf* $p(x_k|Y_{k-1})$

$$p\left(x_k|Y_{k-1}\right) = \int\limits_{\mathbb{R}^n} p\left(x_k|x_{k-1}\right) \cdot p\left(x_{k-1}|Y_{k-1}\right) dx_{k-1}, \tag{15}$$

with $p(x_{k-1}|Y_{k-1})$ known from the previous iteration step and the transition density $p(x_k|x_{k-1})$ determined by (13). When a new measurement y_k becomes available, Bayes' rule gives the posterior *pdf* $p(x_k|Y_k)$

$$p\left(x_k|Y_k\right) = \frac{p\left(y_k|x_k\right) \cdot p\left(x_k|Y_{k-1}\right)}{\int_{\mathbb{R}^n} p\left(y_k|x_k\right) \cdot p\left(x_k|Y_{k-1}\right) dx_k}, \tag{16}$$

with likelihood density $p(y_k|x_k)$ determined by (14). PFs approximate the densities $p(x_k|Y_{k-1})$ and $p(x_k|Y_k)$ with a sum of N Dirac functions centered in $\{x_k^{(i)}\}_{i=1,\ldots,N}$ as

$$p \approx \sum_{i=1}^{N} w_k^{(i)} \delta\left(x_k - x_k^{(i)}\right). \tag{17}$$

The derivation of constrained PFs is now straightforward. The unconstrained prior and posterior *pdf*, $p(x_k|Y_{k-1})$ and $p(x_k|Y_k)$ have to be approximated by GMMs and new constrained sample sets $\{x_k^{(i)}, w_k^{(i)}|\Psi\}_{i=1,\ldots,N}$ representing the constrained prior and posterior *pdf*

$$p\left(x_k|Y_{k-1},\Psi\right) = \begin{cases} \dfrac{p\left(x_k|Y_{k-1}\right)}{\int_{x_k\in\Psi} p\left(x_k|Y_{k-1}\right) dx_k}, & x_k \in \Psi; \\ 0, & \text{else.} \end{cases} \tag{18}$$

and

$$p\left(x_k|Y_k,\Psi\right) = \begin{cases} \dfrac{p\left(x_k|Y_k\right)}{\int_{x_k\in\Psi} p\left(x_k|Y_k\right) dx_k}, & x_k \in \Psi; \\ 0, & \text{else.} \end{cases} \tag{19}$$

can be drawn by the approach presented in the previous section. This procedure guarantees the samples to satisfy $x_k \in \Psi$ and the PFs to maintain physical correct state samples.

4. EXAMPLE

The best available approximate solution of a specific PI problem, i.e. the best available filter, cannot be predicted (in most cases) without (simulatively) testing and evaluating the filter algorithms. Depending on the system characteristics, e.g. (non)linearity, noise distributions, etc., the performance of the filter algorithms may differ substantially. This inspired the development of a new software ReBEL-IoN, designed at *Lehrstuhl für Regelungssysteme und Steuerungstechnik* (RUS) to carry out quantitative comparative evaluation of stochastic state estimators. ReBEL-IoN is a Matlab based toolbox with graphical user interface for testing, evaluating, and selecting probabilistic inference algorithms in a general state space model framework. The implemented evaluation concept incorporates several quantitative criteria to be examined, amongst others relative estimation performance and filter robustness. The following simulation studies have been carried out using the toolbox. For details see (Beyer et al. (2008)).

4.1 Performance metrics for evaluation

Estimation performance

Filter performance is a key issue of the comparative evaluation. In this contribution the performance criteria which compare the algorithms are the root mean squared error (RMSE) and the root time averaged mean squared error ($\mathcal{RT}$) which are defined by

$$RMSE = \sqrt{\frac{1}{MC}\sum_{i=1}^{MC}\left(\hat{x}_k^{(i)} - x_k^{(i)}\right)^2} \tag{20}$$

and

$$\mathcal{RT} = \sqrt{\frac{1}{K \cdot MC}\sum_{k=1}^{K}\sum_{i=1}^{MC}\left(\hat{x}_k^{(i)} - x_k^{(i)}\right)^2}, \tag{21}$$

where K denotes the final time instant, i.e. the total number of observations and MC denotes the total number of Monte Carlo simulations.

Filter robustness

The robustness and stability of the filter algorithms against modeling errors and noises is a decisive criterion of the usability of a filter. To make a statement about the robustness of a filter the number of diverged tracks per simulation study can be used to determine a robustness index as follows:

$$J_{Rob} = \left(1 - \frac{\#_{Divergence}}{MC}\right) \times 100\,\% \tag{22}$$

A track is diverged if the algorithms numerically diverge or if the deviation between reference and estimated trajectory exceeds a user determined threshold. Although this approach yields no proposition of the distance to the stability margin, it yields an empirical robustness index which is bounded within $0\,\% \leq J_{Rob} \leq 100\,\%$, where $J_{Rob} = 100\,\%$ denotes a robust filter.

4.2 Example: Gas-phase reversible reaction

Consider the gas-phase, reversible reaction (Haseltine et al. (2003))

$$2A \xrightarrow{\bar{k}} B \quad \bar{k} = 0.16 \tag{23}$$

with stoichiometric matrix $\nu = [-2\ 1]$ and reaction rate $r = \bar{k} \cdot \phi_A^2$. States are $x = [\phi_A\ \phi_B]^T$, where ϕ_j denotes the partial pressure of species j. Assuming that the ideal gas law holds, and that the reaction occurs in a well-mixed isothermal batch reactor, the time-discrete state space model is given by

$$x_k = x_{k-1} + \tau \cdot \nu^T r_{k-1} = x_{k-1} + \tau \cdot \nu^T \bar{k} x_{1,k-1}^2 \tag{24}$$

with measurement equation

$$y_k = [1\ 1]x_k \tag{25}$$

and discretization step size $\tau = 0.1\,\mathrm{s}$. For stochastic state filtering, consider the following stochastic state space model:

$$\begin{aligned} x_k &= x_{k-1} + \tau \cdot \nu^T \bar{k} x_{1,k-1}^2 + \kappa_{k-1} \\ &= \mathcal{F}\left(x_{k-1}, u_{k-1}, \kappa_{k-1}\right) \end{aligned} \tag{26}$$

$$y_k = [1\ 1]x_k + \gamma_k = \mathcal{H}\left(x_k, u_k, \gamma_k\right), \tag{27}$$

where $\kappa_{k-1} \sim \mathcal{N}(\kappa; 0, \mathrm{diag}([0.001^2\ \ 0.001^2]))$ and $\gamma_k \sim \mathcal{N}(\gamma; 0, 0.1^2)$ are zero-mean Gaussian distributed random variables. The initial assumed state distribution (used in the filters) accords with $\hat{x}_0 \sim \mathcal{N}(\hat{x}_0; [0.1\ 4.5]^T, \mathrm{diag}([6^2\ 6^2]))$ whereas the true initial state is $x_0 = [3\ 1]^T$. Note, that the initial guess is poor. To test the proposed approach, four PFs, namely GPF, APF, RPF and GMSPPF were equipped with the truncated projection sampling approach ($G = 5$) and tested on the system described above. All PFs possess a particle cloud size of 200. The system is stochastically simulated with a total number of $K = 300$ observations. To allow a fair comparison, the estimation results were averaged over 2000 independent Monte Carlo runs. In order to show the influence of different state constraints information on the estimation error and robustness, two scenarios with different information about the constraints were investigated and compared to the state estimates of their unconstrained counterparts and the Extended Kalman Filter (EKF) which is most widely used in control design [1]:

(1) **Scenario 1:** The partial pressures of the system are nonnegative. In this scenario (Ψ_1) this nonnegativity information was supplied to the PFs.
(2) **Scenario 2:** In this particular example the partial pressures also never exceed an upper threshold of 3 Pa. In this scenario (Ψ_2) the full information of the bounds was available to the PFs.

In Fig. 3 the RMSEs of the four filters with different constraints information are plotted. As expected, the RMSE decreases significantly the more information about the constraints is supplied to the filters. Moreover, the plots show that the estimation performance varies substantially from filter to filter which is due to the general curse of the approximate RBE solutions for nonlinear state estimation problems.

Table 4.2 presents a comparison of the four algorithms and the EKF in terms of estimation performance, now indicated by the $\mathcal{RT}$, and filter robustness. Again, it is obvious that the more information about the state constraints is available to the PFs, the better state estimates with higher robustness are gained. The unconstrained $\mathcal{RT}$s are higher with lower robustness. The incorporation of the nonnegativity information leads to robust filters (robustness indices close to 100 %) with improved estimation performance. This is noteworthy since the diverged tracks are not accounted for the $\mathcal{RT}$. If the upper bound is also supplied to the PFs, the estimation performance of all filters can be improved further, while maintaining robustness.

5. CONCLUSIONS & FUTURE WORK

In this contribution a method to incorporate hard state constraints of the form $\Psi_L \leq x \leq \Psi_U,\ \Psi_U > \Psi_L,\ \Psi \in \mathbb{R}^{n \times 2}$ in particle filters was proposed. The approach is based on the idea that the prior and posterior densities in the particle filter framework can be approximated arbitrarily well by unconstrained Gaussian mixture models from which constrained sample sets can be generated via projection sampling, representing the constrained prior and posterior *pdf*. This procedure guarantees the state

[1] Note that the state estimates of the EKF were "clipped" for scenarios 1 and 2, i.e. partial pressures rendered negative (or greater 3 Pa) by the filter update were zeroed (or set to 3 Pa).

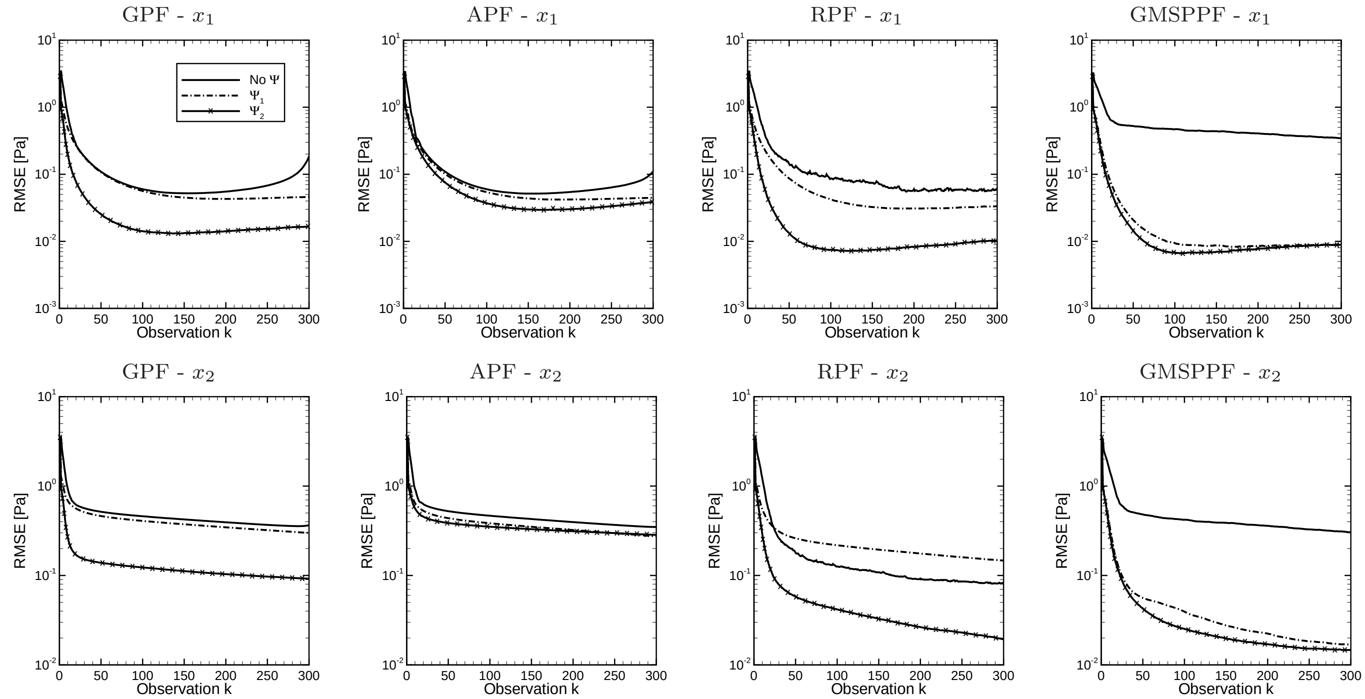

Fig. 3. RMSEs of the investigated PFs with(out) constraints information for the gas-phase reversible reaction example.

Algorithm	Without Constraints $\mathcal{RT}$	Without Constraints J_{Rob}	Ψ	With Constraints $\mathcal{RT}$	With Constraints J_{Rob}
EKF	2.26\2.10 Pa	64.85 %	Ψ_1 :	0.55\0.64 Pa	100.0 %
			Ψ_2 :	0.22\0.25 Pa	100.0 %
GPF	0.39\0.59 Pa	64.40 %	Ψ_1 :	0.24\0.47 Pa	100.0 %
			Ψ_2 :	0.19\0.26 Pa	100.0 %
APF	0.39\0.59 Pa	65.70 %	Ψ_1 :	0.23\0.44 Pa	100.0 %
			Ψ_2 :	0.22\0.41 Pa	100.0 %
RPF	0.44\0.46 Pa	78.80 %	Ψ_1 :	0.23\0.33 Pa	100.0 %
			Ψ_2 :	0.19\0.23 Pa	100.0 %
GMSPPF	0.63\0.61 Pa	100.0 %	Ψ_1 :	0.20\0.24 Pa	100.0 %
			Ψ_2 :	0.19\0.23 Pa	100.0 %

$\Psi_1: \ x > 0$, $\Psi_2: \ 0 < x \leq 3$

Table 2. Performance metrics comparison of the investigated PFs and the EKF with(out) constraints information for the gas-phase reversible reaction example.

particle cloud to maintain a physical correct approximation of the state distribution. The simulation studies show, that the developed PFs with state constraints consistently outperform their unconstrained counterparts in estimation performance and robustness against filter divergence. Both performance metrics improve the more information about the constraints is supplied to the filter algorithms.

Future work will deal with the incorporation of linear state constraints of the form $\mathbf{A}x \leq b$ via GMM approximation and sampling techniques that allow the GMMs to possess components with full covariance matrices.

REFERENCES

A. Doucet, N. de Freitas, and N. Gordon. Sequential Monte Carlo Methods in Practice. Springer, New York, 2001.

B. Ristic, S. Arulampalam, and N. Gordon. Beyond the Kalman Filter - Particle filters for tracking applications. Artech House, 2004.

R. van der Merwe, N. de Freitas, A. Doucet, and E. Wan. The Unscented Particle Filter. Advances in Neural Information Processing Systems, volume 13, 2001.

R. van der Merwe, and E. Wan. Gaussian Mixture Sigma-Point Particle Filters for Sequential Probabilistic Inference in Dynamic State-Space Models. Proceedings of the International Conference on Acoustics, Speech and Signal Processing, 2003.

J.H. Kotecha, and P.M. Djurić Gaussian Sum Particle Filtering. IEEE Transactions on Signal Processing, volume 51, number 10, 2003.

B. Anderson and J. Moore. Optimal Filtering. Prentice Hall, 1979.

G. McLachlan and T. Krishnan. The EM-algorithm and Extensions Wiley, 1997

A. Bhattacharyya. On a measure of divergence between two statistical populations defined by probability distributions. Bulletin of the Calcutta Mathematical Society, volume 35, pp. 99-109, 1943.

G.S. Fishman. Concepts and methods in discrete event digital simulation. Wiley & Sons, 1973.

M.-A. Beyer and G. Reinig. Matlab based Toolbox ReBEL-IoN. http://www.rebel-ion.de

E.L. Haseltine and J.B. Rawlings. A Critical Evaluation of Extended Kalman Filtering and Moving Horizon Estimation. TWMCC Technical Report 2002-2003, 2003