

CONSTRAINED STATE ESTIMATION IN PARTICLE FILTERS

Bradley Ebinger, Nidhal Bouaynaya and Robi Polikar^{*}

Roman Shterenberg

Rowan University
Department of Electrical and Computer Engineering
Glassboro, NJ

University of Alabama at Birmingham
Department of Mathematics
Birmingham, AL

ABSTRACT

Dynamical systems are often required to satisfy certain constraints arising from basic physical laws, mathematical properties or geometric considerations. Incorporating constraints improves the performance of state estimation and increases the accuracy compared to unconstrained estimation.

Particle filters (PF) have gained popularity within the signal processing community, thanks to their asymptotically optimal estimation for nonlinear and non-Gaussian state-space models. However, their constrained formulation has emerged only very recently; and the developments to incorporate state constraints in particle filters have mainly relied on constraining all particles of the PF. This approach is termed Pointwise or Particle Density Truncation (PDT).

In this paper, we show that PDT constrains the posterior density of the state rather than the conditional mean estimate, which leads to more stringent and possibly completely different or even irrelevant conditions than the original constraints. Subsequently, we introduce an alternative novel solution to constrained particle filtering, which enforces the constraints on the conditional mean without further restricting the state posterior density. The proposed approach is termed Mean Density truncation (MDT) and is compared to PDT and projection methods for a severely nonlinear model.

Index Terms— Constrained Particle Filtering; State Estimation; Constrained Bayesian Estimation.

1. INTRODUCTION

The state of many dynamical systems is often required to satisfy certain constraints arising from basic physical laws, mathematical properties or geometric considerations, e.g., maximum power or transmission capacity, energy conservation laws and bounded parameters. In fact, constrained systems are already omnipresent in many real-world applications including camera tracking [1], fault diagnosis [2], chemical

processes [3], vision-based systems [4], target tracking [5, 6], biomedical systems [7], robotics [8] and navigation [9].

Particle Filters (PF) are a broad class of Monte Carlo algorithms, which provide approximate solutions to analytically intractable inference problems, which can include nonlinear and non-Gaussian modeling scenarios. PFs can solve these problems by using *particles*, which sample the state space of the system. These particles are then weighted to estimate the state posterior density. The estimation converges, in the mean-square error, to the true posterior density of the state. PFs have become a viable alternative to more traditional techniques, such as the Extended Kalman Filter (EKF) due to the PF's ability to calculate posterior densities without using functional approximation such as local linearization techniques or assume Gaussian noise.

However, the very numerical nature of the particle filters, which constitutes their strength for multidimensional numerical integration, becomes their major weakness in handling constraints on the state. The main difficulty of the constrained PF problem stems from the fact that every particle in the particle approximation of the state posterior density is a local representation of the density, and thus cannot characterize global properties of the density, such as constraints on the conditional mean or any other functional expectation. The current trend in constrained particle filtering simply enforces the constraints on all particles of the PF. This approach, however, constrains the posterior density of the state rather than its mean, which leads to more stringent conditions and possibly a completely different condition than the original constraints (see Fig. 1). We refer to the approach of constraining all particles as the *Pointwise Density Truncation* (PDT) method.

In this paper, we introduce a new approach called, *Mean Density Truncation* (MDT), that imposes the state constraints on the conditional mean estimate without further restraining the posterior distribution of the state. The paper is organized in the following way: the unconstrained PF framework is reviewed in Section 2, the PDT and MDT approaches are advanced in Section 3, simulation results that compare PDT, MDT and projection approaches are presented in Section 4, and concluding remarks and future directions are discussed in Section 5.

^{*}This project is supported by Award Number R01GM096191 from the National Institute Of General Medical Sciences (NIH/NIGMS). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute Of General Medical Sciences or the National Institutes of Health.

2. THE UNCONSTRAINED PARTICLE FILTER

We consider a discrete-time state-space model defined by the following state and measurement equations:

$$\begin{aligned} \mathbf{x}_k &= \mathbf{f}_k(\mathbf{x}_{k-1}) + \mathbf{w}_k, \\ \mathbf{y}_k &= \mathbf{h}_k(\mathbf{x}_k) + \mathbf{v}_k, \end{aligned} \quad (1)$$

where $\mathbf{x}_k \in \mathbb{R}^{n_x}$ and $\mathbf{y}_k \in \mathbb{R}^{n_y}$ represent the system state and the system output, respectively, \mathbf{f}_k and \mathbf{h}_k are known, possibly non-linear, mappings, and \mathbf{w}_k and \mathbf{v}_k are zero-mean process and measurement noise with known probability density functions (pdfs), g and r , respectively.

Let $\mathbf{Y}^k = [\mathbf{y}_1, \dots, \mathbf{y}_k]$ denote the history of observations up to time k . In the Bayesian context, inference of \mathbf{x}_k given a realization of the observations \mathbf{Y}^k relies upon the posterior density $p(\mathbf{x}_k | \mathbf{Y}^k)$. Using the Bayesian rule, we can obtain the following two-step Bayesian recursion formula:

$$p(\mathbf{x}_k | \mathbf{Y}^{k-1}) = \int g(\mathbf{x}_k | \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} | \mathbf{Y}^{k-1}) d\mathbf{x}_{k-1} \quad (2)$$

$$p(\mathbf{x}_k | \mathbf{Y}^k) = \frac{r(\mathbf{y}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{Y}^{k-1})}{\int r(\mathbf{y}_k | \mathbf{x}) p(\mathbf{x} | \mathbf{Y}^{k-1}) d\mathbf{x}} \quad (3)$$

Equations (2)-(3) are a conceptual solution because the integrals defined are, in general, intractable. For the linear Gaussian model, it is easy to check that $p(\mathbf{x}_k | \mathbf{Y}^k)$ is a Gaussian distribution whose mean and covariance can be computed using the Kalman filter. However, for most nonlinear non-Gaussian models, it is not possible to compute these distributions in closed-form.

The PF approximates the posterior pdf using an ensemble of *particles* $\{\mathbf{x}_k^{(i)}\}_{i=1}^N$ and their associated weights $\{w_k^{(i)}\}$:

$$\hat{p}(\mathbf{x}_k | \mathbf{Y}^k) = \sum_{i=1}^N w_k^{(i)} \delta(\mathbf{x}_k - \mathbf{x}_k^{(i)}) \quad (4)$$

where $\delta(\cdot)$ is the Dirac delta function and N is the number of particles. Ideally, the particles are required to be sampled from the true posterior, $p(\mathbf{x}_k | \mathbf{Y}^k)$, which is not available. Therefore, another distribution, referred to as the *importance distribution* or the *proposal distribution*, $q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{y}_k)$, is used. The particles at time k are sampled from $\mathbf{x}_k^{(i)} \sim q(\mathbf{x}_k | \mathbf{x}_{k-1}^{(i)}, \mathbf{y}_k)$. The importance weight of each particle $\mathbf{x}_k^{(i)}$ is computed as

$$\tilde{w}_k^{(i)} = w_{k-1}^{(i)} \frac{r(\mathbf{y}_k | \mathbf{x}_k^{(i)}) g(\mathbf{x}_k^{(i)} | \mathbf{x}_{k-1}^{(i)})}{q(\mathbf{x}_k^{(i)} | \mathbf{x}_{k-1}^{(i)}, \mathbf{y}_k)}, \quad (5)$$

where \tilde{w}_k are the un-normalized weights [10]. The normalized weights in (4) are given by $w_k^{(i)} = \tilde{w}_k^{(i)} / \sum_{j=1}^N \tilde{w}_k^{(j)}$.

3. THE CONSTRAINED PARTICLE FILTER

We focus on the discrete state-space model in (1) augmented with the following general constraint

$$\mathbf{a}_k \leq \phi_k(\mathbf{x}_k) \leq \mathbf{b}_k, \quad (6)$$

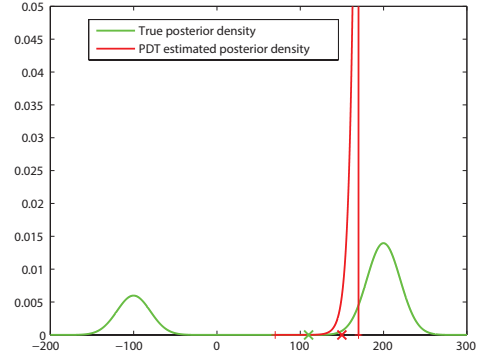


Fig. 1. Illustration of the PDT approach for an interval-constrained system, $x_n \in [70, 170]$ for all n . The true posterior density (green curve) is multimodal with mean 110 (green x-mark). If all particles are constrained to be within the interval $[70, 170]$, then the estimated posterior density (red curve) will be a truncated exponential density, that is dramatically different from the true posterior distribution.

where ϕ_n is the constraint function at time n and the inequality holds for all elements. It is important to emphasize the fact that the constraint needs to only be satisfied by the state estimate given by the conditional mean, i.e., we must have

$$\mathbf{a}_k \leq \phi_k(\hat{\mathbf{x}}_k) = \phi_k(E[\mathbf{x}_k | \mathbf{Y}^k]) \approx \phi_k\left(\sum_{j=1}^N w_k^{(j)} \mathbf{x}_k^{(j)}\right) \leq \mathbf{b}_k$$

This mean constraint is not a local condition, meaning there are many ways to globally constrain the mean. Projection of the unconstrained density onto the constraint set is only one possible option. The widely used approach in constrained sequential Monte Carlo is the *acceptance/rejection* approach, which enforces the constraints by simply rejecting the particles violating them [11]. The acceptance/rejection procedure does not make any assumption on the distributions and therefore maintains the generic property of the particle filter. However, the number of samples will be reduced and hence the estimation accuracy may decrease, especially with a poor choice of the proposal density. An extreme example is when most (or all) of the particle violate the constraint and the algorithm fails [12].

3.1. Pointwise Density Truncation (PDT)

The current practice in the literature constrains the mean of the posterior distribution by imposing the constraints on all particles of the PF [12–20]. However, this is not true. Imposing the constraint on all particles results, in general, in a stronger constraint and possibly a completely different or even irrelevant condition. To see this, let us consider the scalar case with $\mathcal{C}_k = [a, b]$ for all k : the state estimate is constrained in the interval $[a, b]$ or $a \leq x_k \leq b$. Constraining every particle to be within the interval $[a, b]$ is equivalent to constraining the support of the posterior distribution to this interval, which is a much stronger condition than constrain-

ing the mean of the distribution, or any point estimate, to be inside the interval. We refer to this approach as *pointwise density truncation* or *particle density truncation* (PDT). Since the particle filter estimates the posterior density of the state, imposing stronger constraints may, and in general will, result in an erroneous estimation of the density, as illustrated in Fig. 1.

3.2. Mean Density Truncation (MDT)

In the constrained state-space model, the constraints must be satisfied by the estimate of the conditional mean. Unlike the pointwise density truncation approach, which enforces the constraints on all particles, we propose the mean density truncation (MDT) approach, which constrains only one particle in order to confine the estimated mean to the desired constraints. In the MDT approach, $(N - 1)$ unconstrained particles are drawn from the proposal distribution. Then, the N^{th} particle is constrained in order to impose the conditions on the sample mean. A constraint of the form $a_k \leq g(\mathbf{x}_k) \leq b_k$ can be equivalently expressed as

$$a_k \leq g\left(\sum_{j=1}^N \omega_k^{(j)} \mathbf{x}_k^{(j)}\right) \leq b_k. \quad (7)$$

For simplicity, we will assume that the weights are given by the likelihood, i.e., the proposal density is the prior distribution function; the essence of the MDT method remains the same in the general case, where the proposal density is different from the prior distribution. Separating the summation of the $(N - 1)$ unconstrained particles from the N^{th} particle, and taking into account the normalization of the weights, the constraint becomes

$$a_k \leq g\left(\frac{\sum_{j=1}^{N-1} p(\mathbf{y}_k|\mathbf{x}_k^{(j)})\mathbf{x}_k^{(j)} + p(\mathbf{y}_k|\mathbf{x}_k^{(N)})\mathbf{x}_k^{(N)}}{\sum_{j=1}^N p(\mathbf{y}_k|\mathbf{x}_k^{(j)})}\right) \leq b_k$$

Then, conditions on the N^{th} particle can be derived depending on the explicit expression of the constraint function g . For instance, if we consider the interval-type constraint, i.e., g is the identity function, then the above inequality becomes equivalent to the two inequalities,

$$\sum_{j=1}^{N-1} p(\mathbf{y}_k|\mathbf{x}_k^{(j)})(a_k - \mathbf{x}_k^{(j)}) \leq p(\mathbf{y}_k|\mathbf{x}_k^{(N)})(\mathbf{x}_k^{(N)} - a_k), \quad (8)$$

$$\sum_{j=1}^{N-1} p(\mathbf{y}_k|\mathbf{x}_k^{(j)})(b_k - \mathbf{x}_k^{(j)}) \geq p(\mathbf{y}_k|\mathbf{x}_k^{(N)})(\mathbf{x}_k^{(N)} - b_k). \quad (9)$$

Letting $q_1(\mathbf{x}_k^{(N)}) = p(\mathbf{y}_k|\mathbf{x}_k^{(N)})(\mathbf{x}_k^{(N)} - a_k)$ and $q_2(\mathbf{x}_k^{(N)}) = p(\mathbf{y}_k|\mathbf{x}_k^{(N)})(\mathbf{x}_k^{(N)} - b_k)$, we obtain the two inequalities

$$\begin{aligned} q_1(\mathbf{x}_k^{(N)}) &\geq C_1(\{\mathbf{x}_k^{(j)}\}_{j=1}^{N-1}), \\ q_2(\mathbf{x}_k^{(N)}) &\leq C_2(\{\mathbf{x}_k^{(j)}\}_{j=1}^{N-1}), \end{aligned} \quad (10)$$

which have to be satisfied for the N^{th} particle only. C_1 and C_2 are two constants, which depend only on the already sampled $(N - 1)$ unconstrained particles and their weights. Depending on the likelihood function, Eq (10) can be solved analytically or numerically. The solution to (10) may not be

unique. Many " N^{th} particles" can satisfy (10), all of them enforcing the original constraint on the sample mean estimate. These different solutions may lead to different constrained estimates. We found, in our preliminary results, that the solution with the highest weight (here, likelihood) leads to the most accurate estimator among all other solutions having lower weights.

If the proposal distribution is chosen poorly, the $(N - 1)$ unconstrained particles will lie in a low probability region of the posterior density of the state. In this case, it may not be possible to find an N^{th} particle that satisfies (10), thus imposing the constraint on the sample mean. Intuitively, if the initial particle sampling is poor, then one additional particle may not be able to force the mean to satisfy the desired constraints. We advance two solutions to ensure the existence of an N^{th} particle that will enforce the constraint on the sample mean: m^{th} -order MDT and inductive MDT (IMDT).

In the case where one particle may not be sufficient to constrain the mean, it seems reasonable to consider constraining more than one particle, e.g., two, three or up to $m \leq N$ particles. These m constrained particles will ensure that the sample mean satisfies the desired constraint. The MDT method is thus termed 1^{st} -order MDT, and its extension to m constrained particles is called m^{th} -order MDT. In the m^{th} -order MDT, $(N - m)$ unconstrained particles are sampled from the proposal distribution, and the remaining m particles are constrained in order to satisfy the condition on the sample mean. It is important to notice that when $m = N$, the N^{th} -order MDT is very different from the PDT method: In the PDT approach, the original constraint is imposed on all particles. On the other hand, the N^{th} -order MDT constrains the particles, as in Eq. (7), in order to impose the desired condition on the sample mean.

4. SIMULATION RESULTS

We consider the following nonlinear dynamic system

$$\mathbf{x}_{k+1} = \frac{\mathbf{x}_k}{2} + 25 \frac{\mathbf{x}_k}{1 + \mathbf{x}_k^2} + 8 \cos(1.2k) + w_k \quad (11)$$

$$y_k = \frac{\mathbf{x}_k^3}{25} + \mathbf{x}_k + v_k; \quad -5 \leq \mathbf{x}_k \leq 5,$$

where w_k and v_k are zero-mean Gaussian white noise. This example is severely nonlinear, both in the system and the measurement equations. It was shown in [21] that the EKF fails in estimating the true state value of this unconstrained system.

Figure 2(a) shows the true and estimated trajectories using 1^{st} -order MDT, PDT and projection of the unconstrained mean estimate. The results are shown for 1000 Monte Carlo simulations. It is seen that, on average, the 1^{st} -order MDT leads to more accurate estimation of the dynamic state, where both the mean-square error and the variance are smaller. Figures 2 (b),(c),(d) show the posterior density of the constrained state as it evolves over time, for 1^{st} -order MDT, PDT and projection, respectively. First, observe that the PDT approach (Fig. 2(c)) produces posterior distributions with a

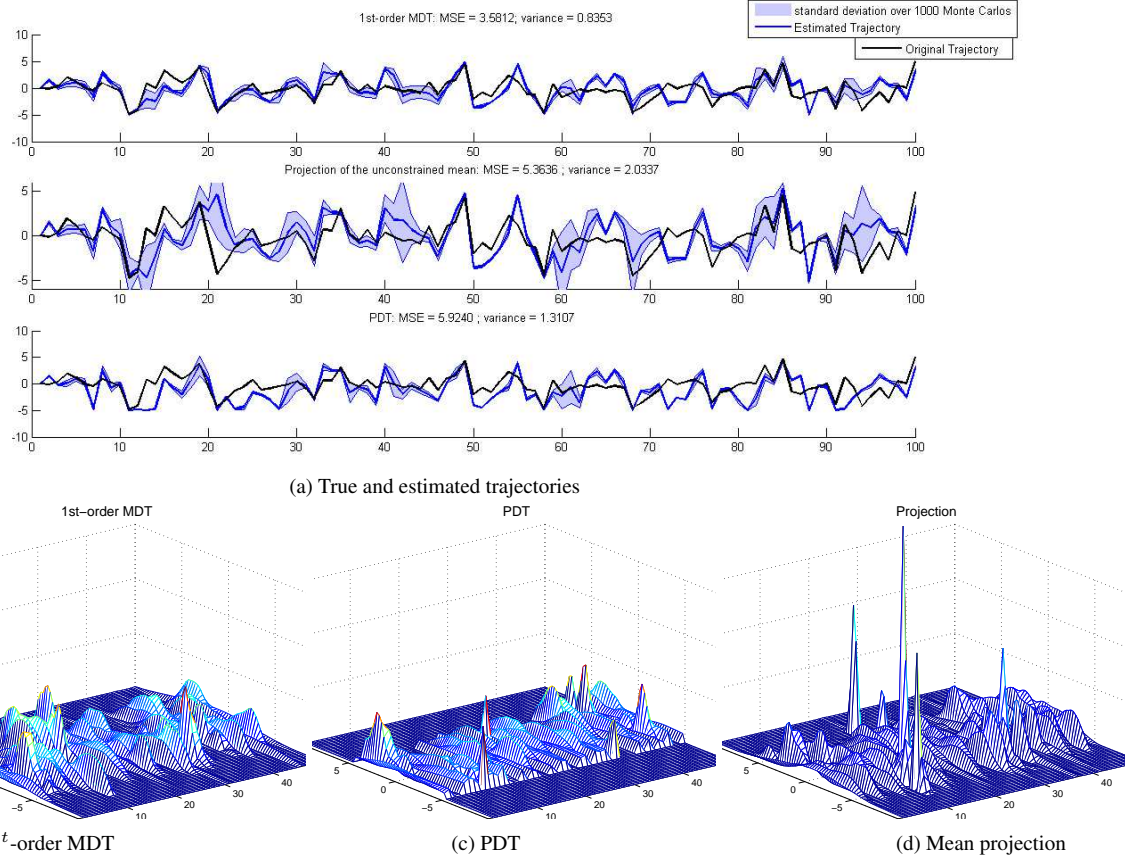


Fig. 2. Constrained state-estimation of the nonlinear dynamic system in (11). (a) State estimation for 1000 Monte Carlo simulations. Shading represents a two standard deviations band. Top row: 1st-order MDT (MSE=3.58, $\sigma = 0.83$); middle row: PDT (MSE=5.36, $\sigma = 2.03$); bottom row: projection of the unconstrained mean estimate (MSE=5.92, $\sigma = 1.31$). (b),(c),(d): State posterior densities evolving over time for 1st-order MDT, PDT and projection, respectively.

bounded support at all time points, whereas the MDT and projection approaches result in proper unbounded support densities. Moreover, the PDT and the projection estimation approaches result in multiple spurious peaks within the densities. These large peaks are located mainly at the boundary of the constraining interval. In the PDT approach, these spurious peaks are due to the fact that sampled particles that do not satisfy the constraint are projected onto the boundary, thus creating a significant positive mass at the boundary of the constraint set and a small density mass elsewhere. In other words, in PDT, the density outside of the interval $[-5, 5]$ is projected onto the boundary points. In the mean projection approach, unconstrained particles are sampled from the proposal density (here, the prior). Because of the highly nonlinear nature of the system, the constraint and the poor choice of the proposal, many of these particles are sampled from low probability regions of the actual posterior; thus having low weight. These low weight particles are replaced, during the resampling procedure, by higher weight boundary particles. On the other hand, the 1st-order MDT method does not suffer from the ‘boundary spurious peaks’ problem and estimates smooth (multinomial) densities over time, which result in

more accurate estimation of the conditional mean.

5. CONCLUSION

Arising from physical principles and process restrictions, constraints are commonly encountered in real-world dynamical systems. Therefore, constraints must be taken into account in order to obtain physically meaningful estimation results. In this paper, we considered the particle filter framework for state estimation in nonlinear and non-Gaussian dynamical systems. We argued that constraining all particles is equivalent to constraining the posterior distribution of the state. This may lead either to a stronger condition or to a different (unrelated) condition; both of which result in incorrect estimation of the posterior distribution of the state. We, subsequently, advanced a new approach, MDT, which imposes the desired constraints on the conditional mean estimate without further restricting the posterior density of the state; and hence preserving the convergence properties of the particle filter towards the optimal posterior density of the state. Future research directions include efficient algorithmic implementation of the MDT approach and its variants.

6. REFERENCES

- [1] S J Julier and J J LaViola, "On Kalman filtering with nonlinear equality constraints," *IEEE Transactions on Signal Processing*, vol. 55, no. 6, pp. 2774 – 2784, June 2007.
- [2] D Simon and D L Simon, "Kalman filtering with inequality constraints for turbofan engine health estimation," *IEE Proceedings in Control Theory and Applications*, vol. 153, no. 3, pp. 371 – 378, May 2006.
- [3] P Vachhani, R Rengaswamy, V Gangwal, and S Narasimhan, "Recursive estimation in constrained nonlinear dynamical systems," *AIChE Journal*, vol. 51, no. 3, pp. 946–959, March 2005.
- [4] J. Porcill, "Optimal combination and constraints for geometrical sensor data," *International Journal of Robotics Research - Special Issue on Sensor Data Fusion*, vol. 7, no. 6, pp. 66 – 77, December 1988.
- [5] L S Wang, Y T Chiang, and F R Chang, "Filtering method for nonlinear systems with constraints," *IEE Proceedings in Control Theory and Applications*, vol. 149, no. 6, pp. 525 – 531, November 2002.
- [6] L Xu, X R Li, Z Duan, and J Lan, "Modeling and state estimation for dynamic systems with linear equality constraints," *IEEE Transactions on Signal Processing*, vol. 61, no. 11, pp. 2927 – 2939, June 2013.
- [7] G Rasool, Kamran Iqbal, N Bouaynaya, and G White, "Neural drive estimation using the hypothesis of muscle synergies and the state-constrained Kalman filter," in *International IEEE EMBS Neural Engineering Conference*, November 2013.
- [8] Mark W. Spong, S. Hutchinson, and M. Vidyasagar, *Robot Modeling and Control*, John Wiley and Sons, Inc., 2005.
- [9] D Simon and T L Chia, "Kalman filtering with state equality constraints," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 1, pp. 128 – 136, January 2002.
- [10] Jun S. Liu and Rong Chen, "Sequential Monte Carlo methods for dynamic systems," *Journal of the American Statistical Association*, vol. 93, no. 443, pp. 1032–1044, September 1998.
- [11] Lixin Lang, Wen-Shiang Chen, Bhavik R. Bakshi, Prem K. Goel, and Sridhar Ungarala, "Bayesian estimation via sequential Monte Carlo sampling Constrained dynamic systems," *Automatica*, vol. 43, no. 9, pp. 1615–1622, September 2007.
- [12] Xinguang Shao, Biao Huang, and Jong Min Lee, "Constrained Bayesian state estimation: A comparative study and a new particle filter based approach," *Journal of Process Control*, vol. 20, no. 2, pp. 143–157, 2010.
- [13] Jagadeesan Prakash, Sachin C Patwardhan, and Sirish L Shah, "Constrained state estimation using particle filters," *oral presentation in IFAC-2008*, 2008.
- [14] O. Straka, J. Dunik, and M. Simandl, "Truncated unscented particle filter," in *American Control Conference*, 2011, pp. 1825–1830.
- [15] Jayesh H. Kotecha and P.M. Djuric, "Gaussian particle filtering," *IEEE Transactions on Signal Processing*, vol. 51, no. 10, pp. 2592–2601, 2003.
- [16] Jayesh H. Kotecha and P.M. Djuric, "Gaussian sum particle filtering," *IEEE Transactions on Signal Processing*, vol. 51, no. 10, pp. 2602–2612, 2003.
- [17] Marc-André Beyer and Gunter Reinig, "Constrained particle filtering using gaussian sum approximations," in *Proceedings of the UKACC International Conference on Control*, 2008.
- [18] Sridhar Ungarala, "A direct sampling particle filter from approximate conditional density function supported on constrained state space," *Computers & Chemical Engineering*, vol. 35, no. 6, pp. 1110–1118, June 2011.
- [19] Ondrej Straka, Jindrich Dunik, and Miroslav Simandl, "Truncation nonlinear filters for state estimation with nonlinear inequality constraints," *Automatica*, vol. 48, no. 2, pp. 273–286, February 2012.
- [20] F Papi, M Podt, Y Boers, and G Battistello, "On constraints exploitation for particle filtering based target tracking," in *International Conference on Information Fusion*, July 2012, pp. 455 – 462.
- [21] N J Gordon, D J Salmond, and A F M Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings in Radar and Signal Processing*, vol. 140, no. 2, pp. 107 – 113, April 1993.