8th International IFAC Symposium on
Dynamics and Control of Process Systems
Preprints Vol.3, June 6-8, 2007, Cancún, Mexico

NONLINEAR MODEL PREDICTIVE CONTROL BASED ON SEQUENTIAL MONTE CARLO STATE ESTIMATION

Satyendra K. Botchu * **and Sridhar Ungarala** *,[1]

* *Department of Chemical and Biomedical Engineering*
Cleveland State University
Cleveland, OH 44115, USA

Abstract: Formulating the nonlinear model predictive control (NMPC) problem using nonlinear differential equations has been gaining attention recently, with the promise of improved performance. NMPC requires a knowledge of the states, which are rarely available directly. Hence, the role of a state estimator is crucial to provide state information from noisy process measurements. Earlier attempts to combine variants of the Kalman filter with NMPC met with limited success due to debilitating effects of linearization. Currently, moving horizon estimation (MHE) is the most popular choice since it is seen as a dual to the control problem. However, MHE typically makes simplifying assumptions about the nature of stochastic variables and lacks an efficient recursive formulation. Most importantly, MHE is an optimization burden in addition to the regulation problem to be solved online. We propose using the sequential Monte Carlo (SMC) filter for state estimation in NMPC since it is significantly faster and at least as accurate as MHE. More accurate and fast estimation results in faster control optimization for realtime use of NMPC and improves the performance. In this paper a comparison of NMPC performance is detailed with MHE and SMC state estimation in a nonlinear CSTR simulation study.

Keywords: Model predictive control, Moving horizon estimation, Sequential Monte Carlo.

1. INTRODUCTION

This paper is concerned with the problem of regulation of nonlinear processes. The first part of the nested control problem is a state estimator whose job is to optimally compute current estimates of measured and unmeasured states from a history of inputs and noisy process measurements. The second part is the model predictive control (MPC) regulator, which computes an optimal control signal to drive the estimated states to their respective setpoints. Figure 1 shows the scheme for feedback MPC regulator with state estimation.

Consider the following discrete time process model and measurement equations,

$$x_k = f(x_{k-1}, u_{k-1}) + w_{k-1}, \quad (1)$$

$$y_k = h(x_k) + \nu_k \quad (2)$$

where f and h are nonlinear vector valued functions, x_k is the state vector, u_k is the control vector, y_k is the measurement vector and w_k and ν_k are white Gaussian noise processes distributed according to $N(0, Q)$ and $N(0, R)$ respectively.

[1] E-mail: `s.ungarala@csuohio.edu`

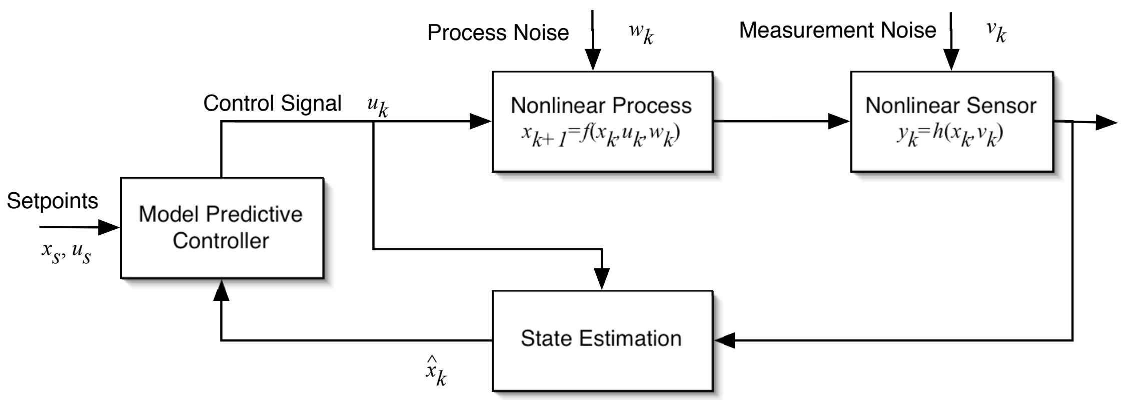

Fig. 1. Closed loop nonlinear model predictive control.

The knowledge about the initial condition is summarized as the estimate $\tilde{x}_0$ with covariance $\tilde{P}_0$.

The Bayesian approach to state estimation is the most general problem statement. However, for nonlinear and non-Gaussian systems, the Bayesian solution is impossible to implement without approximations. Originally the extended Kalman filter (EKF) was used as a state estimator in conjunction with nonlinear model predictive control (NMPC) and it is still popular in industrial applications due to its simplicity (Lee and Ricker, 1994). More recently, moving horizon estimation (MHE) has gained popularity due to its superior estimation properties (Tenny *et al.*, 2004). Both EKF and MHE have been shown as suboptimal Bayesian estimators. The choice of MHE is also motivated by recognizing that MHE is a dual to the MPC control problem. The MHE is a computationally demanding approach, which requires careful tuning for good performance.

In recent years, advances in Monte Carlo techniques revived the interest in implementing the true Bayesian solution without simplifying assumptions. Sequential Monte Carlo (SMC) filter has been shown to be superior to MHE and orders of magnitude faster than the online optimization based methods (Chen *et al.*, 2004). To date SMC studies have been limited to open loop estimation. In this paper, for the first time we demonstrate the efficacy of the SMC filter as an alternative to MHE for closed loop NMPC regulation. Two important points motivate the use of SMC, (1) accurate estimates close to true states can help the control optimization converge faster and (2) fast estimation takes NMPC one step closer to realtime implementation.

In the following sections the estimators are briefly discussed along with NMPC. A CSTR case study compares state estimation accuracy, minimized control cost and speed of execution, all of which favor the SMC over MHE.

2. STATE ESTIMATION

2.1 Sequential Monte Carlo Filter

The Bayesian formulation of state estimation aims to construct the conditional pdf of the state,

$$p(x_k|y_k) \propto p(y_k|x_k)p(x_k|y_{k-1}), \tag{3}$$

where $p(x_k|y_{k-1})$ is *a priori* knowledge, which is modified into the *a posteriori* pdf $p(x_k|y_k)$ with the likelihood function $p(y_k|x_k)$. The prior is computed using the Chapman-Kolmogorov equation,

$$p(x_k|y_{k-1}) = \int p(x_k|x_{k-1})p(x_{k-1}|y_{k-1})\ dx_{k-1} \tag{4}$$

where the transition probability density is obtained from the process model as,

$$p(x_k|x_{k-1}) = \int \delta(x_k - f(x_{k-1}, u_{k-1})) \times p(w_{k-1})\ dw_{k-1}, \tag{5}$$

and the likelihood function is,

$$p(y_k|x_k) = \int \delta(y_k - h(x_k))p(\nu_k)\ d\nu_k. \tag{6}$$

where $p(w_{k-1})$ and $p(\nu_k)$ are the probability density functions of the noise processes. A point estimate such as the mean (minimum variance) or the mode (maximum *a posteriori*) of the posterior pdf is typically taken as the state estimate.

Sequential Monte Carlo methods are recursive algorithms, where N random samples of the states are recursively generated from which summary statistics of $p(x_k|y_k)$ are readily obtained,

$$\int \phi(x_k)p(x_k|y_k) \approx \frac{1}{N}\sum_{i=1}^{N}\phi(x_k^i), \tag{7}$$

where, x_k^i are samples drawn according to $p(x_k|y_k)$. The filter algorithm is summarized here (Gordon *et al.*, 1993):

Algorithm: Sequential Monte Carlo

(1) *Initialize*: start with N samples distributed according to posterior $p(x_{k-1}|y_{k-1})$.
(2) *Predict*: pass each particle through the system model along with samples drawn from noise pdf $p_w(w_{k-1})$, to generate particles distributed according to the prior pdf $p(x_k|y_{k-1})$.
(3) *Update*: using the datum y_k, evaluate the likelihood of each particle and assign weights. Resample N times from the discrete distribution of the weights to generate samples according to the posterior pdf $p(x_k|y_k)$.
(4) *Infer*: compute moments and the associated confidence intervals. Return to step 2

The great advantage of this simple algorithm is that no restrictions are placed on the model or the noise processes, which means linearization is not necessary and non-Gaussian systems are readily tackled. The algorithm is fully recursive and computationally efficient. There is no need for optimization in realtime and the computational burden is largely due only to N computations of the model at each time. In the open loop setting the SMC has been compared with MHE both for estimation accuracy and speed of execution, both of which were shown to favor the SMC (Chen et al., 2004).

It is possible to restrict the samples of the states and noise processes to obey constraints since the ability to impose constraints is one of the primary motivating factors for the use of MPC (Chen et al., 2007). The convergence of the SMC methods has also been well established and extended to constrained systems.

The SMC is a recursive method to provide samples distributed according to the unknown posterior pdf. Typically the mean, the median or the mode (from kernel density estimators) are chosen as the point estimates according to various loss functions. This choice has no influence on the estimator performance unlike in EKF and MHE where the mode is the only choice, which also affects the recursion equations causing stability problems.

2.2 Moving Horizon Estimation

As an alternative to constructing the a posteriori conditional density of the current state and solving for its minimum variance or maximum a posteriori estimate, a conditional joint density may be constructed for a sequence of the discrete state trajectory and then the joint density can be maximized. In this manner the integration in the Chapman-Kolmogorov equation is avoided. However, the optimization problem becomes much larger since the joint probability is a function of the state trajectory instead of one point. From a probabilistic interpretation of moving horizon estimation, the state trajectory of a dynamic system is estimated by solving an optimization problem with the following objective (Haseltine and Rawlings, 2005),

$$\max_{x_{k-H,\dots,\ k}} p(x_{k-H,\dots,\ k}|y_{0,\dots,\ k}), \quad (8)$$

subject to appropriate constraints on the decision variables. The function $p(.)$ is the joint probability of a sequence of discrete trajectory over the most recent time horizon H, conditioned on all available measurements.

In order to establish a tractable convex objective function, which is amenable for local optimization, the MHE problem formulation makes several simplifying assumptions. The MHE problem statement in its most popular format is as follows,

$$\min_{\substack{x_{k-H},\\ \{w\}}} (x_{k-H} - \tilde{x}_{k-H})^T \tilde{P}^{-1}(x_{k-H} - \tilde{x}_{k-H}) + \sum_{j=k-H}^{k-1} w_j^T Q^{-1} w_j + \sum_{j=k-H}^{k} \nu_j^T R^{-1} \nu_j, \quad (9)$$

where $\tilde{x}_{k-H}$ and $\tilde{P}$ summarize the past information represented as the arrival cost. It is common to use inequality constraints on the noise variables. State inequality constraints are also more commonly invoked than equality constraints on the states.

A quadratic form for the arrival cost function is typically desired for ease of implementation using algorithms for solving convex optimization problems. At worst, the arrival cost is not known, corresponding to a uniform prior pdf, and at best it is represented by the first two moments of variants of Gaussians such as truncated Gaussian and half-Gaussian pdfs due to upper and lower bounds imposed on states (Robertson and Lee, 2002). Although one may use any pdf as prior information as long as it generates a convex arrival cost function, a poor choice can easily lead to loss of accuracy and even divergence.

As the horizon is moved forward, the arrival cost must be updated at each time instance for the new horizon. Since the transition probability density is not known, analytical expressions are not possible to recursively update the arrival cost. Generally a filter such as the extended Kalman filter is used based on a linear time-varying process model with additive Gaussian noise terms. The arrival cost is recursively approximated by relying on the shape-invariance of the Gaussian prior pdf subject to linearized dynamics. This is a source of potential instability of the estimator. The problem of accurately summarizing the past information as arrival cost remains an open issue in MHE. The stability can be assured under certain restrictive technical conditions (Rao et al., 2003). In practice, the use of forgetting factors and smoothing over large horizons are reported to improve the estimator's stability properties.

3. MPC REGULATION

The MPC regulation is implemented by a control profile computed from an open loop optimal control problem solved online in a finite horizon. Only the first control action is implemented at a given time instance and a new control profile is obtained the next time instance based on the available state information. The optimal control profile is generally obtained by minimizing the deviations of chosen variables from their setpoints. The objective function for the typical MPC formulation is chosen as a sum of quadratic functions of the form,

$$\min_{u_k,\ldots,u_{k+M-1}} (x_{k+P} - x_{sp})^T Q_c (x_{k+P} - x_{sp}) \sum_{j=0}^{P-1} J_{k+j}(x_{k+j}, u_{k+j}) \quad (10)$$

where M is the control horizon, P is the prediction horizon and,

$$J = (x - x_{sp})^T Q_c (x - x_{sp}) + (u - u_{sp})^T R_c (u - u_{sp}) + \Delta u^T S_c \Delta u \quad (11)$$

where x_{sp} and u_{sp} are the setpoints and Q_c, R_c and S_c are symmetric positive definite penalty matrices. The optimization may be subjected to constraints or bounds on state and control vectors. The prediction of future states is carried out by equation 1 with the assumption that the stochastic variable w_k takes the mean value.

4. SIMULATION EXAMPLE

Consider a continuous stirred tank reactor (CSTR) with an irreversible first order exothermic reaction with the following nonlinear model of the reactor concentration and temperature,

$$\frac{dC}{dt} = \frac{q}{V}(C_f - C) - k_0 C e^{\frac{-E}{RT}} \quad (12)$$

$$\frac{dT}{dt} = \frac{q}{V}(T_0 - T) - \frac{k_0 \Delta H}{\rho c_p} C e^{\frac{-E}{RT}} + \frac{\rho_c c_{pc} q_c}{\rho c_p V}\left(1 - e^{\frac{-hA}{q_c \rho_c c_{pc}}}\right)(T_{c0} - T) \quad (13)$$

The reactor is cooled by a coolant stream q_c and the parameters are shown in Table 1 (Biagiola et al., 2005).

Table 1. CSTR parameters.

Parameter	Value
Process flow rate, q, l/min	100
Feed concentration, C_f, mol/l	1
Feed temperature, T_0, K	350
Inlet coolant temperature, T_{c0}, K	350
Reactor volume, V, l	100
Heat transfer term, hA, cal/min K	7×10^5
Reaction rate constant, k_0, /min	7.2×10^10
Activation energy term, E/R, K	10^4
Heat of reaction, ΔH, cal/mol	-2×10^5
Liquid densities, ρ, ρ_c g/l	1000
Specific heats, c_p, c_{pc}, cal/g K	1

Only the reactor temperature is sampled at $\Delta t = 0.05$ min, which is corrupted by zero mean Gaussian noise with variance $R = 4.5^2$.

4.1 Open Loop Results

The performance of MHE and SMC estimators are compared in the open loop case first. The system is simulated from the initial condition $C = 0$ mol/l and $T = 300$ K. The prior information available to the estimators is the initial estimate $[\tilde{C},\ \tilde{T}]_0^T = [0,\ 300]^T$ and the associated covariance matrix $\tilde{P}_0 = \text{diag}(0.5, 1)$. The discrete time model available to the estimators is of the form,

$$\begin{bmatrix} C \\ T \end{bmatrix}_{k+1} = \begin{bmatrix} C \\ T \end{bmatrix}_k + \begin{bmatrix} dC/dt \\ dT/dt \end{bmatrix} \Delta t + w_k \quad (14)$$

where $w_k \sim N(0, Q)$, with $Q = \text{diag}(0.1^2, 2.5^2)$. The MHE is formulated in a horizon $H = 5$ with the error covariance matrices Q and R and the initial state covariance matrix $\tilde{P}_0$. The SMC filter is initialized with 500 samples. Simulations were performed in MatLab on a 3GHz Intel Linux computer. Average results over twenty realizations are reported in Table 2, which indicate that the SMC filter provides greater accuracy faster than MHE.

Table 2. Average estimation MSE and CPU time.

Estimator	MSE	CPU sec
MHE	1.14	23
SMC	0.16	7

4.2 Closed Loop Results

Given the initial condition $C = 0.06$ mol/l and $T = 449$ K, the objective of the controller is to drive the concentration to setpoint $C_{sp} = 0.1$ mol/l and to keep the manipulated variable at $q_c = 103.411$ l/min, where the corresponding reactor temperature is $T = 438.54$ K. The objective function for control signal optimization is,

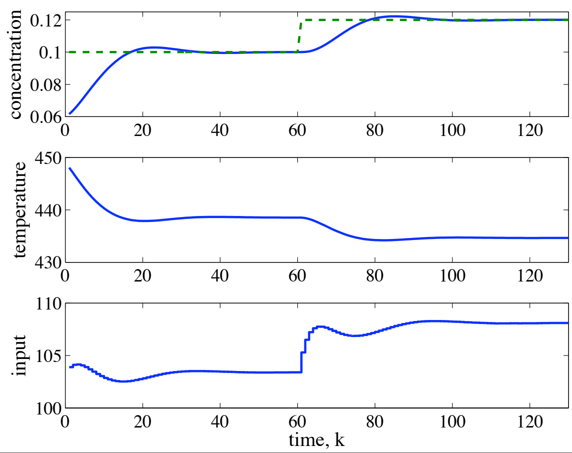

Fig. 2. NMPC reference case

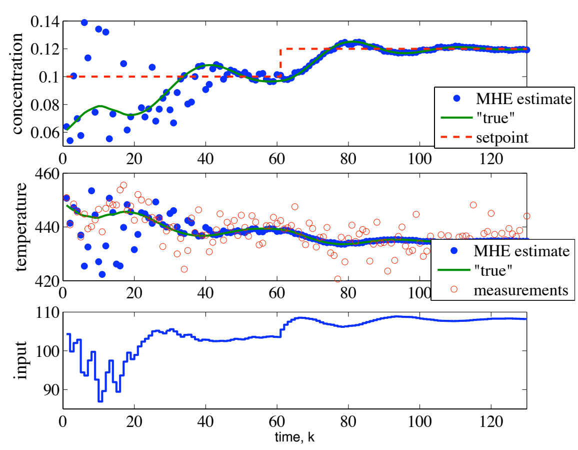

Fig. 3. NMPC based on MHE state estimates

$$J = \sum_{i=1}^{p} 400(C_i - C_{sp})^2 + \sum_{i=1}^{m} 0.005(q_{ci} - q_{csp})^2$$
$$\sum_{i=1}^{m} 0.05(q_{ci} - q_{c(i-1)})^2 \quad (15)$$

where C_{sp} and q_{csp} are the setpoints. A step change is introduced in the setpoint at time instance $k = 60$ with $C_{sp} = 0.12$ mol/l and $q_{csp} = 108.1$ l/min for evaluating setpoint tracking performance. The control horizon is $M = 2$ and the prediction horizon is $P = 5$. The state estimators estimate both the concentration and temperature to pass to the control signal optimizer. The prior information available to the estimators is the initial estimate $[\tilde{C},\ \tilde{T}]_0^T = [0.06,\ 449]^T$ and the associated covariance matrix $\tilde{P}_0 = \mathrm{diag}(0.1^2, 1)$. The mean of the prior pdf is the same as the initial condition, which is not an unreasonable assumption because control loops are almost never closed without a reasonable certainty about the initial condition.

A reference cumulative cost is computed for the ideal case of perfect temperature measurements and noiseless process model. The cumulative cost is simply the sum of the minimized costs at each stage for the length of the simulation, $J_{\text{cumu}}^{\text{ref}} = 6.0883$. Figure 2 shows the states and the control signal for the reference case. If the states are perfectly known, this is the best control profile.

Figure 3 shows setpoint tracking by the controller based on state estimates provided by MHE. The estimator shows slow convergence initially but tracks the true states successfully later. When the estimates are far from true values it is important to note that, (1) the NMPC optimizer takes longer to minimize the cost function and (2) the minimized cost is far from the reference cost computed with true states.

Figure 4 illustrate the superior performance of NMPC based on SMC state estimation. Since the estimates track the true states more closely than MHE, the optimized control signal is closer to the reference case. Average results over twenty realizations are reported in Table 3. The SMC based NMPC runs noticeably faster with smaller estimation error and results in a minimized cost closer to the reference cumulative cost of 6.0883.

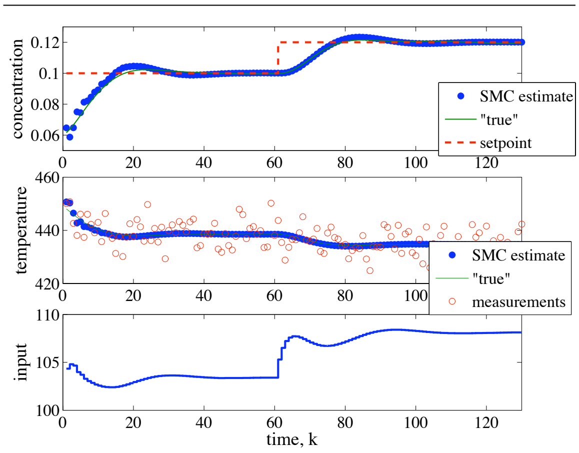

Fig. 4. NMPC based on SMC state estimates

Table 3. Average cumulative cost, estimation MSE and CPU time.

Estimator	$J_{\text{cumu}}^{\text{Estimator}}$	MSE	CPU sec
MHE	88	9.5	346
SMC	7.3	0.5	228

5. CONCLUSIONS

In this paper we propose to use the sequential Monte Carlo (SMC) filter as a true Bayesian state estimator for nonlinear model predictive control (NMPC). In comparison with the currently used moving horizon estimation (MHE), the SMC filter is more accurate and runs significantly faster than MHE. The SMC filter does not use simplifying assumptions about the process model and no real-time optimization is necessary. The results of the CSTR case study indicate that SMC is a viable state estimator for NMPC for realtime implementation. Current work is focused on formulating efficient SMC filter for cases where prior information about initial condition are poorly known, which may lead to poor convergence. Future research will be focused on disturbance estimation using SMC.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant No. CTS-0522864.

REFERENCES

Biagiola, S. I., J. A. Solsona and L. L. Figueroa (2005). Use of state estimation for inferential nonlinear MPC: a case study. Chem. Eng. J **106**, 13–24.

Chen, W. S., B. R. Bakshi, P. K. Goel and S. Ungarala (2004). Bayesian estimation via sequential monte carlo sampling: unconstrained nonlinear dynamic systems. Ind. Eng. Chem. Res **43**, 4012–4025.

Chen, W. S., L. Lang, B. R. Bakshi, P. K. Goel and S. Ungarala (2007). Bayesian estimation via sequential Monte Carlo sampling - constrained dynamic systems. Automatica.

Gordon, N., D. J. Salmond and A. F. M. Smith (1993). Novel approach to onlinear/on-Gaussian Bayesian state estimation. IEE Proc. F **140**(2), 107–113.

Haseltine, E. L. and J. B. Rawlings (2005). Critical evaluation of extended Kalman filtering and moving horizon estimation. Ind. Eng. Chem. Res. **44**, 2451–2460.

Lee, J. H. and N. L. Ricker (1994). Extended Kalman filter based nonlinear model predictive control. Ind. Eng. Chem. Res. **33**, 1530–1541.

Rao, C. V., J. B. Rawlings and D. Q. Mayne (2003). Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations. IEEE Trans. Automatic Control **48**(2), 246–258.

Robertson, D. G. and J. H. Lee (2002). On the use of constraints in least squares estimation and control. Automatica **38**, 1113–1123.

Tenny, M. J., J. B. Rawlings and S. J. Wright (2004). Closed-loop behavior of nonlinear model predictive control. AIChE J. **50**(9), 2142–2154.