TWMCC ⋆ Texas-Wisconsin Modeling and Control Consortium

Technical report number 2002-03

A Critical Evaluation of Extended Kalman Filtering and Moving Horizon Estimation

Eric L. Haseltine and James B. Rawlings

March 12, 2003

Abstract

The goal of state estimation is to reconstruct the state of a system from process measurements and a model. State estimators for most physical processes often must address many different challenges, including nonlinear dynamics, states subject to hard constraints (e.g. nonnegative concentrations), and local optima. In this article, we compare the performance of two such estimators: the extended Kalman filter (EKF) and moving horizon estimation (MHE). We illustrate conditions that lead to estimation failure in the EKF when there is no plant-model mismatch and demonstrate such failure via several simple examples. We then examine the role that constraints, the arrival cost, and the type of optimization (global versus local) play in determining how MHE performs on these examples. In each example, the two estimators are given exactly the same information, namely tuning parameters, model, and measurements; yet MHE consistently provides improved state estimation and greater robustness to both poor guesses of the initial state and tuning parameters in comparison to the EKF.

1 Introduction

It is well established that the Kalman filter is the optimal state estimator for unconstrained, linear systems subject to normally distributed state and measurement noise. Many physical systems, however, exhibit nonlinear dynamics and have states subject to hard constraints, such as nonnegative concentrations or pressures. Hence Kalman filtering is no longer directly applicable. As a result, many different types of nonlinear state estimators have been proposed, such as extended Kalman filters, moving horizon

estimation, model inversion, and Bayesian estimation. Soroush reviews current nonlinear state estimation techniques [16]. Chen et al. present a new method for Bayesian maximum likelihood estimation [15]. Of these methods, extended Kalman filtering has garnered the most interest due to its relative simplicity and demonstrated efficacy in handling nonlinear systems. Examples of implementations include estimation for the production of silicon/germanium alloy films [6], polymerization reactions [7], and fermentation processes [5]. However, the extended Kalman filter, or EKF, is at best an *ad hoc* solution to a difficult problem, and hence there exist many barriers to the practical implementation of EKFs (see, for example, Wilson et al. [21]). Some of these problems include the inability to accurately incorporate physical state constraints and poor use of the nonlinear model. In order to overcome these problems, we propose the use of moving horizon estimation (MHE) as a computationally feasible online solution for state estimation. MHE is an online optimization strategy that accurately employs the nonlinear model and incorporates constraints into the optimization. In this paper, we first outline the basics of nonlinear observability, extended Kalman filtering, and moving horizon estimation. We then present several motivating chemical engineering examples in which the accurate incorporation of both state constraints and the nonlinear model are paramount for obtaining accurate state estimates.

2 Formulation of the Estimation Problem

In chemical engineering systems, most processes consist of continuous processes with discrete measurements. In general, one derives a first principles model by assuming that the continuous process is deterministic, and then one uses Bayesian estimation to estimate the model parameters from process measurements. This model is equivalent to:

$$x_{k+1} = \bar{F}(x_k, u_k, \theta) \tag{1a}$$

$$y_k = h(x_k) + v_k \tag{1b}$$

in which

- x_k is the state of the system at time t_k,
- u_k is the system input at time t_k (assumes a zero order hold over the interval $[t_k, t_{k+1})$),
- θ is the system parameters,
- y_k is the system measurement at time t_k, and
- v_k is a $\mathcal{N}(0, R_k)$ noise [1].

[1] $\mathcal{N}(m, P)$ denotes a normal distribution with mean m and covariance P.

The function $\bar{F}(x_k, u_k, \theta)$ is commonly the solution to a set of differential-algebraic equations. If the measurement noises (v_k's) are assumed normally distributed, determining the optimal parameter estimates corresponds to a weighted least squares optimization of the measurement residuals (i.e. $y_k - h(x_k)$) with respect to the model parameters θ.

In contrast to equation (1), many recent models permit random disturbances to affect the model propagation step. Parameter estimation for nonlinear variations of such models is a subject of on-going research. For this work, we choose the discrete stochastic system model

$$x_{k+1} = F(x_k, u_k) + G(x_k, u_k)w_k \tag{2a}$$

$$y_k = h(x_k) + v_k \tag{2b}$$

in which

- w_k is a $\mathcal{N}(0, Q_k)$ noise,
- $F(x_k, u_k)$ is the solution to a first principles, differential equation model, and
- $G(x_k, u_k)$ yields a matrix with full column rank.

We believe that by appropriately choosing both a first principles model and a noise structure, we can identify both the model parameters and the state and measurement noise covariance structures. Such identification will proceed iteratively as follows:

1. Assuming a noise structure, identify the model parameters.
2. Assuming a model, model parameters, and a noise structure, identify the covariance structures.

This identification procedure is an area of current research, but we maintain that such a procedure will yield a rough, yet useful stochastic model from the system measurements.

Ideally, state estimators should solve the problem

$$x_T^+ = \arg\max_{x_T} p(x_T | y_0, \ldots, y_T) \tag{3}$$

in which $p(x_T|y_0, \ldots, y_T)$ is the probability that the state of the system is x_T given measurements $y_0, \ldots, y_T$. Equation (3) is referred to as the maximum likelihood estimate. In the special case that the system is not constrained and in equation (2)

1. $F(x_k, u_k)$ is linear with respect to x_k,
2. $h(x_k)$ is linear with respect to x_k, and
3. $G(x_k, u_k)$ is a constant matrix,

the maximum likelihood estimator is the Kalman filter. The Kalman filter is a recursive estimator whose form is conducive for online implementation. For the more general formulation given by equation (2), online solution of the exact maximum likelihood estimate is impractical, and approximations must be used to obtain state estimates in real time.

3 Nonlinear Observability

The determination of observability for nonlinear systems such as equation (2) is substantially more difficult than for linear systems. For linear systems, either one state is the optimal estimate, or infinitely many states are optimal estimates, in which case the system is unobservable. Nonlinear systems have the additional complication that *finitely many* states may be *locally* optimal estimates. Definitions of nonlinear observability should account for such a condition. Concepts such as output-to-state stability [20] offer promise for a rigorous mathematical definition of nonlinear observability, but currently no easily implemented tests for such determination exist. In lay terms, such a definition for deterministic models should roughly correspond to "for the given model and measurements, if the measurement data are close, the initial conditions generating the measurements are close."

One approximate method of checking nonlinear observability is to examine the time-varying Gramian [3]. This test actually establishes the observability criterion for linear, time-varying systems. By approximating nonlinear systems as linear time-varying systems, we can obtain a rough estimate of the degree of observability for the system by checking the condition number of the time-varying Gramian. In general, ill-conditioned Gramians indicate poor observability because different initial conditions can reconstruct the data arbitrarily closely [6].

4 Extended Kalman Filtering

The extended Kalman filter is one approximation for calculating equation (3). The EKF linearizes nonlinear systems, then applies the Kalman filter (the optimal, unconstrained, linear state estimator) to obtain the state estimates. The tacit approximation here is that the process statistics are multivariate normal distributions. We summarize the algorithm for implementing the EKF presented by Stengel [17], employing the following notation:

- $E[\alpha]$ denotes the expectation of α,
- A_k denotes the value of the function A at time t_k,
- $x_{k|l}$ refers to the value of x at time t_k given measurements up to time t_l,
- $\hat{x}$ denotes the estimate of x, and

- $\bar{x}_0$ denotes the *a priori* estimate of x_0, that is, the estimate of x_0 with knowledge of no measurements.

The assumed prior knowledge is identical to that of the Kalman filter:

$$\bar{x}_0 \text{ given} \tag{4a}$$
$$P_0 = E[(x - \bar{x}_0)(x - \bar{x}_0)^T] \tag{4b}$$
$$R_k = E[v_k v_k^T] \tag{4c}$$
$$Q_k = E[w_k w_k^T] \tag{4d}$$

The inputs u_k are also assumed to be known.

The approximation uses the following linearized portions of equation (2)

$$A_k = \left.\frac{\partial F(x,u)}{\partial x^T}\right|_{x=x_k, u=u_k} \tag{5}$$
$$C_k = \left.\frac{\partial h(x)}{\partial x^T}\right|_{x=x_k} \tag{6}$$

to implement the following algorithm:

1. At each measurement time, compute the filter gain L and update the state estimate and covariance matrix:

$$L_k = P_{k|k-1} C_k^T [C_k P_{k|k-1} C_k^T + R_k]^{-1} \tag{7}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + L_k(y_k - h(\hat{x}_{k|k-1})) \tag{8}$$
$$P_{k|k} = P_{k|k-1} - L_k C_k P_{k|k-1} \tag{9}$$

2. Propagate the state estimate and covariance matrix to the next measurement time via the equations:

$$\hat{x}_{k+1|k} = F(\hat{x}_k, u_k) \tag{10}$$
$$P_{k+1|k} = A_k P_{k|k} A_k^T + G_k Q_k G_k^T \tag{11}$$

3. Let $k \leftarrow k + 1$. Return to step 1.

Until recently, few properties regarding the stability and convergence of the EKF have been proven. Recent publications present bounded estimation error and exponential convergence arguments for the continuous and discrete EKF forms given detectability, small initial estimation error, small noise terms, and perfect correspondence between the plant and the model [12, 13, 14]. However, depending on the system, the bounds on initial estimation error and noise terms may be unreasonably small. Also, initial estimation error may result in bounded estimate error but not exponential convergence, as illustrated by Chaves and Sontag [2].

5 Moving Horizon Estimation

This section briefly recaps the points made by Rao et al. [9]. One alternative to solving the maximum likelihood estimate is to maximize a joint probability for a trajectory of state values, i.e.,

$$\left\{x_0^*, \ldots, x_T^*\right\} = \arg \max_{x_0,\ldots,x_T} p(x_0, \ldots, x_T | y_0, \ldots, y_T) \tag{12}$$

For unconstrained, linear systems, maximizing the joint probability given in equation (12) yields the equivalent state estimate as maximizing the desired marginal distribution of equation (3):

$$x_T^+ = x_T^* \tag{13}$$

For nonlinear systems or systems with constraints, equation (13) does not hold.

Computationally, it is easiest to consider the logarithmic transformation of equation (12)

$$\arg \min_{x_0,\ldots,x_T} -\log p(x_0, \ldots, x_T | y_0, \ldots, y_T) = \arg \max_{x_0,\ldots,x_T} p(x_0, \ldots, x_T | y_0, \ldots, y_T) \tag{14}$$

because, by assuming that the *a priori* state estimate is a $\mathcal{N}(\bar{x}_0, \Pi_0)$ distributed noise, minimization (14) subject to the nonlinear model (2) gives rise to a least-squares optimization:[2]

$$\Phi_T = \min_{x_0,\ldots,x_T} \Gamma(x_0) + \sum_{k=0}^{T-1} w_k^T Q_k^{-1} w_k + \sum_{k=0}^{T} v_k^T R_k^{-1} v_k \tag{15a}$$

$$\text{s.t.: } \Gamma(x_0) = (x_0 - \bar{x}_0)^T \Pi^{-1} (x_0 - \bar{x}_0) \tag{15b}$$

$$x_{k+1} = F(x_k, u_k) + G(x_k, u_k) w_k \tag{15c}$$

$$y_k = h(x_k) + v_k \tag{15d}$$

Equation (15) is known as the full information problem. Up to here, we have carefully maintained correspondence with the probabilistic interpretation of our problem. One salient feature of virtually all physical models necessitates a departure from this interpretation: constraints. Physically motivated constraints include, for example, nonnegative concentrations, saturation conditions, etc. Using optimization (15) as a suitable starting place, we propose to estimate the state via

$$\Phi_T = \min_{x_0,\ldots,x_T} \Gamma(x_0) + \sum_{k=0}^{T-1} w_k^T Q_k^{-1} w_k + \sum_{k=0}^{T} v_k^T R_k^{-1} v_k \tag{16a}$$

[2] See Appendix 10.3 for details.

$$\text{s.t.: } \Gamma(x_0) = (x_0 - \bar{x}_0)^T \Pi^{-1} (x_0 - \bar{x}_0) \tag{16b}$$
$$x_{k+1} = F(x_k, u_k) + G(x_k, u_k) w_k \tag{16c}$$
$$y_k = h(x_k) + v_k \tag{16d}$$
$$x \in \mathbf{X}, w \in \mathbf{W}, v \in \mathbf{V} \tag{16e}$$

in which the sets **X**, **W**, and **V** contain all feasible values of the system state, state disturbances, and measurement disturbances, respectively. As more data come online, this problem increases in size. One way to overcome computational limitations is to reformulate the problem over a fixed-size estimation horizon, N. This estimation technique is known as moving horizon estimation, or MHE. The constrained MHE optimization is:

$$\Phi_T = \min_{x_{T-N+1},\ldots,x_T} \bar{\Phi}_{T-N} + \sum_{k=T-N+1}^{T-1} w_k^T Q_k^{-1} w_k + \sum_{k=T-N+1}^{T} v_k^T R_k^{-1} v_k \tag{17a}$$

$$\text{s.t.: } x_{k+1} = F(x_k, u_k) + G(x_k, u_k) w_k \tag{17b}$$
$$y_k = h(x_k) + v_k \tag{17c}$$
$$x \in \mathbf{X}, w \in \mathbf{W}, v \in \mathbf{V} \tag{17d}$$

The arrival cost, $\bar{\Phi}_{T-N}$, summarizes the past information up to the observer horizon. Rao et al. [10] explore estimating this cost for constrained linear systems with the corresponding cost for an unconstrained linear system. More specifically, the following two schemes are examined:

1. a "filtering" scheme, in which the optimization accounts for the effects of past data by penalizing deviations of the initial estimate in the horizon from an *a priori* estimate; and

2. a "smoothing" scheme, in which the optimization accounts for the effects of past data by penalizing deviations of the trajectory of states in the estimation horizon from an *a priori* estimate.[3]

For these schemes, MHE is roughly equivalent to maximizing the following probability:

$$\max_{x_{T-N+1},\ldots,x_T} p(x_{T-N+1}, \ldots, x_T | y_0, \ldots, y_T) \tag{18}$$

For unconstrained, linear systems, the MHE optimization (17) collapses to the Kalman filter for both of these schemes. For nonlinear systems, Tenny and Rawlings [19] estimate the arrival cost by approximating the constrained, nonlinear system as an unconstrained, linear time-varying system and applying the corresponding filtering and smoothing schemes. They conclude that the smoothing scheme is superior to the filtering scheme because the filtering scheme induces oscillations in the state estimates due

[3] See Appendix sections 10.1 and 10.2 for derivations of the smoothing and filtering formulations.

to unnecessary propagation of initial error. Here, the tacit assumption is that the probability distribution around the optimal estimate is a multivariate normal. The problem with this assumption is that nonlinear systems may exhibit multiple peaks (i.e. local optima) in this probability distribution. Approximating the arrival cost with either the smoothing or filtering scheme in the presence of multiple local optima will likely skew all future estimates.

Rao [8] further considers several optimal and suboptimal approaches for estimating the arrival cost via a series of optimizations. These approaches stem from the property that, in a deterministic setting (no state or measurement noise), MHE is an asymptotically stable observer as long as the arrival cost is underbounded. One simple way of estimating the arrival cost, therefore, is to implement a uniform prior. Computationally, a uniform prior corresponds to not penalizing deviations of the initial state from the *a priori* estimate; that is, $\hat{\Phi}_{T-N}$ is a constant in optimization (17). The effect of different choices of arrival cost upon the performance of MHE will be illustrated later in this paper.

From a theoretical perspective, Rao et al. showed that MHE is an asymptotically stable observer in a deterministic modeling framework [11, 8]. Furthermore, recent advances in numerical computation have ensured that real-time implementation of MHE strategies for the local optimization of problems such as (17) are possible [18, 19]. We now seek to demonstrate by simulation examples that MHE is a necessary and practical tool for state estimation of chemical process systems.

6 Example 1

Consider the gas-phase, reversible reaction

$$2\mathrm{A} \xrightarrow{\bar{k}} \mathrm{B} \qquad \bar{\mathrm{k}} = 0.16 \tag{19}$$

with stoichiometric matrix

$$\nu = \begin{bmatrix} -2 & 1 \end{bmatrix} \tag{20}$$

and reaction rate

$$r = \bar{k} P_A^2 \tag{21}$$

We define the state and measurement to be

$$x = \begin{bmatrix} P_A \\ P_B \end{bmatrix}, \quad y_k = \begin{bmatrix} 1 & 1 \end{bmatrix} x_k \tag{22}$$

where P_j denotes the partial pressure of species j. We assume that the ideal gas law holds (high temperature, low pressure), and that the reaction occurs in a well-mixed, isothermal batch reactor. From first principles, the model for this system is

$$\dot{x} = f(x) = \nu^T r, \quad x_0 = \begin{bmatrix} 3 & 1 \end{bmatrix}^T \tag{23}$$

For state estimation, consider the following parameters:

$$\Delta t = t_{k+1} - t_k = 0.1, \quad \Pi_0 = \text{diag}(6^2, 6^2), \quad G_k = \text{diag}(1, 1),$$
$$Q_k = \text{diag}(0.001^2, 0.001^2), \quad R_k = 0.1^2, \quad \bar{x}_0 = \begin{bmatrix} 0.1 & 4.5 \end{bmatrix}^T \tag{24}$$

Note that the initial guess, $\bar{x}_0$, is poor. The actual plant experiences $\mathcal{N}(0, Q_k)$ noise in the state and $\mathcal{N}(0, R_k)$ noise in the measurements. We now examine the estimation performance of both the EKF and MHE for this system.

6.1 Comparison of Results

Figure 1 demonstrates that the EKF converges to incorrect estimates of the state (the partial pressures). In addition, the EKF estimates that the partial pressures are negative, which is physically unrealizable. To explain why this phenomenon occurs, we examine the probability density $p(x_k|y_0, \ldots, y_k)$. Recall that the goal of the maximum likelihood estimator is to determine the state x_k that maximizes this probability density. Since we know the statistics of the system, we can calculate this density by successively

1. using the discretized version of the nonlinear model

$$x_{k+1} = F(x_k, w_k) = \begin{bmatrix} \dfrac{x_{k,1}}{2\bar{k}\Delta t x_{k,1} + 1} \\ x_{k,2} + \dfrac{\bar{k}\Delta t x_{k,1}^2}{2\bar{k}\Delta t x_{k,1} + 1} \end{bmatrix} + w_k \tag{25}$$

   to propagate the probability density from $p(x_k|y_0, \ldots, y_k)$ to $p(x_{k+1}|y_0, \ldots, y_k)$ via

$$p(x_{k+1}, w_k|y_0, \ldots, y_k) = p(x_k|y_0, \ldots, y_k)p(w_k) \left| \begin{matrix} \frac{\partial F(x_k, w_k)}{\partial x_k^T} & \frac{\partial F(x_k, w_k)}{\partial w_k^T} \\ \frac{\partial w_k}{\partial x_k^T} & \frac{\partial w_k}{\partial w_k^T} \end{matrix} \right|^{-1} \tag{26}$$

   and then

2. using measurements to update $p(x_k|y_0, \ldots, y_{k-1})$ to $p(x_k|y_0, \ldots, y_k)$

$$p(x_k|y_0, \ldots, y_k) = \frac{p(x_k|y_0, \ldots, y_{k-1})p_{v_k}(y_k - Cx_k)}{\int_{-\infty}^{\infty} p(x_k|y_0, \ldots, y_{k-1})p_{v_k}(y_k - Cx_k)dx_k} \tag{27}$$

Therefore, the expression for the probability density we are interested in is

$$p(x_k|y_0, \ldots, y_k) = \frac{\int_{-\infty}^{\infty} \ldots \int_{-\infty}^{\infty} \Omega_k dw_0 \ldots dw_{k-1}}{\int_{-\infty}^{\infty} \ldots \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Omega_k dw_0 \ldots dw_{k-1} dx_k} \tag{28}$$

in which

$$\Omega_k = \left[\prod_{j=0}^{k-1}\left(2\bar{k}\Delta t x_{j,1}+1\right)^2\right]\exp\left[-\frac{1}{2}\left((x_0-\bar{x}_0)^T\Pi_0^{-1}(x_0-\bar{x}_0)+\sum_{j=0}^{k}v_j^T R^{-1}v_j+\sum_{j=0}^{k-1}w_j^T Q^{-1}w_j\right)\right] \tag{29}$$

We can numerically evaluate equation (28) using the integration package Bayespack [4]. Figure 2 presents a contour plot of the results for $p(x_1|y_0,y_1)$ with transformed axes

$$t=\sqrt{2}\begin{bmatrix}1 & -1\\ 1 & 1\end{bmatrix}^{-1}x$$

This plot clearly illustrates the formation of two peaks in the probability density. However, only one of these peaks corresponds to a region where both the partial pressures for species A and B are positive. The real problem is that the process prohibits negative partial pressures, whereas unconstrained estimators permit updating of the state to regions where partial pressures may be negative. Since the EKF falls into the unconstrained estimator category with a local optimization (at best), the estimation behavior in Figure 1 is best explained as a poor initial guess leading to an errant region of attraction.

One method of preventing negative estimates for the partial pressure is to "clip" the EKF estimates. In this strategy, partial pressures rendered negative by the filter update are zeroed. As seen in Figure 3, this procedure results in an improved estimate in that the EKF eventually converges to the true state, but estimation during the initial dynamic response is poor. Also, only the estimates are "clipped", not the covariance matrix. Thus the accuracy of the approximate covariance matrix is now rather questionable.

Alternatively, we can optimally constrain the partial pressures by applying MHE. Figure 4 presents the MHE results for a horizon length of one time unit ($N = 11$ measurements). These results indicate significant improvement over those of either the EKF or the clipped EKF.

To explore further the differences between the full information and maximum likelihood estimates, we examine contour plots of the projection

$$\max_{x_0,\ldots,x_{k-1}} p(x_0,\ldots,x_k|y_0,\ldots,y_k) \tag{30}$$

noting again the equivalence between this probability and the full information cost function Φ_k given by equation (14). Figure 5 confirms of our previous assertion that the full information and maximum likelihood estimates are not equivalent for nonlinear systems. In fact, the global optima are even different. However, the full information formulation retains the dominant characteristic of the maximum likelihood estimate, namely the formation of two local optima.

6.2 Evaluation of Arrival Cost Strategies

The next logical question is: does MHE retain the same properties as the maximum likelihood estimate? The short answer is: it depends on what approximation one chooses

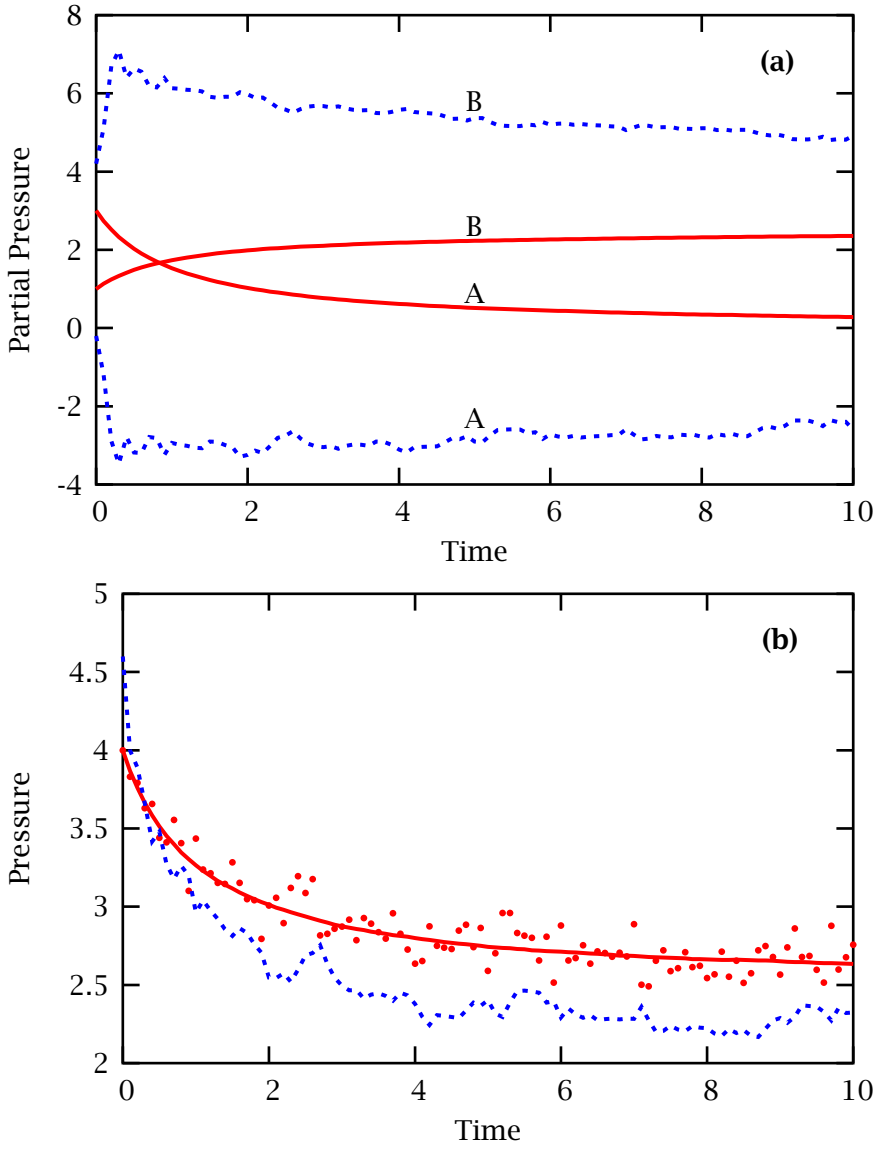

Figure 1: Extended Kalman filter results. (a) plots the evolution of the actual (solid line) and EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and EKF updated (dashed line) pressure estimates.

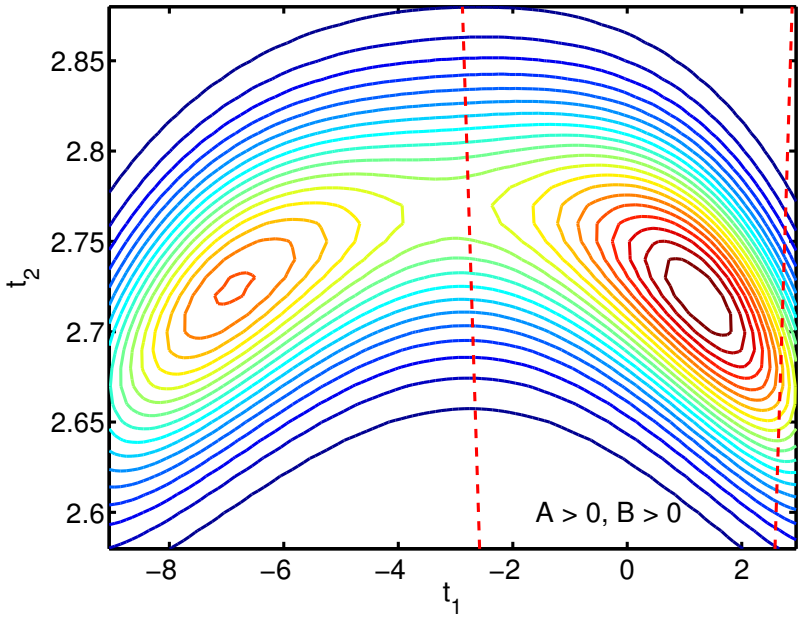

Figure 2: Contours of $p(x_1|y_0, y_1)$

for the arrival cost.

Figures 6 through 8 compare contours of the maximum likelihood estimate, unconstrained MHE with a smoothing update, and unconstrained MHE with a uniform prior, respectively, given five measurements. Figure 7 shows that the smoothing update biases the contours of the state estimate so much that the estimator no longer predicts multiple optima. This biasing occurs because the update has "smoothed" the estimate around only one of the optima in the estimator. Using MHE with a uniform prior, on the other hand, retains the property of multiple optima in the estimator as seen in Figure 8.

Increasing the number of measurements in the estimation horizon can overcome the biasing of the smoothing update. Figure 9 shows the eventual reemergence of multiple optima in the estimator upon increasing the estimation horizon from four (i.e. Figure 7) to ten. However, the optima are still heavily biased by the smoothing update.

We speculate that any approximation of the arrival cost using the assumption that the process is a time-varying linear system may lead to substantial biasing of the estimator. A short estimation horizon further compounds such biasing because the information contained in the data can no longer overcome the prior information (i.e. the arrival cost). This situation is analogous to cases in Bayesian inference when the prior dominates and distorts the information contained in the data [1]. We expect the EKF to demonstrate similar biasing since it is essentially a suboptimal MHE with a short estimation horizon and an arrival cost approximated by a filtering update. For such approximations to work well, one must have a system that does not exhibit multiple

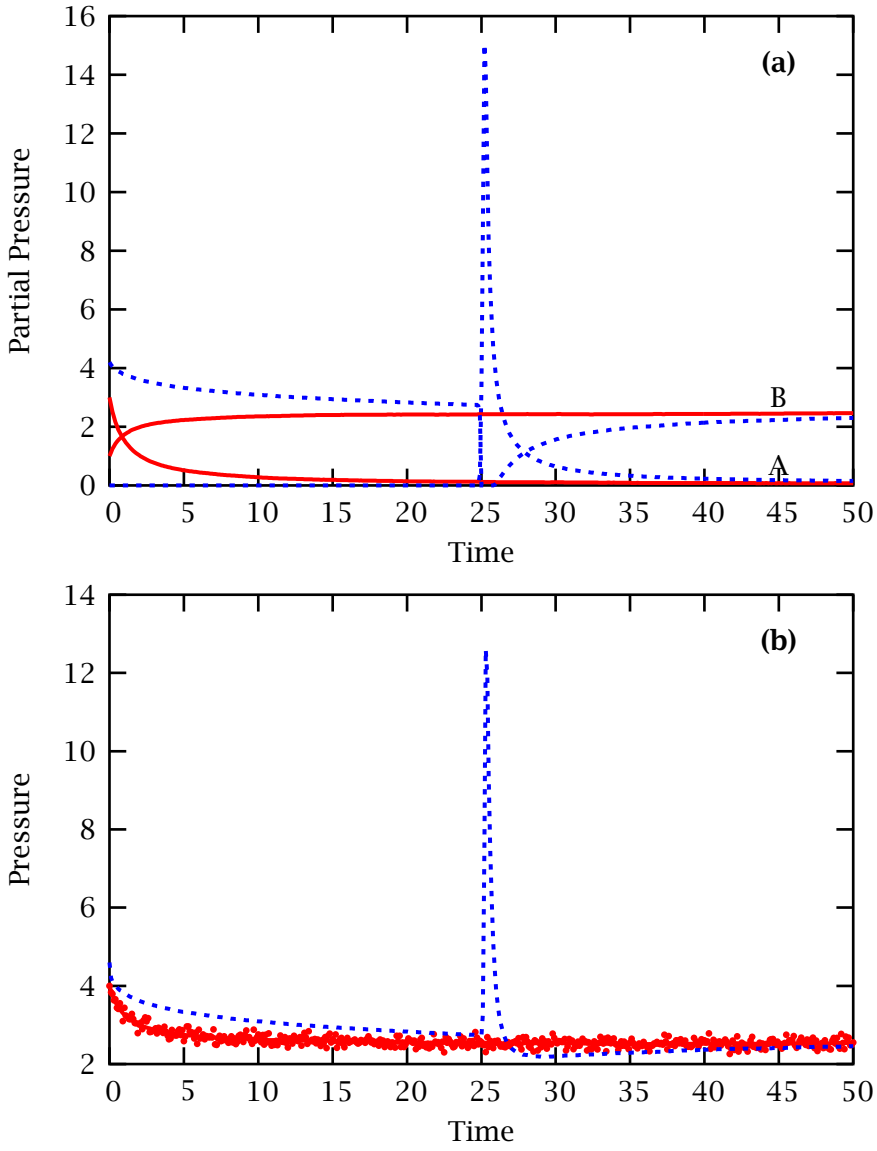

Figure 3: Clipped extended Kalman filter results. (a) plots the evolution of the actual (solid line) and clipped EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and clipped EKF updated (dashed line) pressure estimates.

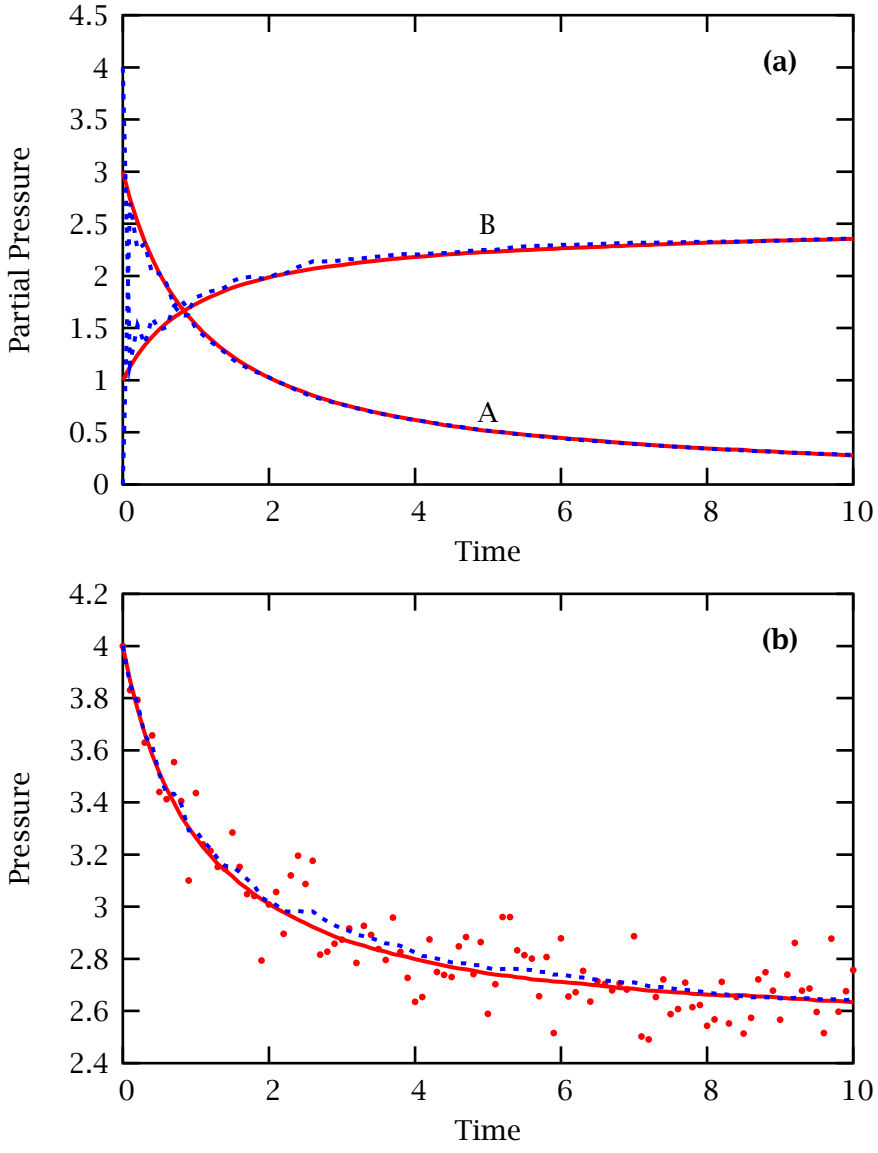

Figure 4: Moving horizon estimation results, states constrained to $x \geq 0$, smoothing initial covariance update, and horizon length of one time unit ($N = 11$ measurements). (a) plots the evolution of the actual (solid line) and MHE updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) pressure estimates.

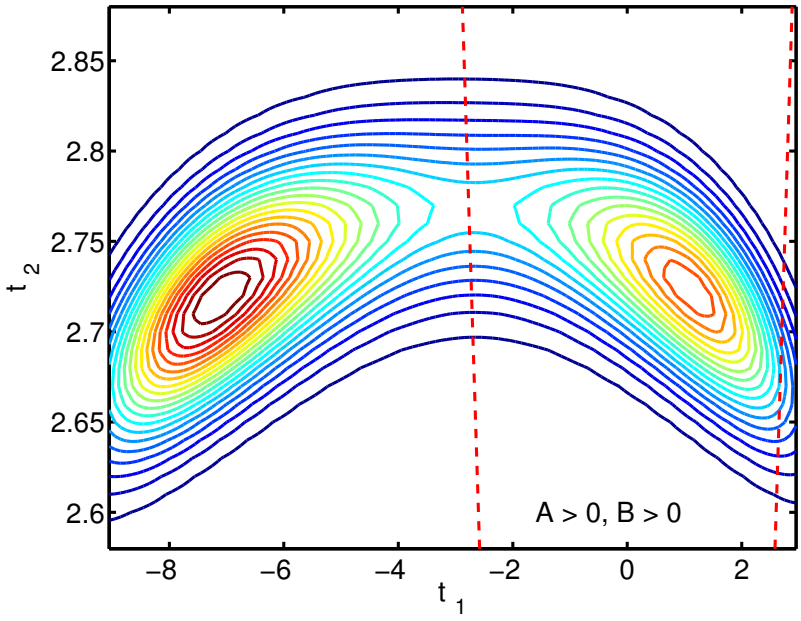

Figure 5: Contours of $\max_{x_0} p(x_1, x_0 | y_0, y_1)$.

local optima in the probability distribution.

The optimization strategy further obfuscates the issue of whether or not to approximate the arrival cost via linearization (e.g. the smoothing and filtering updates). Ideally, one would implement a global optimizer so that MHE could then distinguish between local optima. With global optimization, approximating the arrival cost with a uniform prior and making the estimation horizon reasonably long is preferable to approximating the arrival cost as a multivariate normal because of the observed biasing effect. Currently, though, only local optimization strategies can provide the computational performance required to perform the MHE calculation in real time. For this case, it may be preferable to use a linear approximation of the arrival cost and then judiciously apply constraints to prevent multiple optima in the estimator. The examples considered next examine the estimator performance of this type of MHE.

7 EKF Failure

The results of the preceding example indicate that multiple optima may arise in the estimation problem. In this section, we outline the conditions that generate this phenomenon in two classes of chemical reactors. We then present several examples that demonstrate failure of the EKF as an estimator.

If there is no plant-model mismatch, measurement noise, or state noise, one defini-

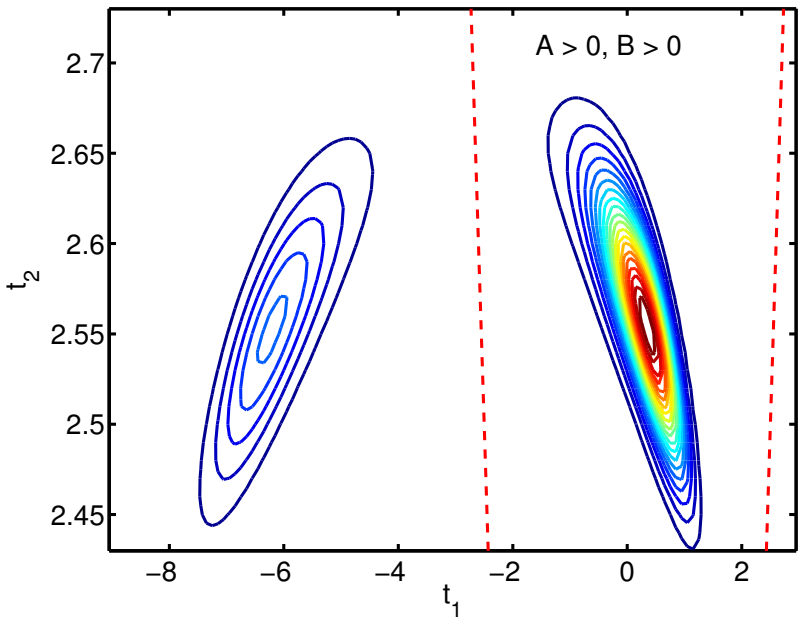

Figure 6: Contours of $p(x_4|y_0,\ldots,y_4)$.

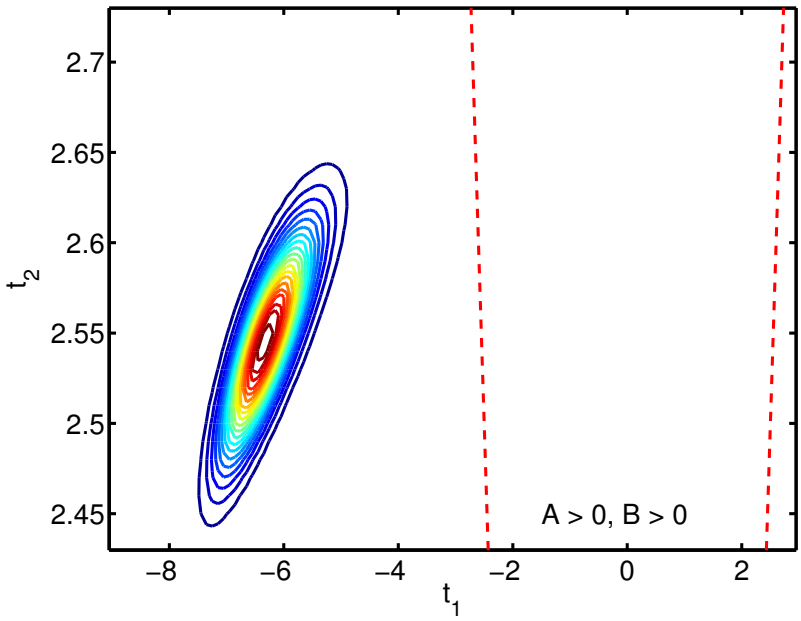

Figure 7: Contours of $\max_{x_1,\ldots,x_3} p(x_1,\ldots,x_4|y_0,\ldots,y_4)$ with the arrival cost approximated using the smoothing update.

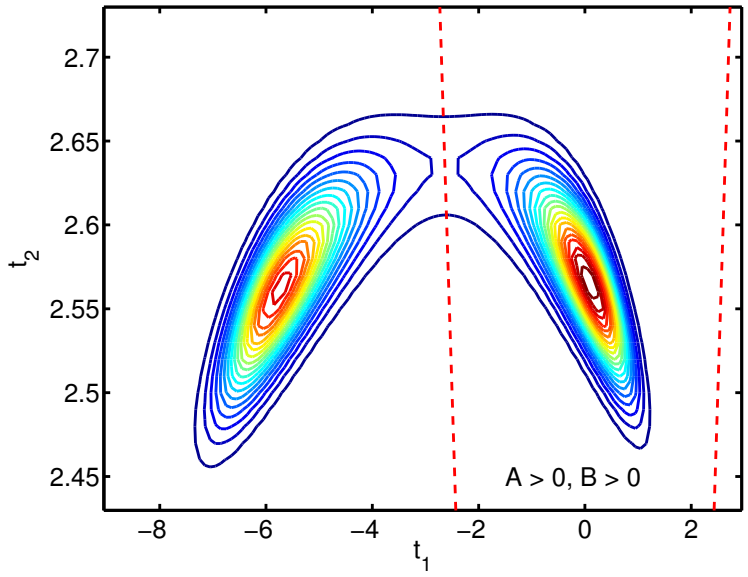

Figure 8: Contours of $\max_{x_1,\ldots,x_3} p(x_1,\ldots,x_4|y_0,\ldots,y_4)$ with the arrival cost approximated as a uniform prior.

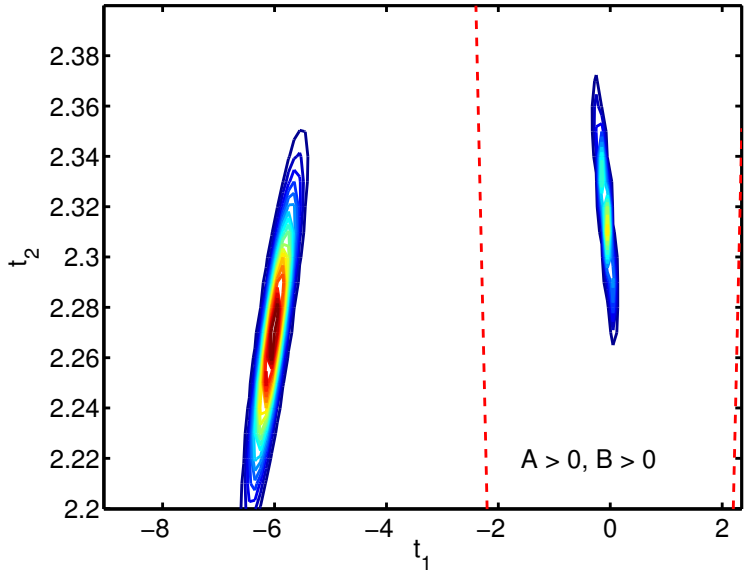

Figure 9: Contours of $\max_{x_1,\ldots,x_9} p(x_1,\ldots,x_{10}|y_0,\ldots,y_{10})$ with the arrival cost approximated using the smoothing update.

tion of estimator failure is

$$\lim_{k\to\infty} |\hat{x}_{k|k} - x_k| > \epsilon \tag{31}$$

for some $\epsilon > 0$ ($|x|$ is a norm of x). That is, the estimator is unable to reconstruct the true state no matter how many measurements it processes. For stable systems, i.e. those systems tending to a steady state, we expect that

$$\hat{x}_{k|k} = \hat{x}_{k-1|k-1} \tag{32}$$

in the same limit as equation (31). We now examine the discrete EKF given such conditions. Recall that the following equations dictate the propagation and update steps:

$$\hat{x}_{k|k-1} = F(\hat{x}_{k-1|k-1}, u_{k-1}) \tag{33a}$$
$$P_{k|k-1} = A_{k-1}P_{k-1|k-1}A_{k-1}^T + G_{k-1}Q_{k-1}G_{k-1}^T \tag{33b}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + L_k(y_k - h(\hat{x}_{k|k-1})) \tag{33c}$$
$$P_{k|k} = P_{k|k-1} - L_k C_k P_{k|k-1} \tag{33d}$$
$$L_k = P_{k|k-1}C_k^T[C_k P_{k|k-1} C_k^T + R_k]^{-1} \tag{33e}$$

At steady state, the following equalities hold:

$$\hat{x}_{k|k} = \hat{x}_{k-1|k-1} \tag{34a}$$
$$P_{k|k} = P_{k-1|k-1} \tag{34b}$$

Combining expressions (33) and (34) yields:

$$0 = F(\hat{x}_{k-1|k-1}, u_{k-1}) - \hat{x}_{k|k-1} \tag{35a}$$
$$0 = A_{k-1}P_{k-1|k-1}A_{k-1}^T + G_{k-1}Q_{k-1}G_{k-1}^T - P_{k|k-1} \tag{35b}$$
$$0 = \hat{x}_{k|k-1} + L_k(y_k - h(\hat{x}_{k|k-1})) - \hat{x}_{k-1|k-1} \tag{35c}$$
$$0 = P_{k|k-1} - L_k C_k P_{k|k-1} - P_{k-1|k-1} \tag{35d}$$
$$L_k = P_{k|k-1}C_k^T[C_k P_{k|k-1} C_k^T + R_k]^{-1} \tag{35e}$$

If both equations (31) and (35) hold, then the EKF has failed as an estimator.

One solution to equation (35) results when multiple steady states satisfy the steady-state measurement. This phenomenon corresponds to the case that

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} = \hat{x}_{k-1|k-1} \tag{36}$$
$$y_k = h(\hat{x}_{k|k-1}) \tag{37}$$
$$\hat{x}_{k|k} \neq x_k \tag{38}$$

We would expect the EKF to fail when

1. the system model and measurement are such that multiple states satisfy the steady-state measurement, and

2. the estimator is given a poor initial guess of the state.

Condition 1 does not imply that the system is unobservable; rather, this condition states that the state cannot be uniquely determined from solely the *steady-state measurement*. For such a case to be observable, the process dynamics must make the system observable. Condition 2 implies that the poor initial guess skews the estimates ($\hat{x}_{k|k}$'s) toward a region of attraction not corresponding to the actual state (x_k's).

For well-mixed systems consisting of reaction networks, the null space of the stoichiometric matrix in combination with the number (and type) of measurements dictate whether or not multiple steady states can satisfy the steady-state measurement. More specifically, the nonlinearity of the system must be present at steady state so that multiple steady states can satisfy the steady-state measurement. Define:

- ν, the stoichiometric matrix of size $r \times s$, in which r is the number of reactions and s is the number of species;
- ρ, the rank of ν ($\rho = r$ if there are no linearly dependent reactions);
- η, the nullity of ν;
- n, the number of measurements; and
- n_m, the number of measurements that can be written as a linear combination of states (e.g. $y = x_1 + x_2$ and $(x_1 + x_2)y = x_1$).

For batch reactors, conservation laws yield a model of the form

$$\frac{d}{dt}(xV_R) = \nu^T r(x) V_R \tag{39}$$

in which

- x is an s-vector containing the concentration of each species in the reactor,
- V_R is the volume of the reactor, and
- $r(x)$ is an r-vector containing the reaction rates.

For this system ρ specifies the number of independent equations at equilibrium. In general, we will require that

1. all reactions are reversible
2. the following inequalities hold:

number of "linear" equations		number of estimated species		number of independent equations
$n_m + \rho$	$<$	s	$\leq$	$n + \rho$

Note that the batch reactor preserves the nonlinearity of the reaction rates in the steady-state calculation. Also, the combination of batch steady-state equations and measurements may or may not be an over-specified problem.

For continuously stirred tank reactors (CSTRs), conservation laws yield a model of the form

$$\frac{d}{dt}(xV_R) = Q_f c_f - Q_o x + \nu^T r(x) V_R \tag{40}$$

where

- x is an s-vector containing the concentration of each species in the reactor,
- V_R is the volume of the reactor,
- Q_f is the volumetric flow rate into the reactor,
- c_f is an s-vector containing the inlet concentrations of each species,
- Q_o is the effluent volumetric flow rate, and
- $r(x)$ is an r-vector containing the reaction rates.

Here η specifies the number of linear algebraic relationships among the s species at equilibrium because the null space represents linear combinations of the material balances that eliminate nonlinear reaction rates. We will require

$$\begin{array}{c}\text{number of ``linear'' equations} \\ n_m + \eta\end{array} \quad < \quad \begin{array}{c}\text{number of estimated species} \\ s\end{array} \tag{41}$$

If equation (41) is an equality instead of an inequality, then determination of the steady state is generally a well-defined, linear problem with a unique solution. Note that the left hand side of equation (41) is actually an upper bound since we could potentially choose a measurement contained within the span of the null space (a linear combination of the null vectors). However, such measurements would be invariant and hence would give no dynamic information. Also, equation (41) is only a necessary condition. EKF failure for CSTRs modeled by equation (40) must be confirmed by verifying that equation (35) holds. This requirement differs from the batch case because in general, the CSTR design equation (40) yields a sufficient number of equations to calculate all possible steady states, whereas the batch design equation (39) does not.

We now examine several examples that illustrate these points.

7.1 Example 2

Consider the gas-phase, reversible reactions

$$\mathrm{A} \underset{k_2}{\overset{k_1}{\rightleftharpoons}} \mathrm{B} + \mathrm{C} \tag{42a}$$

$$2\mathrm{B} \underset{k_4}{\overset{k_3}{\rightleftharpoons}} \mathrm{C} \tag{42b}$$

$$k = \begin{bmatrix} 0.5 & 0.05 & 0.2 & 0.01 \end{bmatrix}^T \tag{42c}$$

with stoichiometric matrix

$$\nu = \begin{bmatrix} -1 & 1 & 1 \\ 0 & -2 & 1 \end{bmatrix} \tag{43}$$

and reaction rates

$$r = \begin{bmatrix} k_1 c_A - k_2 c_B c_C \\ k_3 c_B^2 - k_4 c_C \end{bmatrix} \tag{44}$$

We define the state and measurements to be

$$x = \begin{bmatrix} c_A & c_B & c_C \end{bmatrix}^T \tag{45a}$$

$$y = \begin{bmatrix} RT & RT & RT \end{bmatrix} x \tag{45b}$$

where c_j denotes the concentration of species j, R is the ideal gas constant, and T is the reactor temperature [4]. We assume that the ideal gas law holds (high temperature, low pressure). We consider state estimation for both a batch reactor and a CSTR.

7.1.1 Batch Reactor

From first principles, the model for a well-mixed, constant volume, isothermal batch reactor is

$$\dot{x} = f(x) = \nu^T r \tag{46}$$

$$x_0 = \begin{bmatrix} 0.5 & 0.05 & 0 \end{bmatrix}^T \tag{47}$$

We consider state estimation with the following parameters:

$$\Delta t = t_{k+1} - t_k = 0.25 \tag{48a}$$

$$\Pi_0 = \text{diag}\left(0.5^2, 0.5^2, 0.5^2\right) \tag{48b}$$

$$G_k = \text{diag}\left(1, 1, 1\right) \tag{48c}$$

$$Q_k = \text{diag}\left(0.001^2, 0.001^2, 0.001^2\right) \tag{48d}$$

$$R_k = 0.25^2 \tag{48e}$$

$$\bar{x}_0 = \begin{bmatrix} 0 & 0 & 4 \end{bmatrix}^T \tag{48f}$$

Note that the initial guess, $\bar{x}_0$, is poor. The actual plant experiences $\mathcal{N}(0, Q_k)$ noise in the state and $\mathcal{N}(0, R_k)$ noise in the measurements. We now examine the estimation performance of both the EKF and MHE for this system.

Figure 10 demonstrates that the EKF cannot reconstruct the evolution of the state for this system. In fact, the EKF appears to converge to incorrect steady-state estimates

[4] For the simulations, $RT = 32.84$.

of the state. Table 1 presents the results of solving the equations in (35) for this system. Note that the concentrations of components A and B are negative, indicating that the EKF has converged to an unphysical state estimate. To prevent negative concentrations, we next implement an *ad hoc* clipping strategy in which negative filtered values of the state are set to zero (i.e. if $\hat{x}_{k|k} < 0$, set $\hat{x}_{k|k} = 0$). Figure 11 plots these clipped EKF results. Here, the clipped EKF drives the predicted pressure three orders of magnitude larger than the measured pressure before eventually converging to the actual states. Figure 12 presents the results of applying MHE. For these results, we have constrained the state to prevent estimation of negative concentrations. The figures demonstrate that MHE swiftly converges to the correct state estimates.

A little algebraic analysis reveals that multiple steady states satisfy the steady-state measurement for this system. At steady state, the model and measurement equations yield one linear equation (assuming no noise in the steady-state measurement y_{ss})

$$c_A + c_B + c_C = \frac{y_{ss}}{RT} \tag{49}$$

and two nonlinear equations

$$k_1 c_A = k_{-1} c_B c_C \tag{50}$$

$$k_2 c_B^2 = k_{-2} c_C \tag{51}$$

Solving for the steady-state solution using equations (49)-(51):

$$c_C = \frac{k_2}{k_{-2}} c_B^2 = K_2 c_B^2 \tag{52}$$

$$c_A = \frac{k_{-1} k_2}{k_1 k_{-2}} c_B^3 = \frac{K_2}{K_1} c_B^3 \tag{53}$$

$$0 = \frac{K_2}{K_1} c_B^3 + K_2 c_B^2 + c_B - \frac{y_{ss}}{RT} \tag{54}$$

Descartes' rule of signs states that for polynomials with real coefficients, the number of positive, real roots is either the number of sign changes between consecutive coefficients or two less than this number. Since equilibrium constants and the steady-state measurement are positive, equation (54) has at most one positive root. Thus there is only one physically realizable steady state. MHE is a natural estimation tool for this system since its incorporation of constraints can thus prevent the estimator from converging to unphysical steady states.

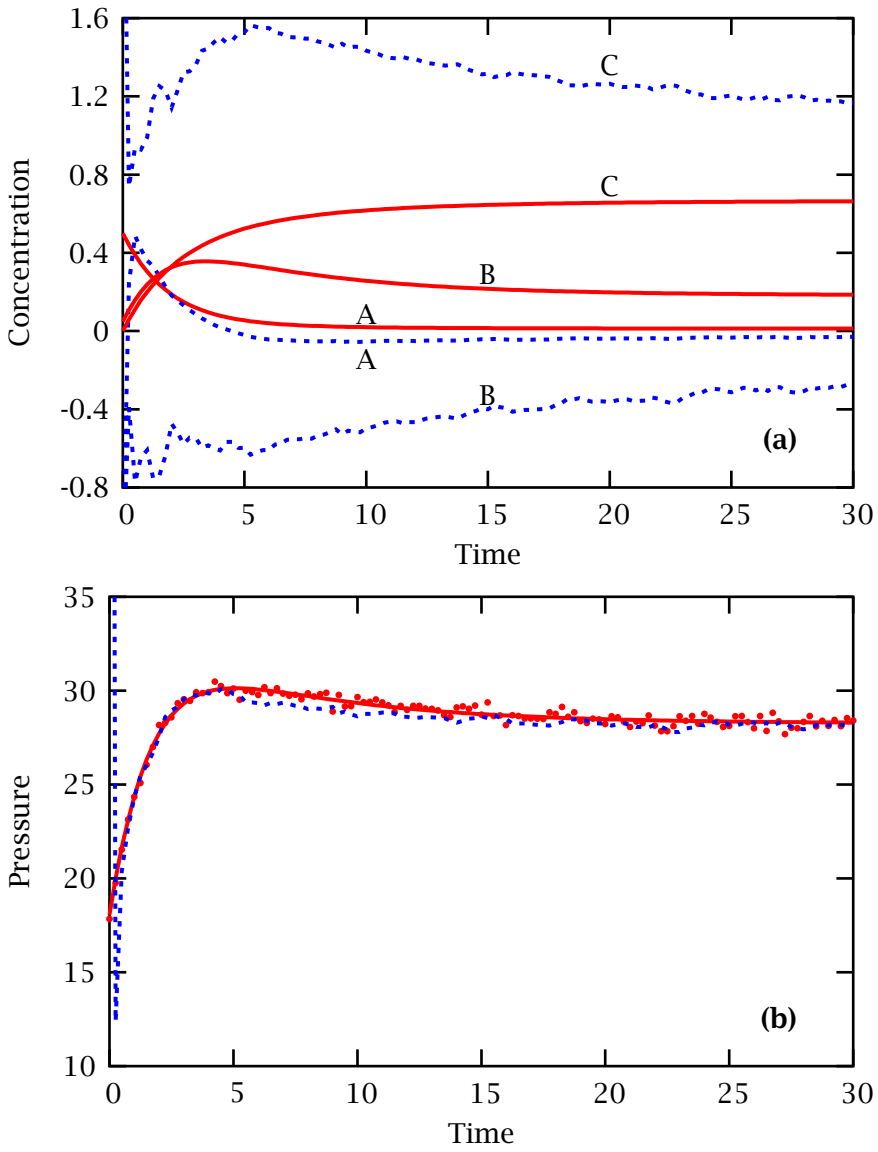

Figure 10: Extended Kalman filter results. (a) plots the evolution of the actual (solid line) and EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and EKF updated (dashed line) pressure estimates.

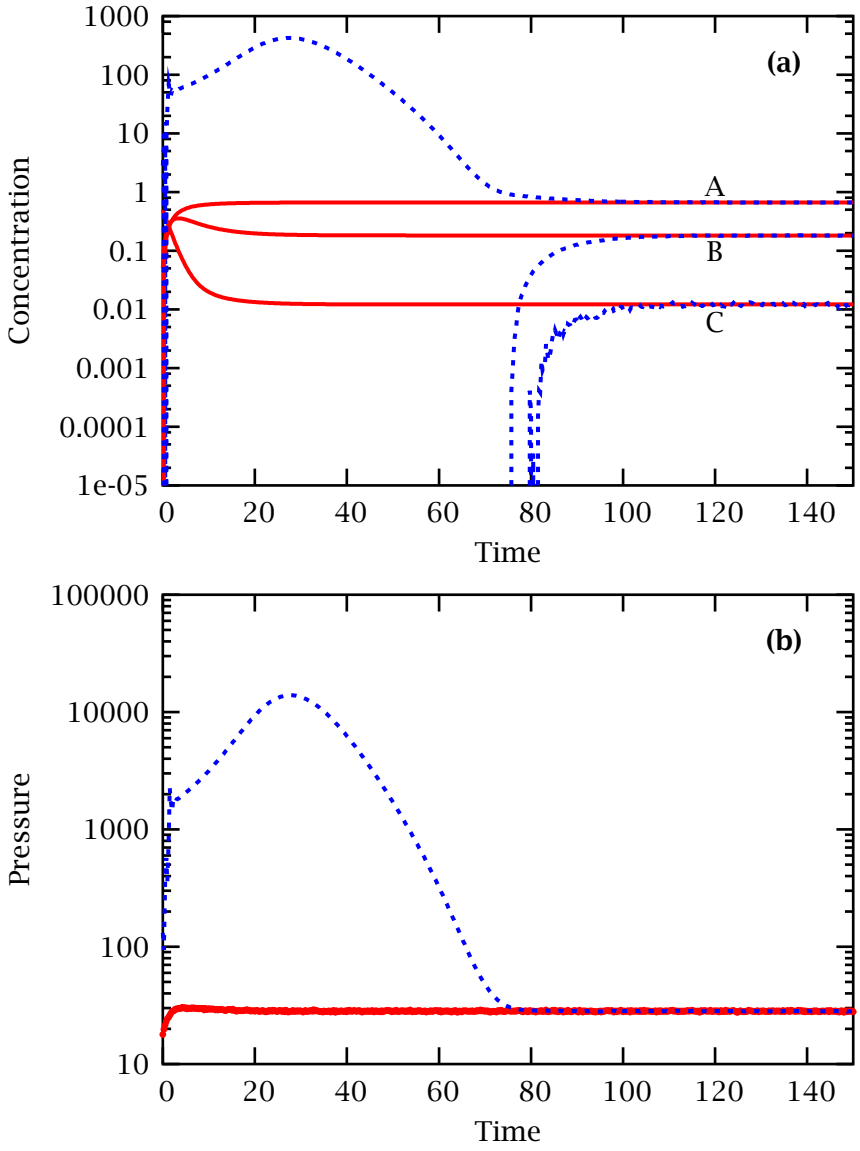

Figure 11: Clipped extended Kalman filter results. (a) plots the evolution of the actual (solid line) and clipped EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and clipped EKF updated (dashed line) pressure estimates.

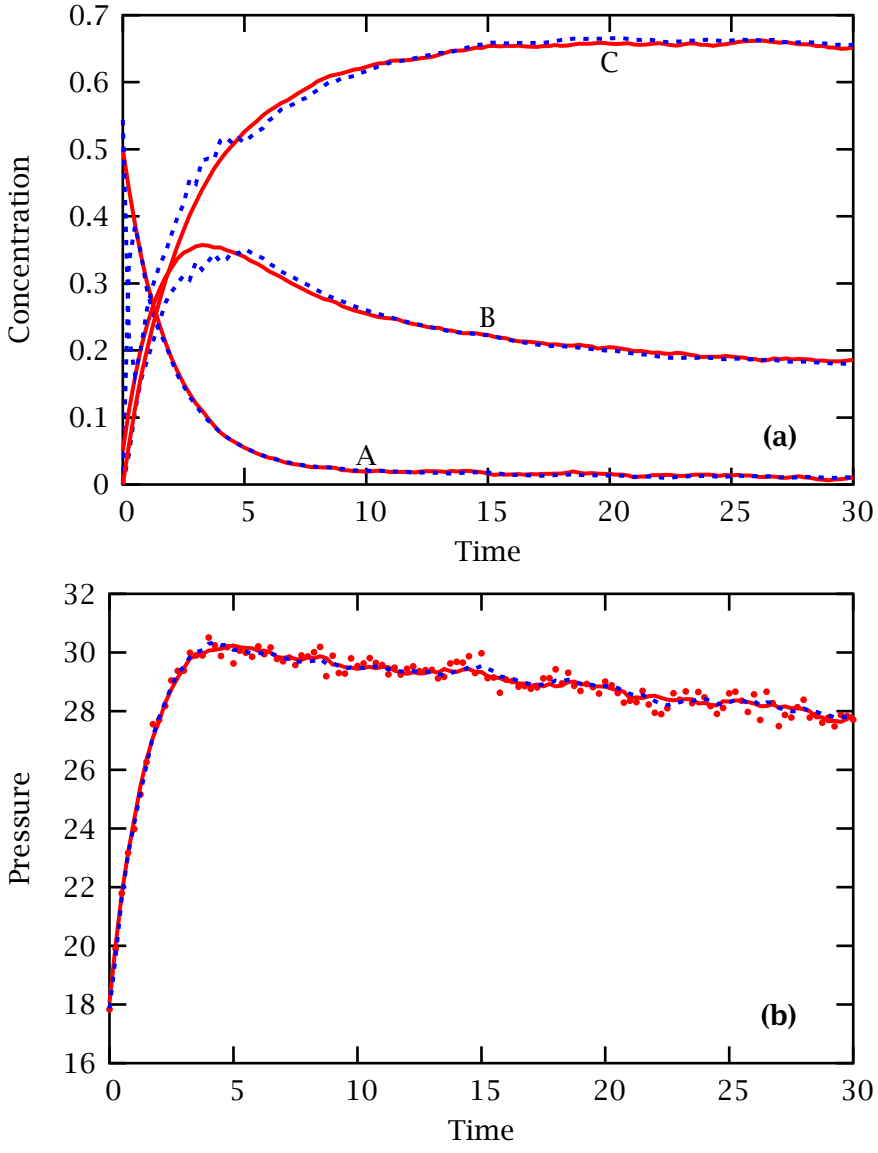

Figure 12: Moving horizon estimation results, states constrained to $x \geq 0$, smoothing initial covariance update, and horizon length of 2.5 time units ($N = 11$ measurements). (a) plots the evolution of the actual (solid line) and MHE updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) pressure estimates.

Component	Predicted EKF Steady State	Actual Steady State
A	−0.0274	0.01241
B	−0.2393	0.1837
C	1.1450	0.6753

Table 1: EKF steady-state behavior, no measurement or state noise

7.1.2 CSTR

From first principles, the model for a well-mixed, isothermal CSTR reactor is

$$\dot{x} = \frac{Q_f}{V_R} c_f - \frac{Q_o}{V_R} x + \nu^T r \tag{55}$$

$$c_f = \begin{bmatrix} 0.5 & 0.05 & 0 \end{bmatrix}^T \tag{56}$$

$$x_0 = \begin{bmatrix} 0.5 & 0.05 & 0 \end{bmatrix}^T \tag{57}$$

$$Q_f = Q_o = 1 \tag{58}$$

$$V_R = 100 \tag{59}$$

We consider state estimation with the following measurement and parameters:

$$y_k = \begin{bmatrix} RT & RT & RT \end{bmatrix} x_k \tag{60a}$$

$$\Delta t = t_{k+1} - t_k = 0.25 \tag{60b}$$

$$\Pi_0 = \text{diag}\left(4^2, 4^2, 4^2\right) \tag{60c}$$

$$G_k = \text{diag}\left(1, 1, 1\right) \tag{60d}$$

$$Q_k = \text{diag}\left(0.001^2, 0.001^2, 0.001^2\right) \tag{60e}$$

$$R_k = 0.25^2 \tag{60f}$$

$$\bar{x}_0 = \begin{bmatrix} 0 & 0 & 3.5 \end{bmatrix}^T \tag{60g}$$

Again, the initial guess, $\bar{x}_0$, is poor. The actual plant experiences $\mathcal{N}(0, Q_k)$ noise in the state and $\mathcal{N}(0, R_k)$ noise in the measurements. We now examine the estimation performance of both the EKF and MHE for this system.

Figure 13 demonstrates that, similarly to the batch case, the EKF appears to converge to an incorrect steady-state estimate. This observation is confirmed by determining the EKF steady state assuming no state or measurement noise. Calculating the EKF steady state via equations (35) and assuming no state or measurement noise yields the results in Table 2. Some steady-state analysis of the system sheds light on the cause of this phenomenon. Assuming no noise in the steady-state measurement, the system has one

linear steady-state measurement y_{ss}

$$c_A + c_B + c_C = \frac{y_{ss}}{RT} \tag{61}$$

and one linear combination resulting from ρ, the null space of the stoichiometric matrix

$$\rho = \begin{bmatrix} 3 & 1 & 2 \end{bmatrix} \tag{62}$$

$$3c_A + c_B + 2c_C = 3c_{Af} + c_{Bf} + 2c_{Cf} \tag{63}$$

Therefore the steady-state calculation is a nonlinear problem, and this system satisfies both conditions required for EKF failure.

Figure 14 presents the EKF estimation results for implementation of a clipping strategy. Although clipping eliminates estimation error, this strategy causes a lengthy period of overestimation of the pressure, in some cases by two orders of magnitude.

Figure 15 presents the results of applying MHE. For these results, we have constrained the state to prevent estimation of negative concentrations. These figures demonstrate that MHE swiftly converges to the correct state estimates.

Component	Predicted EKF Steady State	Actual Steady State
A	−0.0122	0.0224
B	−0.1364	0.2006
C	1.1746	0.6411

Table 2: EKF steady-state behavior, no measurement or state noise

7.2 Example 3

Reconsider the batch model given in section 7.1, but with the following updated parameters

$$k = \begin{bmatrix} 0.5 & 0.4 & 0.2 & 0.1 \end{bmatrix}^T \tag{64a}$$

$$R_k = 0.1^2 \tag{64b}$$

and new measurement

$$y_k = \begin{bmatrix} -1 & 1 & 1 \end{bmatrix} x_k \tag{65}$$

Note that the measurement has no physical meaning. Solving for the steady-state solution in terms of c_B yields

$$0 = -\frac{K_2}{K_1} c_B^3 + K_2 c_B^2 + c_B - y_{ss} \tag{66}$$

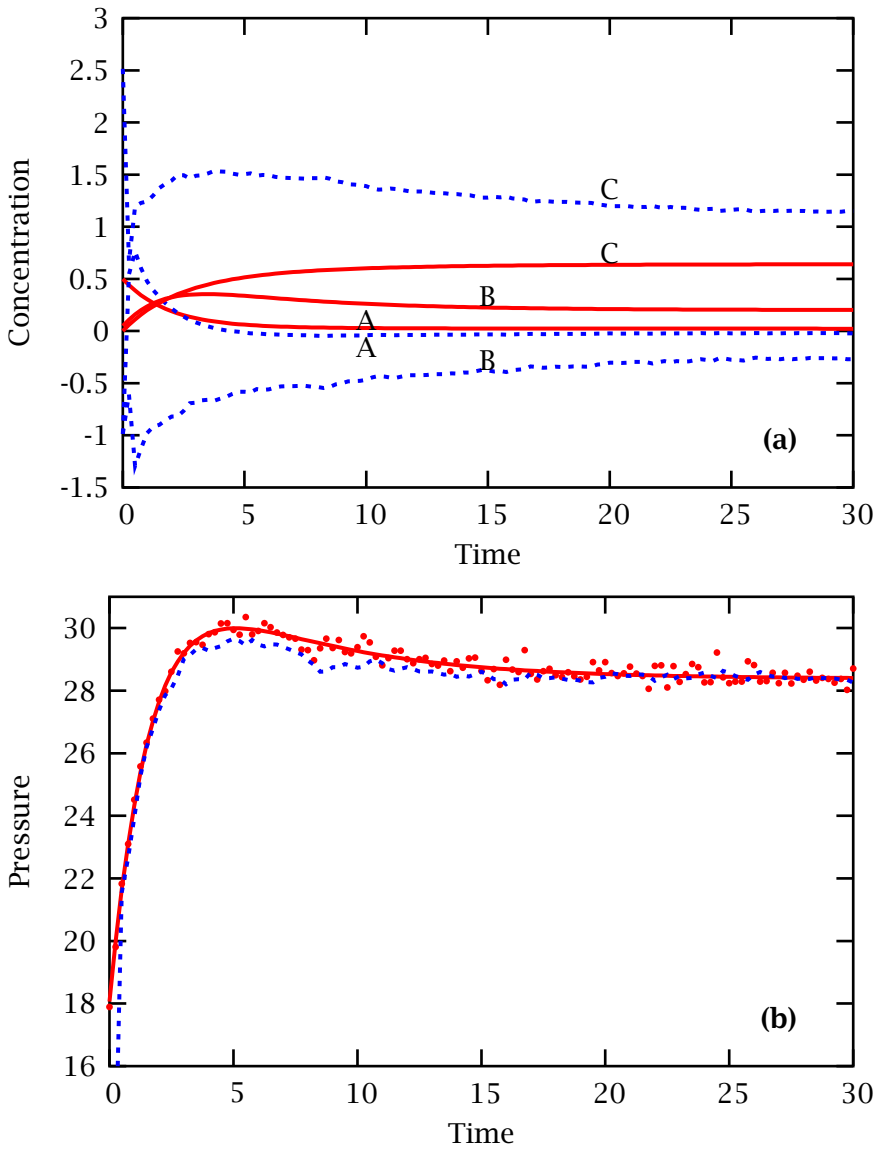

Figure 13: Extended Kalman filter results. (a) plots the evolution of the actual (solid line) and EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and EKF updated (dashed line) pressure estimates.

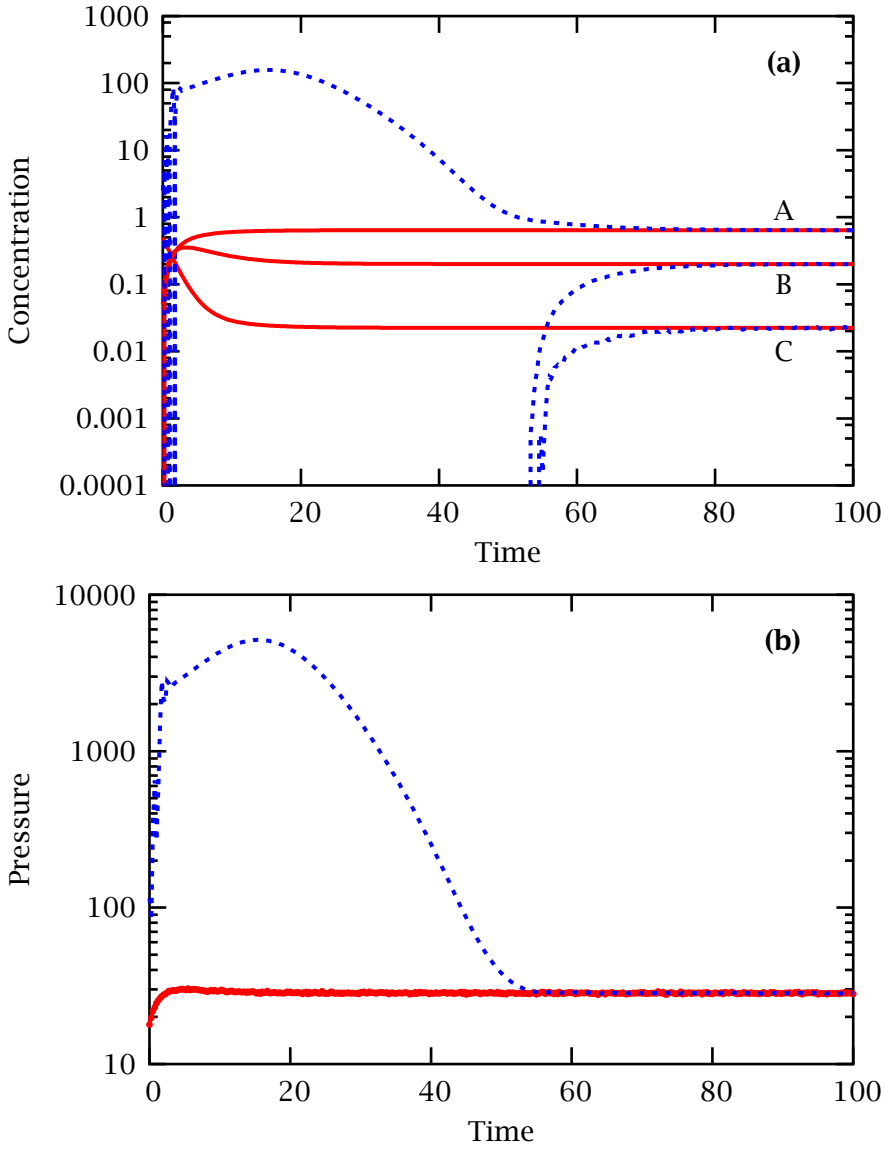

Figure 14: Clipped extended Kalman filter results. (a) plots the evolution of the actual (solid line) and clipped EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and clipped EKF updated (dashed line) pressure estimates.

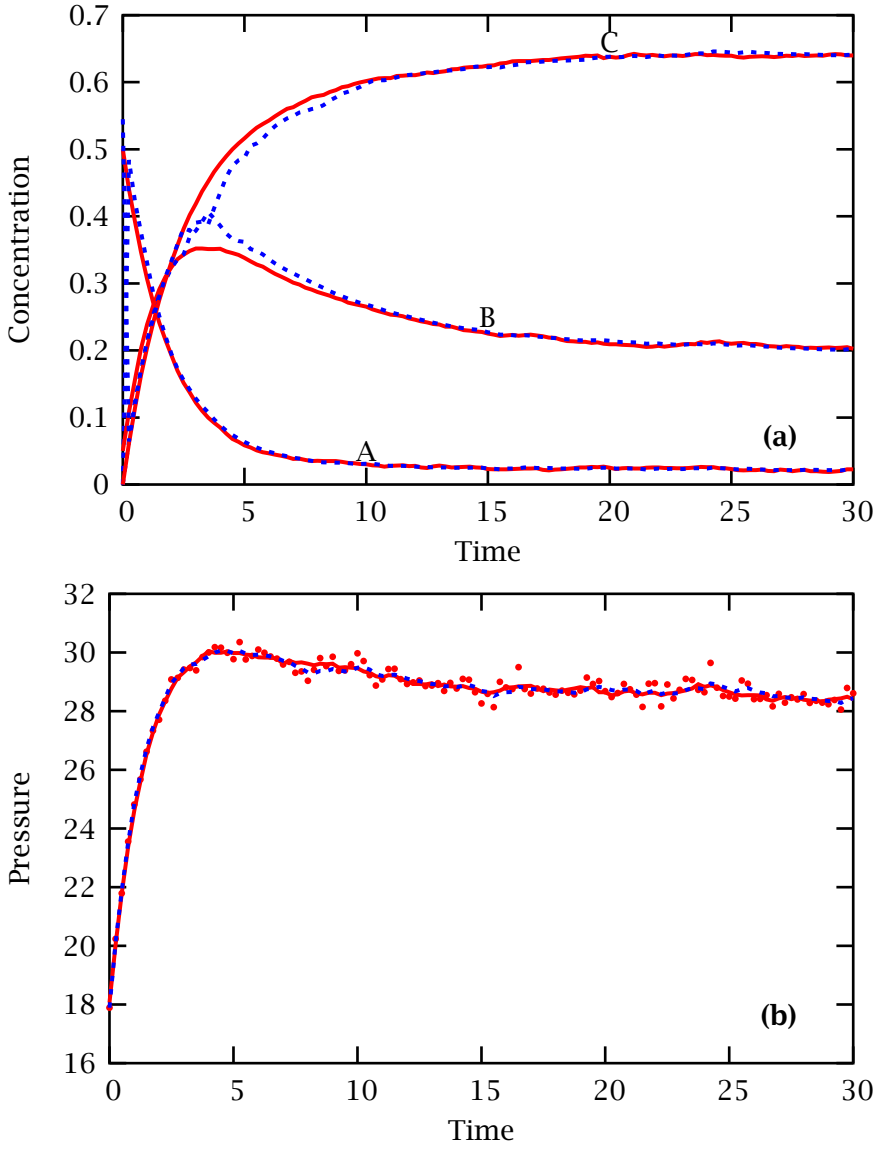

Figure 15: Moving horizon estimation results, states constrained to $x \geq 0$, smoothing initial covariance update, and horizon length of 2.5 time units ($N = 11$ measurements). (a) plots the evolution of the actual (solid line) and MHE updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) pressure estimates.

Again using Descartes' rule of signs and taking into account the specified parameters, equation (66) has two positive roots and one negative root. In contrast to the previous example, there are multiple physically realizable steady states. We now examine the effect of poor initial conditions upon the estimation behavior of the EKF and MHE.

Table 3 presents the *a priori* initial conditions for state estimation. Comparison of Figures 16 and 17 demonstrates that given a poor estimate of the initial state, the EKF cannot reconstruct the evolution of the state while MHE can. Figures 18 and 19 show that given an even poorer estimate of the initial state, both the EKF and MHE fail to reconstruct the evolution of the state. To improve the quality of the estimates, we constrain the concentrations in the estimators so that

$$0 \leq c_j \leq 4.5, \quad j = \text{A, B, C} \tag{67}$$

Figures 20 and 21 demonstrate that with this extra knowledge, MHE converges to the true state estimates while the clipped EKF estimates are trapped on the constraint. Finally, we relax the concentration constraints to

$$0 \leq c_j \leq 5.5, \quad j = \text{A, B, C} \tag{68}$$

Not surprisingly, the clipped EKF estimates remain trapped on the constraint, as shown in Figure 22. The quality of the MHE estimates is a function of the estimation horizon, as seen in Figure 23. If the estimation horizon is too short, the MHE estimates are pinned against the state constraint; increasing the horizon remedies this problem. For short horizons, we suspect that the data in the estimation horizon cannot overcome the biasing of the arrival cost approximation (with the smoothing scheme), hence resulting in state estimates pinned against the constraint. Changing arrival cost approximations (e.g. switching from the smoothing scheme to a uniform prior) when constraints are active may constitute one way of addressing this problem without having to increase the estimation horizon.

Table 4 summarizes the estimation results examined in this section.

Figures	$\bar{x}_0$
16, 17	$\begin{bmatrix} 3 & 0.1 & 3 \end{bmatrix}^T$
18-23	$\begin{bmatrix} 4 & 0 & 4 \end{bmatrix}^T$

Table 3: *A priori* initial conditions for state estimation

8 Conclusions

Virtually all chemical engineering systems contain nonlinear dynamics and/or state constraints. The need to incorporate this information into state estimation is illustrated by the examples presented in this paper. These examples demonstrate that even with

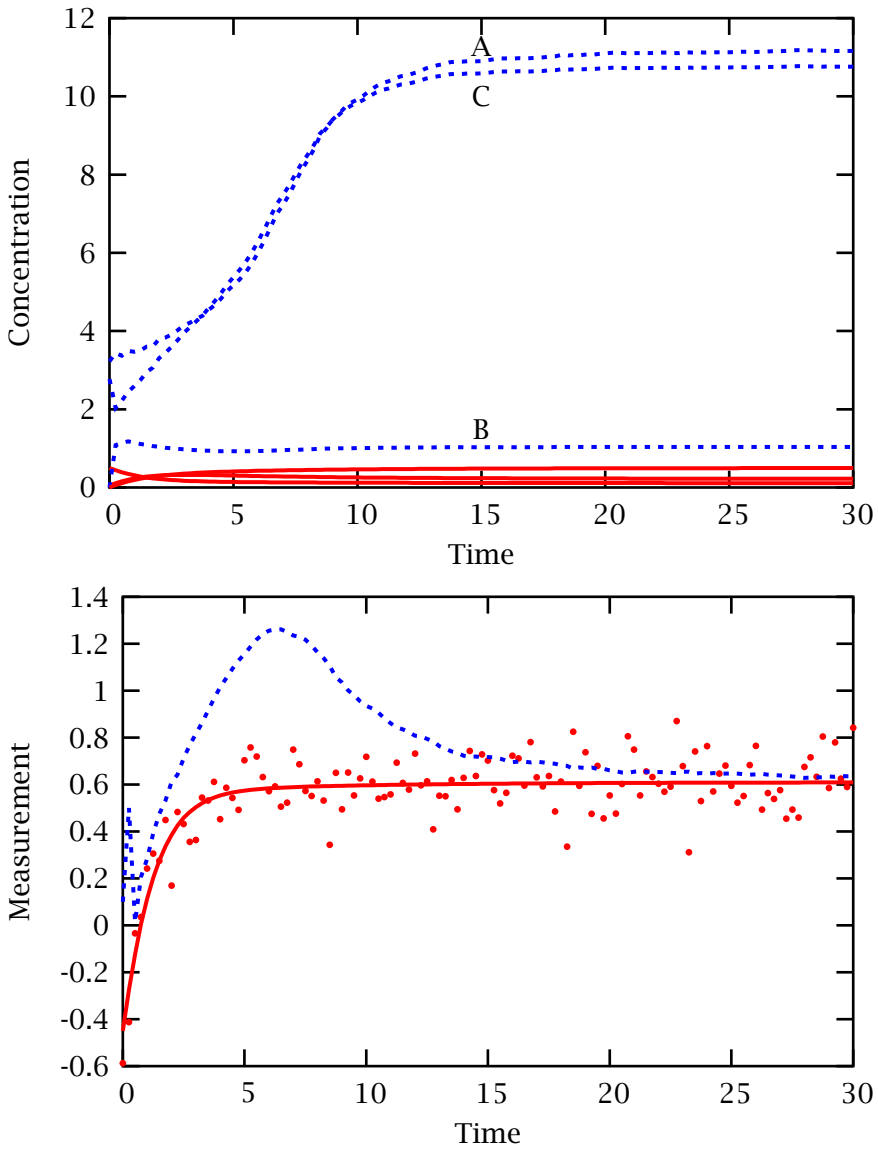

Figure 16: Extended Kalman filter results. (a) plots the evolution of the actual (solid line) and EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and EKF updated (dashed line) pressure estimates.

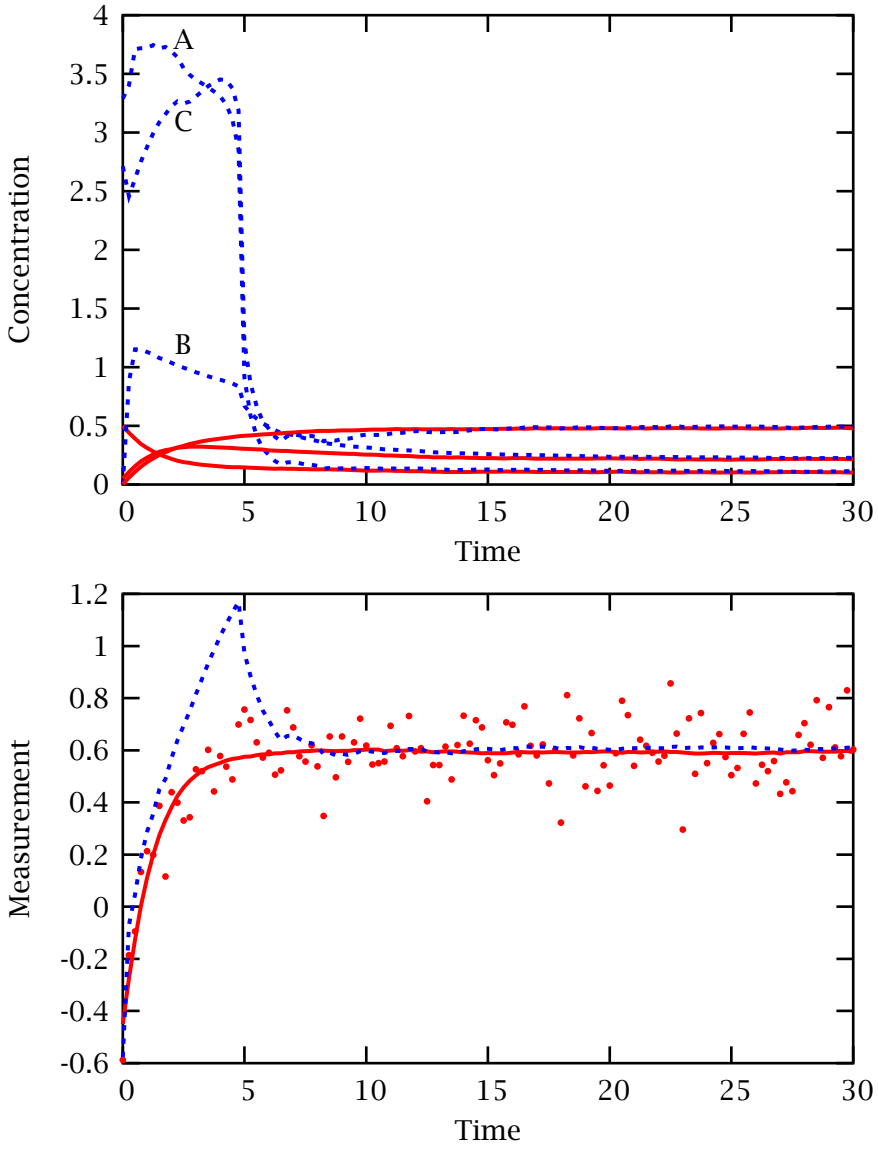

Figure 17: Moving horizon estimation results, states constrained to $x \geq 0$, smoothing initial covariance update, and horizon length of 2.5 time units ($N = 11$ measurements). (a) plots the evolution of the actual (solid line) and MHE updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) pressure estimates.

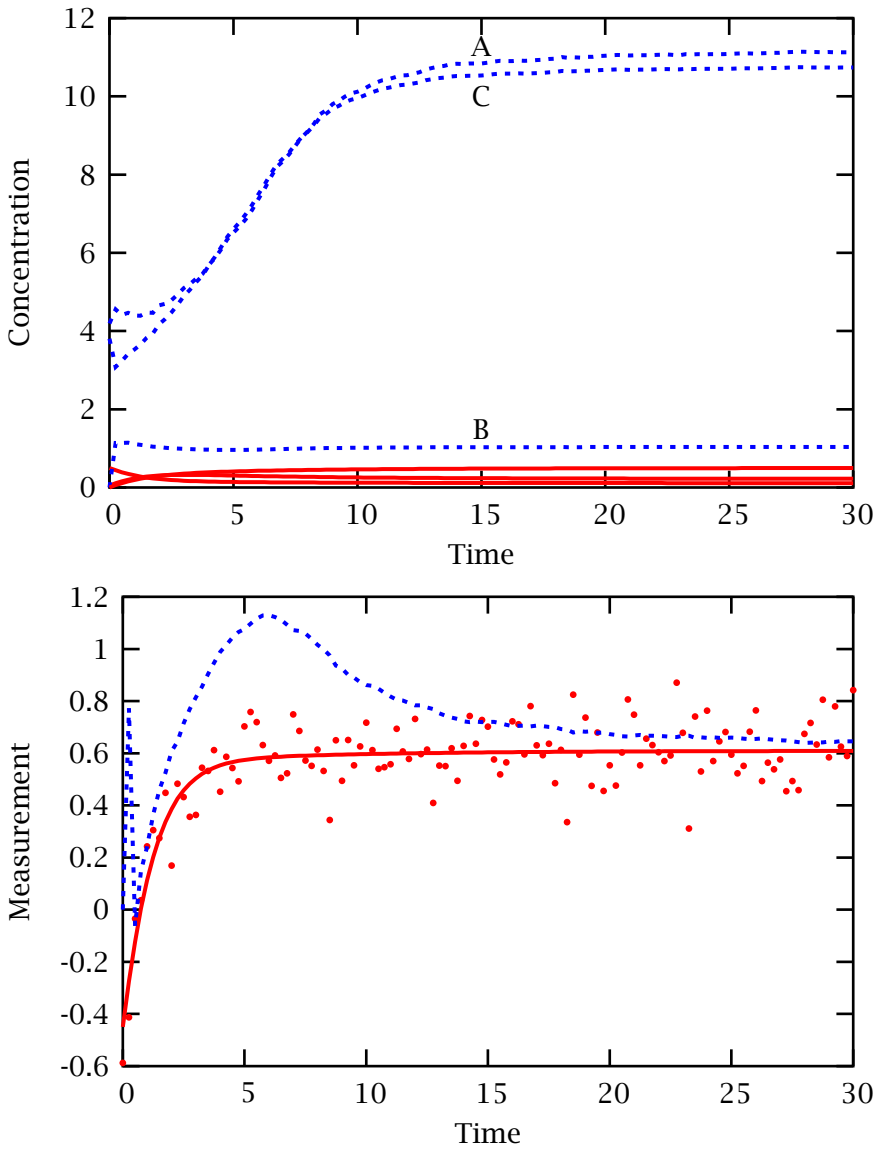

Figure 18: Extended Kalman filter results. (a) plots the evolution of the actual (solid line) and EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and EKF updated (dashed line) pressure estimates.

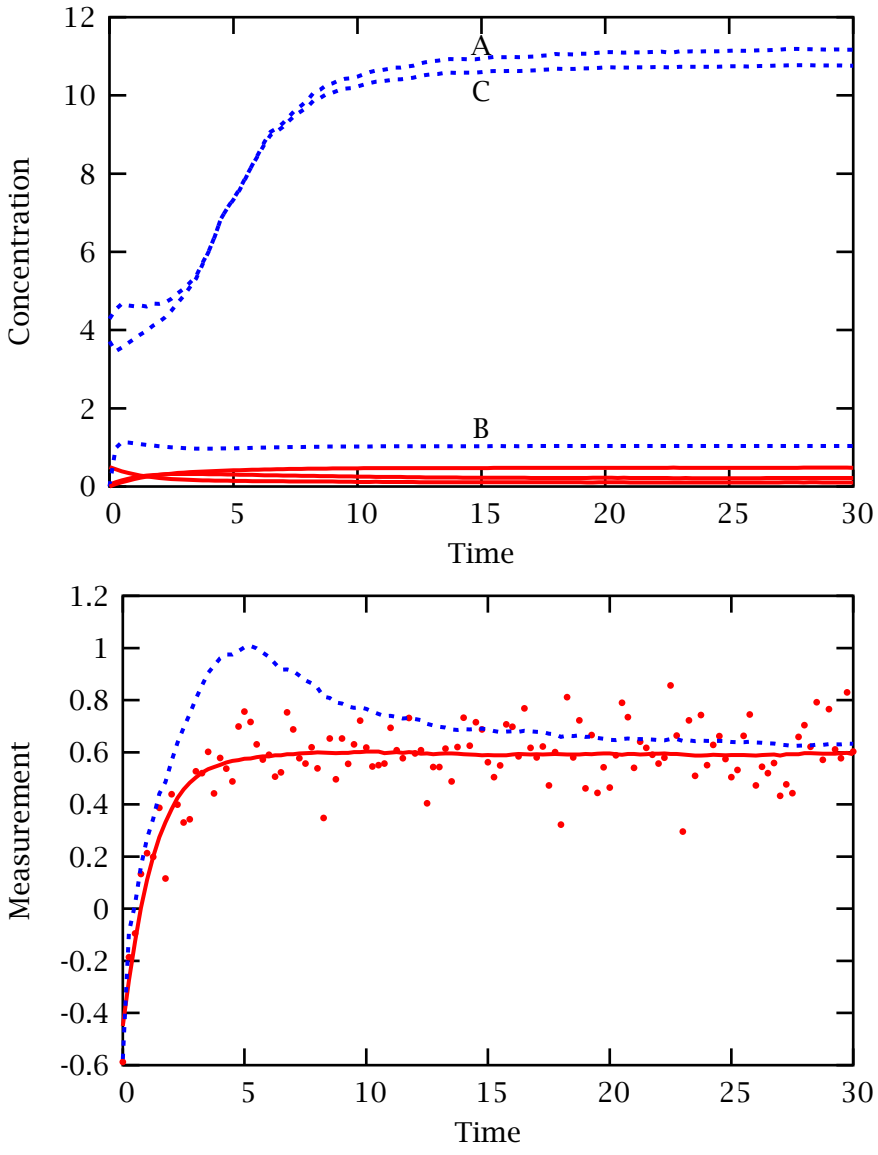

Figure 19: Moving horizon estimation results, states constrained to $x \geq 0$, smoothing initial covariance update, and horizon length of 2.5 time units ($N = 11$ measurements). (a) plots the evolution of the actual (solid line) and MHE updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) estimates.

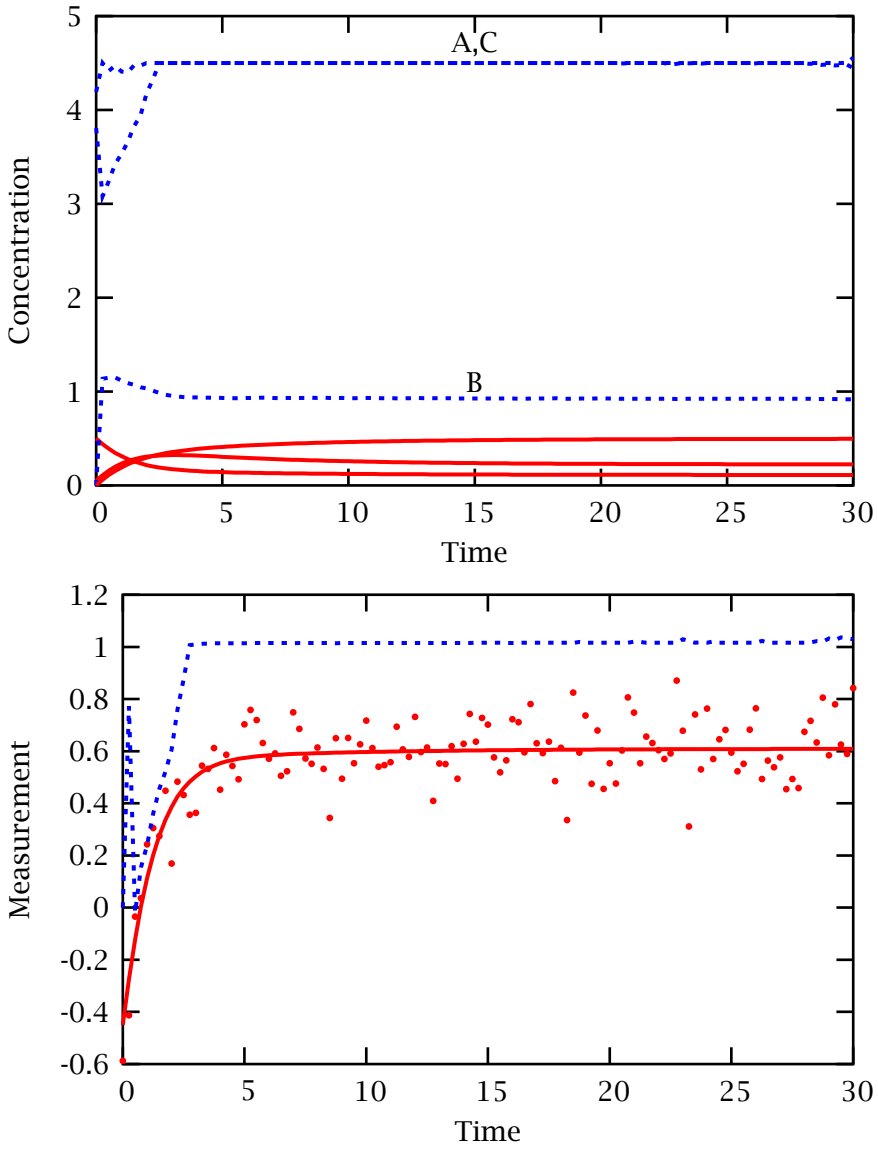

Figure 20: Clipped extended Kalman filter results, states clipped to $0 \leq x \leq 4.5$. (a) plots the evolution of the actual (solid line) and clipped EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and clipped EKF updated (dashed line) estimates.

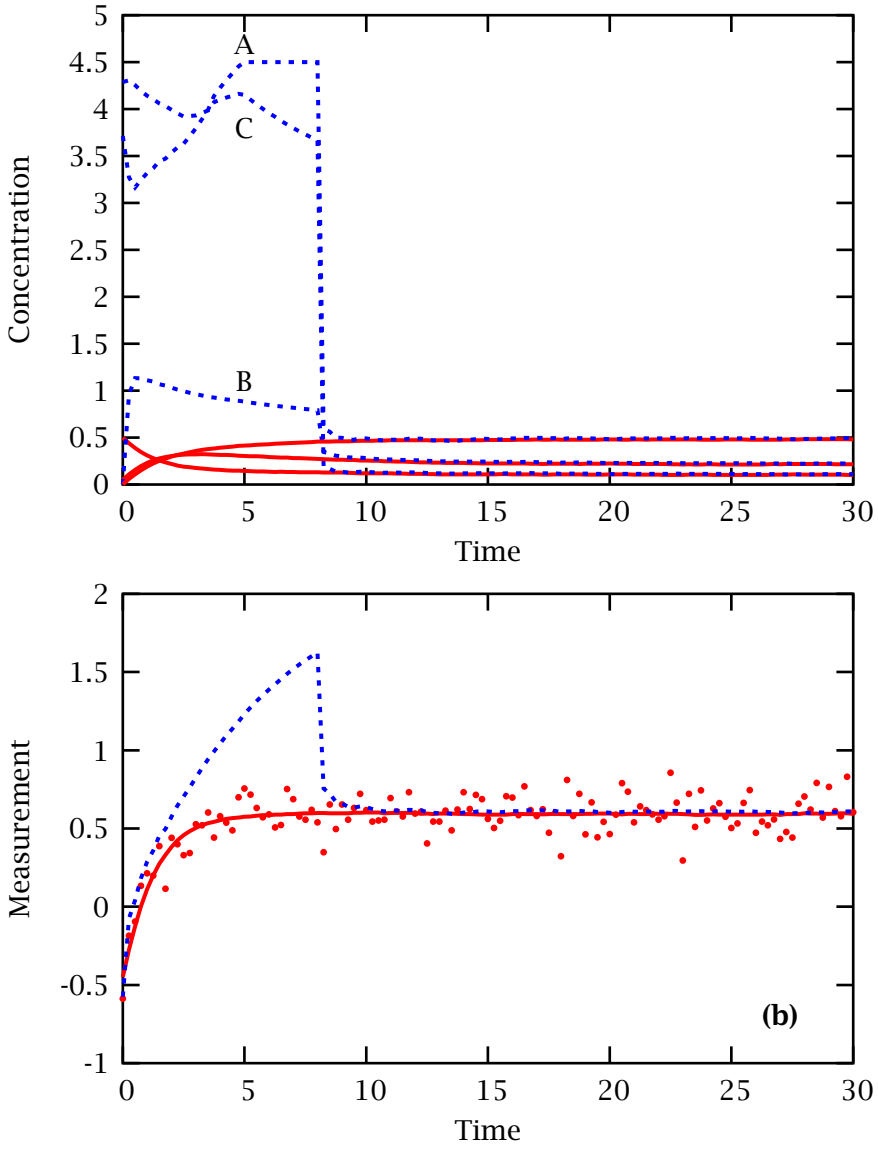

Figure 21: Moving horizon estimation results, states constrained to $0 \leq x \leq 4.5$, smoothing initial covariance update, and horizon length of 2.5 time units ($N = 11$ measurements). (a) plots the evolution of the actual (solid line) and MHE updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) estimates.

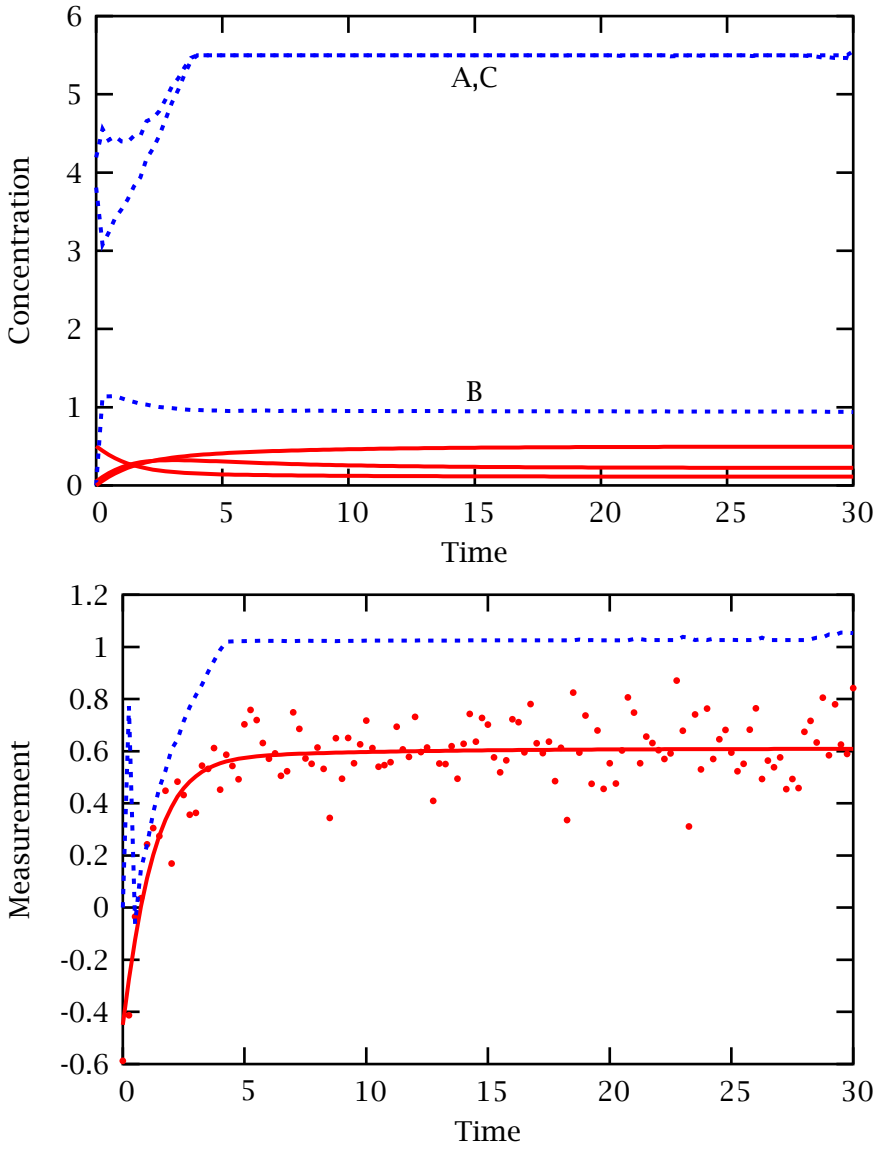

Figure 22: Clipped extended Kalman filter results, states clipped to $0 \leq x \leq 5.5$. (a) plots the evolution of the actual (solid line) and clipped EKF updated (dashed line) concentrations. (b) plots the evolution of the actual (solid line), measured (points), and clipped EKF updated (dashed line) estimates.

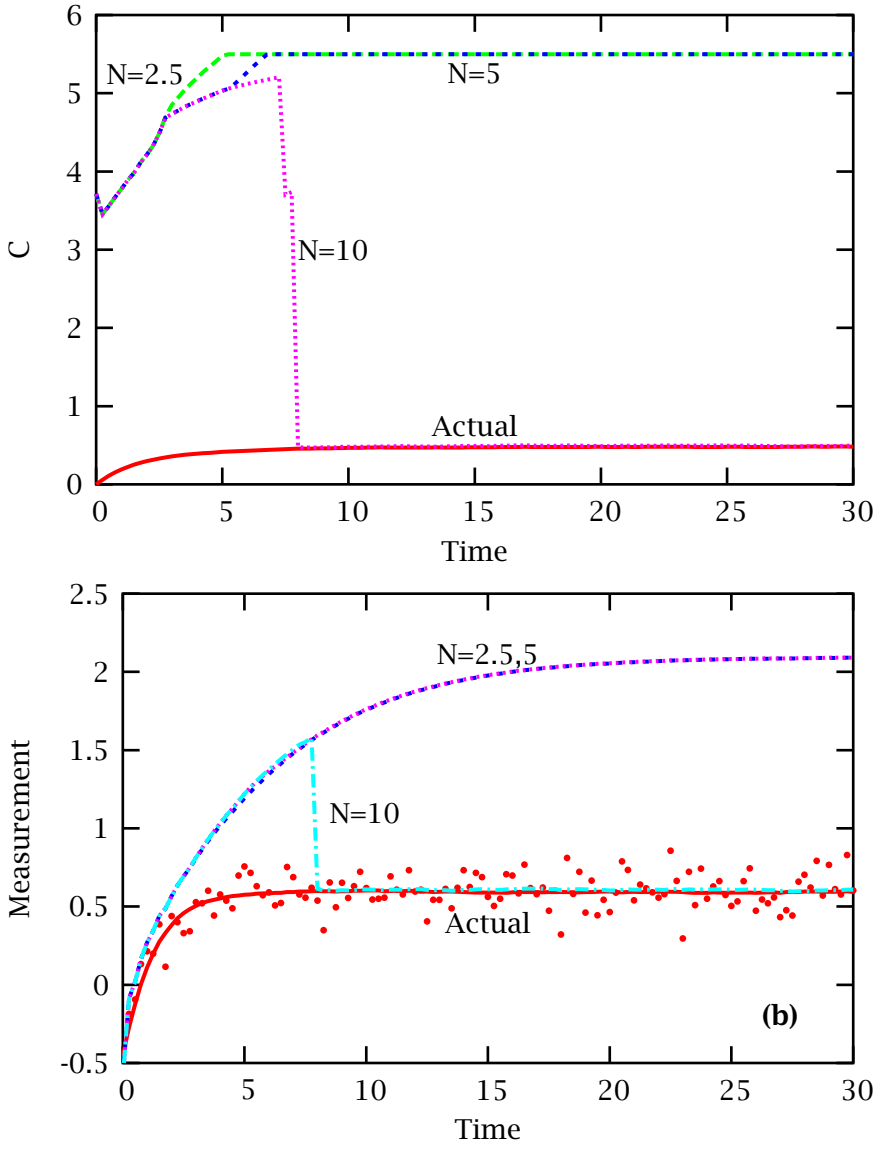

Figure 23: Moving horizon estimation results, states constrained to $0 \leq x \leq 5.5$, and smoothing initial covariance update. (a) plots the effect of horizon length on the evolution of the actual (solid line) and MHE updated (dashed line) C concentration. (b) plots the evolution of the actual (solid line), measured (points), and MHE updated (dashed line) estimates. Values of N on the plots correspond to the horizon length in time units.

Estimator	$\bar{x}_0$	Constraints	Horizon Length [5]	Estimates Converge?
EKF	$\begin{bmatrix}3 & 0.1 & 3\end{bmatrix}^T$	$x \geq 0$	NA	No
MHE	$\begin{bmatrix}3 & 0.1 & 3\end{bmatrix}^T$	$x \geq 0$	2.5 time units ($N = 11$)	Yes
EKF	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$x \geq 0$	NA	No
MHE	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$x \geq 0$	2.5 time units ($N = 11$)	No
EKF	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$0 \leq x \leq 4.5$	NA	No
MHE	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$0 \leq x \leq 4.5$	2.5 time units ($N = 11$)	Yes
EKF	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$0 \leq x \leq 5.5$	NA	No
MHE	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$0 \leq x \leq 5.5$	2.5 time units ($N = 11$)	No
MHE	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$0 \leq x \leq 5.5$	5 time units ($N = 21$)	No
MHE	$\begin{bmatrix}4 & 0 & 4\end{bmatrix}^T$	$0 \leq x \leq 5.5$	10 time units ($N = 41$)	Yes

Table 4: Effects of *a priori* initial conditions, constraints, and horizon length on state estimation. N denotes the number of measurements in the estimation horizon.

perfect concordance between the model and the physical plant, it is possible for the nominal EKF to fail to converge to the true state when

1. the system model and measurement are such that multiple states satisfy the steady-state measurement, and

2. the estimator is given a poor initial guess of the state.

Given the same estimator tuning, model, and measurements as the EKF, MHE provides improved state estimation and greater robustness to poor guesses of the initial state. These benefits arise because MHE incorporates physical state constraints into an optimization, accurately uses the nonlinear model, and optimizes over a trajectory of states and measurements.

The issue of global versus local optimization and the selection of an arrival cost also have substantial impact on the behavior of MHE. If one could implement a global optimization strategy in real time, approximating the arrival cost with a uniform prior and making the estimation horizon reasonably long is preferable to an approximate multivariate normal arrival cost because of the latter's biasing effect on the state estimates. With local optimization, our results indicate that multivariate normal approximations to the arrival cost combined with judicious use of constraints can prevent multiple optima in the estimator and generate acceptable estimator performance.

It is reasonable to expect that more complicated models than the ones proposed here may yield multiple optima corresponding to both physically realizable and unrealizable states. Since MHE permits incorporation of constraints into its optimization, it is the

natural choice for preventing estimation of physically unrealizable states. Since MHE employs a trajectory of measurements as opposed to measurements at only a single time, it is better suited than the EKF for distinguishing among the remaining physically realizable states.

9 Acknowledgments

The financial support of the National Science Foundation through grant #CTS-0105360 and the industrial members of the Texas-Wisconsin Modeling and Control Consortium is gratefully acknowledged. Helpful discussions with Gabriele Pannocchia and Matt Tenny are appreciated as well. All simulations were performed using Octave (`http://www.octave.org/`) and the NMPC toolbox (`http://www.che.wisc.edu/~tenny/nmpc/`). Both Octave and the NMPC toolbox are freely distributed under the terms of the GNU General Public License.

10 Appendix

10.1 Derivation of the MHE Smoothing Formulation

Assume that the system is Markov, that is,

$$p(x_{k+1}|x_0,\ldots,x_k) = p(x_{k+1}|x_k)$$

$$p(x_{T-N+1},\ldots,x_T|y_0,\ldots,y_T) \tag{69}$$

$$=p(x_{T-N+1}|y_0,\ldots,y_T)\boxed{p(x_{T-N+2},\ldots,x_T|y_0,\ldots,y_T,x_{T-N+1})} \tag{70}$$

$$=p(x_{T-N+1}|y_0,\ldots,y_T)\frac{\boxed{p(x_{T-N+2},\ldots,x_T,y_0,\ldots,y_T|x_{T-N+1})}}{p(y_0,\ldots,y_T|x_{T-N+1})} \tag{71}$$

$$=\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{p(y_0,\ldots,y_T|x_{T-N+1})}p(x_{T-N+2},\ldots,x_T|x_{T-N+1})\boxed{p(y_0,\ldots,y_T|x_{T-N+1},\ldots,x_T)} \tag{72}$$

$$=\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{p(y_0,\ldots,y_T|x_{T-N+1})}p(y_0,\ldots,y_{T-N}|x_{T-N+1},\ldots,x_T,y_{T-N+1},\ldots,y_T)\times$$
$$\boxed{p(x_{T-N+2},\ldots,x_T|x_{T-N+1})p(y_{T-N+1},\ldots,y_T|x_{T-N+1},\ldots,x_T)} \tag{73}$$

$$=\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{\boxed{p(y_0,\ldots,y_T|x_{T-N+1})}}p(y_0,\ldots,y_{T-N}|x_{T-N+1},\ldots,x_T,y_{T-N+1},\ldots,y_T)\times$$
$$\left(\prod_{k=T-N+1}^{T-1}p(x_{k+1}|x_k)\right)\left(\prod_{k=T-N+1}^{T}p(y_k|x_k)\right) \tag{74}$$

$$=\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})}\left(\prod_{k=T-N+1}^{T-1}p(x_{k+1}|x_k)\right)\left(\prod_{k=T-N+1}^{T}p(y_k|x_k)\right)\times$$
$$\boxed{\frac{p(y_0,\ldots,y_{T-N}|x_{T-N+1},\ldots,x_T,y_{T-N+1},\ldots,y_T)}{p(y_0,\ldots,y_{T-N}|x_{T-N+1},y_{T-N+1},\ldots,y_T)}} \tag{75}$$

$$=\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})}\left(\prod_{k=T-N+1}^{T-1}p(x_{k+1}|x_k)\right)\left(\prod_{k=T-N+1}^{T}p(y_k|x_k)\right)\times$$
$$\boxed{\frac{p(x_{T-N+2},\ldots,x_T|x_{T-N+1},y_0,\ldots,y_T)}{p(x_{T-N+2},\ldots,x_T|x_{T-N+1},y_{T-N+1},\ldots,y_T)}} \tag{76}$$

$$=\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})}\left(\prod_{k=T-N+1}^{T-1}p(x_{k+1}|x_k)\right)\left(\prod_{k=T-N+1}^{T}p(y_k|x_k)\right) \tag{77}$$

The corresponding probabilistic manipulations are:

From Equation	To Equation	Manipulation of Boxed Quantity
(69)	(70)	$p(a,b\|c) = p(a\|b,c)p(b\|c)$
(70)	(71)	$p(a,b\|c) = p(a\|b,c)p(b\|c)$
(71)	(72)	$p(a\|b,c) = \frac{p(a,b\|c)}{p(b\|c)}$
(72)	(73)	$p(a,b\|c) = p(a\|b,c)p(b\|c)$
(73)	(74)	$p(a,b\|c) = p(a\|b,c)p(b\|c)$ and the Markov property
(74)	(75)	$p(a,b\|c) = p(a\|b,c)p(b\|c)$
(75)	(76)	$\frac{p(a\|b,c,d)}{p(a\|b,d)} = \frac{p(c\|a,b,d)}{p(c\|b,d)}$
(76)	(77)	1 by the Markov property

10.2 Derivation of the MHE Filtering Formulation

From the smoothing formulation, we can recover the filtering formulation by manipulating the first term of (77):

$$\frac{p(x_{T-N+1}|y_0,\ldots,y_T)}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})} \tag{78}$$

$$= \frac{\boxed{p(x_{T-N+1},y_0,\ldots,y_T)}}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})p(y_0,\ldots,y_T)} \tag{79}$$

$$= \frac{\boxed{p(x_{T-N+1},y_{T-N+1},\ldots,y_T|y_0,\ldots,y_{T-N})}\,p(y_0,\ldots,y_{T-N})}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})p(y_0,\ldots,y_T)} \tag{80}$$

$$= \boxed{\frac{p(y_{T-N+1},\ldots,y_T|x_{T-N+1},y_0,\ldots,y_{T-N})}{p(y_{T-N+1},\ldots,y_T|x_{T-N+1})}}\,\frac{p(x_{T-N+1}|y_0,\ldots,y_{T-N})p(y_0,\ldots,y_{T-N})}{p(y_0,\ldots,y_T)} \tag{81}$$

$$= \boxed{\frac{p(y_0,\ldots,y_{T-N})}{p(y_0,\ldots,y_T)}}\,p(x_{T-N+1}|y_0,\ldots,y_{T-N}) \tag{82}$$

$$= \frac{p(x_{T-N+1}|y_0,\ldots,y_{T-N})}{p(y_{T-N+1},\ldots,y_T|y_0,\ldots,y_{T-N})} \tag{83}$$

$$\tag{84}$$

The corresponding probabilistic manipulations are:

From Equation	To Equation	Manipulation of Boxed Quantity
(78)	(79)	$p(a\|b) = \frac{p(a,b)}{p(b)}$
(79)	(80)	$p(a,b,c) = p(a,b\|c)p(c)$
(80)	(81)	$p(a,b\|c) = p(a\|b,c)p(b\|c)$
(81)	(82)	1 by the Markov property
(82)	(83)	$\frac{p(a,b)}{p(b)} = p(a\|b)$

The filtering formulation of MHE is thus

$$p(x_{T-N+1},\ldots,x_T|y_0,\ldots,y_T) = \frac{p(x_{T-N+1}|y_0,\ldots,y_{T-N})}{p(y_{T-N+1},\ldots,y_T|y_0,\ldots,y_{T-N})}\left(\prod_{k=T-N+1}^{T-1} p(x_{k+1}|x_k)\right)\left(\prod_{k=T-N+1}^{T} p(y_k|x_k)\right) \quad (85)$$

10.3 Equivalence of the Full Information and Least Squares Formulations

Starting with the maximum likelihood estimate, we rewrite equation (14) using equation (85):

$$\max_{x_0,\ldots,x_T} p(x_0,\ldots,x_T|y_0,\ldots,y_T) \quad (86)$$

$$= \min_{x_0,\ldots,x_T} -\log p(x_0,\ldots,x_T|y_0,\ldots,y_T) \quad (87)$$

$$= \min_{x_0,\ldots,x_T} -\log\left\{p(x_0)\left(\prod_{k=T-N+1}^{T-1} p(x_{k+1}|x_k)\right)\left(\prod_{k=T-N+1}^{T} p(y_k|x_k)\right)\right\} \quad (88)$$

$$= \min_{x_0,\ldots,x_T} -\log p(x_0) - \sum_{k=T-N+1}^{T-1} \log p(x_{k+1}|x_k) - \sum_{k=T-N+1}^{T} \log p(y_k|x_k) \quad (89)$$

We can then calculate the conditional probabilities in equation (89) by first rewriting the joint distributions as functions of independent random variables

$$p(x_{k+1},x_k) = p(G_k w_k)p(x_k) \quad (90)$$

$$p(y_k,x_k) = p(v_k)p(x_k) \quad (91)$$

and then calculating the conditional probabilities

$$p(x_{k+1}|x_k) = p(G_k w_k), \quad p(Gw_k) \sim \mathcal{N}(0, G_k Q_k G_k^T) \quad (92)$$

$$p(y_k|x_k) = p(v_k), \qquad p(v_k) \sim \mathcal{N}(0, R_k) \quad (93)$$

$$(94)$$

$\mathcal{N}(m,P)$-distributed multivariate normals have probability functions of the form

$$p(x) = \frac{1}{(2\pi)^{n/2}|P|^{1/2}} \exp\left[-\frac{1}{2}(x-m)^T P^{-1}(x-m)\right] \tag{95}$$

where n is the number of elements of the variable x. Therefore

$$\min_{x_k, x_{k+1}} -\log p(x_{k+1}|x_k) = \min_{x_k, x_{k+1}} -\log p(G_k w_k) \tag{96}$$

$$= \min_{x_k, x_{k+1}} \frac{1}{2} w_k^T Q_k^{-1} w_k \tag{97}$$

and

$$\min_{x_k} -\log p(y_k|x_k) = \min_{x_k} -\log p(v_k) \tag{98}$$

$$= \min_{x_k} \frac{1}{2} v_k^T R_k^{-1} v_k \tag{99}$$

Plugging these values into equation (89) yields the minimization presented in equation (15).

10.4 Evolution of a Nonlinear Probability Density

We are interested in determining formulas for the evolution of the probability density $p(x_k|y_0,\ldots,y_k)$ for the system

$$x_{k+1} = \begin{bmatrix} \dfrac{x_{k,1}}{2k\Delta t x_{k,1}+1} \\ x_{k,2} + \dfrac{k\Delta t x_{k,1}^2}{2k\Delta t x_{k,1}+1} \end{bmatrix} + w_k \tag{100a}$$

$$y_k = \begin{bmatrix} 1 & 1 \end{bmatrix} x_k + v_k \tag{100b}$$

We view future states (x_k's) as functions of the random variables with known statistics (x_0, w_k's, and v_k's). First update the *a priori* estimate, $\bar{x}_0$, with the first measurement, y_0, by

1. writing the joint probability density $p(x_0,y_0)$ as a function of $p(x_0,v_0)$

$$p(x_0,y_0) = p(x_0,v_0)\begin{vmatrix} I & 0 \\ C & I \end{vmatrix}^{-1} \tag{101a}$$

$$= p(x_0)p(v_0) \qquad (x_0 \text{ and } v_0 \text{ are independent}) \tag{101b}$$

2. calculating the conditional probability density $p(x_0|y_0)$

$$p(x_0|y_0) = \frac{p(x_0, y_0)}{p(y_0)} \tag{102a}$$

$$= \frac{p(x_0, y_0)}{\int_{-\infty}^{\infty} p(x_0, y_0) dx_0} \tag{102b}$$

$$= \frac{p(x_0) p_{v_0}(y_0 - Cx_0)}{\int_{-\infty}^{\infty} p(x_0) p_{v_0}(y_0 - Cx_0) dx_0} \tag{102c}$$

$$= \frac{\exp\left[-\frac{1}{2}(x_0 - \bar{x}_0)^T \Pi_0^{-1}(x_0 - \bar{x}_0) - \frac{1}{2}(y_0 - Cx_0)^T R_k^{-1}(y_0 - Cx_0)\right]}{\int_{-\infty}^{\infty} \exp\left[-\frac{1}{2}(x_0 - \bar{x}_0)^T \Pi_0^{-1}(x_0 - \bar{x}_0) - \frac{1}{2}(y_0 - Cx_0)^T R_k^{-1}(y_0 - Cx_0)\right] dx_0} \tag{102d}$$

Now propagate $p(x_0|y_0)$ to the next measurement time to obtain $p(x_1|y_0)$:

$$p(x_1, w_0|y_0) = p(x_0, w_0|y_0) \left| \begin{matrix} \frac{\partial x_1}{\partial x_0^T} & \frac{\partial x_1}{\partial w_0^T} \\ \frac{\partial w_0}{\partial x_0^T} & \frac{\partial w_0}{\partial w_0^T} \end{matrix} \right|^{-1} \tag{103a}$$

$$= p(x_0, w_0|y_0) \left(2\bar{k}\Delta x_{0,1} + 1\right)^2 \tag{103b}$$

$$= p(x_0|y_0) p(w_0) \left(2\bar{k}\Delta x_{0,1} + 1\right)^2 \tag{103c}$$

$$p(x_1|y_0) = \int_{-\infty}^{\infty} p(x_1, w_0|y_0) dw_0 \tag{104}$$

$$= \int_{-\infty}^{\infty} p(x_0|y_0) p(w_0) \left(2\bar{k}\Delta x_{0,1} + 1\right)^2 dw_0 \tag{105}$$

$$p(x_1|y_0, y_1) = \frac{p(x_1|y_0) p_{v_1}(y_1 - Cx_1)}{\int_{-\infty}^{\infty} p(x_1|y_0) p_{v_1}(y_1 - Cx_1) dx_1} \tag{106a}$$

$$= \frac{\int_{-\infty}^{\infty} \Omega_1 dw_0}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Omega_1 dw_0 dx_1} \tag{106b}$$

where

$$\Omega_1 = \int_{-\infty}^{\infty} \left(2\bar{k}\Delta x_{0,1} + 1\right)^2 \exp\left[-\frac{1}{2}\left((x_0 - \bar{x}_0)^T \Pi_0^{-1}(x_0 - \bar{x}_0) + w_0^T Q_k^{-1} w_0 + \sum_{j=0}^{1} v_j^T R_k^{-1} v_j\right)\right] dw_0 \tag{107}$$

For future times, it is straightforward to derive equations (28) and (29).

References

[1] George E. P. Box and George C. Tiao. *Bayesian Inference in Statistical Analysis*. Addison–Wesley Publishing Company, Reading, Massachusetts, 1st edition, 1973.

[2] Madalena Chaves and Eduardo Sontag. State-estimators for chemical reaction networks of Feinberg-Horn-Jackson zero deficiency type. *Euro. J. Cont.*, 8(2002):343–359, 2002.

[3] Chi-Tsong Chen. *Linear System Theory and Design*. Oxford University Press, 3rd edition, 1999.

[4] Alan Genz and Robert E. Kass. A collection of numerical integration software for Bayesian analysis. Available from http://www.sci.wsu.edu/math/faculty/genz/homepage, 1998.

[5] R. Gudi, S. Shah, and M. Gray. Multirate state and parameter estimation in an antibiotic fermentation with delayed measurements. *Biotechnology and Bioengineering*, 44:1271–1278, 1994.

[6] Scott A. Middlebrooks. *Modelling and Control of Silicon and Germanium Thin Film Chemical Vapor Deposition*. PhD thesis, University of Wisconsin–Madison, 2001.

[7] Vinay Prasad, Matthias Schley, Louis P. Russo, , and B. Wayne Bequette. Product property and production rate control of styrene polymerization. *J. Proc. Cont.*, 12(3):353–372, 2002.

[8] Christopher V. Rao. *Moving Horizon Strategies for the Constrained Monitoring and Control of Nonlinear Discrete-Time Systems*. PhD thesis, University of Wisconsin–Madison, 2000.

[9] Christopher V. Rao and James B. Rawlings. Constrained process monitoring: moving-horizon approach. *AIChE J.*, 48(1):97–109, January 2002.

[10] Christopher V. Rao, James B. Rawlings, and Jay H. Lee. Constrained linear state estimation – a moving horizon approach. *Automatica*, 37(10):1619–1628, 2001.

[11] Christopher V. Rao, James B. Rawlings, and David Q. Mayne. Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations. *IEEE Trans. Auto. Cont.*, 48(2):246–258, February 2003.

[12] Konrad Reif, Stefan Günther, Engin Yaz, and Rolf Unbehauen. Stochastic stability of the discrete-time extended Kalman filter. *IEEE Trans. Auto. Cont.*, 44(4):714–728, April 1999.

[13] Konrad Reif, Stefan Günther, Engin Yaz, and Rolf Unbehauen. Stochastic stability of the continuous-time extended Kalman filter. *IEE Proceedings-Control Theory and Applications*, 147(1):45–52, January 2000.

[14] Konrad Reif and Rolf Unbehauen. The extended Kalman filter as an exponential observer for nonlinear systems. *IEEE Transactions on Signal Processing*, 47(8):2324–2328, August 1999.

[15] Wen shiang Chen, Sridhar Ungarala, Bhavik Bakshi, and Prem Goel. Bayesian rectification of nonlinear dynamic processes by the weighted bootstrap. In *AIChE Annual Meeting, Reno, Nevada*, 2001.

[16] Masoud Soroush. State and parameter estimations and their applications in process control. *Comput. Chem. Eng.*, 23(2):229–245, December 1998.

[17] Robert F. Stengel. *Optimal Control and Estimation*. Dover Publications, Inc., 1994.

[18] Matthew J. Tenny and James B. Rawlings. State estimation strategies for nonlinear model predictive control. AIChE Annual Meeting, Reno, November 2001.

[19] Matthew J. Tenny and James B. Rawlings. Efficient moving horizon estimation and nonlinear model predictive control. In *Proceedings of the American Control Conference*, pages 4475–4480, Anchorage, Alaska, May 2002.

[20] Y. Wang and E. D. Sontag. Output-to-state stability and detectability of nonlinear systems. *Sys. Cont. Let.*, 29:279–290, 1997.

[21] D. I. Wilson, M. Agarwal, and D.W.T. Rippin. Experiences implementing the extended Kalman filter on an industrial batch reactor. *Comput. Chem. Eng.*, 22(11):1653–1672, 1998.