Nonlinear Model Predictive Control for Stochastic Differential Equation Systems

Brok, Niclas Laursen; Madsen, Henrik; Jørgensen, John Bagterp

Published in:
I F A C Workshop Series

Link to article, DOI:
10.1016/j.ifacol.2018.11.071

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 51-20 (2018) 430–435

Nonlinear Model Predictive Control for Stochastic Differential Equation Systems

Niclas Laursen Brok, Henrik Madsen, John Bagterp Jørgensen

Department of Applied Mathematics and Computer Science, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

Abstract: Using the Van der Pol oscillator model as an example, we provide a tutorial introduction to nonlinear model predictive control (NMPC) for systems governed by stochastic differential equations (SDEs) that are observed at discrete times. Such systems are called continuous-discrete systems and provides a natural representation of systems evolving in continuous-time. Furthermore, this representation directly facilities construction of the state estimator in the NMPC. We provide numerical details related to systematic model identification, state estimation, and optimization of dynamical systems that are relevant to the NMPC.

Keywords: Nonlinear Model Predictive Control, Continuous-Discrete Extended Kalman Filter, Maximum Likelihood Estimation, Stochastic Differential Equations, Van der Pol Oscillator

1. INTRODUCTION

This paper provides a tutorial of how to use the nonlinear model predictive control (NMPC) principle to regulate a stochastic system governed by stochastic differential equations (SDEs). The systems considered in this paper are continuous-discrete systems of the form (Jazwinski, 1970)

$$\mathrm{d}\boldsymbol{x}(t) = f(\boldsymbol{x}(t), u(t); p)\,\mathrm{d}t + g(\boldsymbol{x}(t), u(t); p)\,\mathrm{d}\boldsymbol{\omega}(t), \quad (1a)$$

$$\boldsymbol{y}(t_k) = h(\boldsymbol{x}(t_k)) + \boldsymbol{v}_k, \quad (1b)$$

where $\boldsymbol{x}$, u and p are the states, inputs and time-invariant parameters. $\boldsymbol{v}_k \sim N_{\mathrm{iid}}(0, R_k)$ is the measurement noise and $\boldsymbol{\omega}(t)$ is a standard Brownian motion. Brownian motion is defined by its independent increments which satisfies that for each $s, t \in \mathbb{R}$, $\boldsymbol{\omega}(t) - \boldsymbol{\omega}(s)$ is normally distributed with zero mean and covariance $I(t-s)$; i.e. $\mathrm{d}\boldsymbol{\omega}(t) \sim N_{\mathrm{iid}}(0, I\,\mathrm{d}t)$. The SDE model representation (1a) provides a natural way to represent physical systems as they evolve in continuous-time. In contrast to discrete-time models, a priori knowledge about the system can be included and the estimated parameters do not depend on the sampling time. The representation of noise in continuous time also allow for a parsimonious representation that is independent of the sampling time. While these advantages of the continuous-discrete model (1) are well-known in the systems identification community (Garnier and Young, 2012; Kristensen et al., 2004; Rao and Unbehauen, 2006), most NMPC methods rely on either 1) a deterministic discrete-time model, 2) a stochastic discrete-time model, or 3) a deterministic continuous-time model for which the noise terms in the estimators are added in an ad hoc manner. The key insight is that the continuous-discrete model (1) provides a systematic way to obtain an estimation (filtering and prediction) algorithm that is used in the offline system identification, the online state- and parameter-estimation, and the prediction of the dynamic optimization. The diffusion model, $g(\boldsymbol{x}(t), u(t); p)$, represents a convenient and powerful way of representing complex stochastic processes and model-plant mismatch as needed for the filtering and prediction algorithm in NMPC. Boiroux et al. (2016a,b,c, 2010) and Mahmoudi et al. (2016a, 2017, 2016b) demonstrate systematic use of the continuous-discrete model (1) for system identification, nonlinear filtering and prediction, fault detection, and NMPC in an artificial pancreas for people with type 1 diabetes. Mahmood and Mhaskar (2012) uses a continuous-discrete model (1) and a Lyapunov-based NMPC to stabilize a reaction model around an unstable equilibrium. Buehler et al. (2016) uses a similar setup where a bioreactor model is controlled according to a desired set-point probability density. Jørgensen and Jørgensen (2007a,b) apply the continuous-discrete stochastic model (1) for linear MPC using transfer function representations of the model.

1.1 Components and software of the NMPC

Allgöwer et al. (1999), Johansen (2011), Grüne and Pannek (2011), and Rawlings et al. (2017) describe state-of-the-art NMPC technology. A system for NMPC consists of an offline method for identification of the model as well as an online part - where the online part consists of a state- and parameter estimation algorithm and an algorithm for dynamic optimization. Fig. 1 schematically illustrates this structure. This tutorial fills a missing gap in existing NMPC literature, by systematically formulating all components in the NMPC software system based on the continuous-discrete model (1). All components in the NMPC, presented in this paper, use the same continuous-discrete extended Kalman filter (CDEKF) for (1) in the one-step prediction of the offline system identification, in the online state- and parameter-estimation algorithm, and in the prediction of the dynamic optimization algorithm.

⋆ Innovation Fund Denmark is acknowledged for partly funding the work as a part of Center for IT-Intelligent Energy Systems (CITIES - IFD 1035-00027B).

10.1016/j.ifacol.2018.11.071

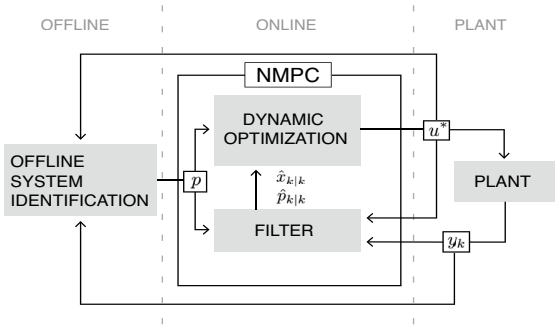

Fig. 1. Overview of the closed-loop control structure.

The rigorous solution to the filtering and prediction problem is obtained by solving the Fokker-Planck equations (Jazwinski, 1970). However, for systems with more than a couple of states, this solution is computationally intractable. The CDEKF is a computationally tractable alternative for filtering and prediction in (1). While other filters such as the unscented Kalman filter (UKF), the ensemble Kalman filter (EnKF), the particle filter (PF), and the moving horizon estimator (MHE) can also be used instead of the (CDEKF), the CDEKF represents the best balance between performance and computational tractability for many processes (Simon, 2006). This is particularly true, when the maximum-likelihood method is used for estimation of the parameters in the filter and predictor. For good performance of the CDEKF, it is also important to notice that we implement it using a differential equation solver with adaptive time step and using the Joseph stabilization scheme (Schneider and Georgakis, 2013). By including a disturbance model as part of the model, the CDEKF is used for online estimation of the states as well as selected rapidly varying parameters.

The offline system identification is based on a maximum likelihood (ML) formulation, where the conditional densities of the state equations are approximated by Gaussian densities. Using this assumption, it is possible to derive an optimization problem which uses the CDEKF to compute the likelihood of the parameters, p, given a set of observations (Kristensen et al., 2004).

The dynamic optimization component of the NMPC consists of the solution of a deterministic open-loop optimal control problem. The optimal control problem considered is a Bolza problem with input constraints, i.e.

$$\min_{x,u} \quad \int_{t_k}^{t_k+T} l(x(t),u(t))\,\mathrm{d}t \;\; + l_f(x(t_k+T)), \tag{2a}$$
$$s.t. \quad x(t_k) = \hat{x}_{k|k}, \tag{2b}$$
$$\dot{x}(t) = f(x(t),u(t);p),\; t \in [t_k, t_k+T]\,, \tag{2c}$$
$$u(t) \in \mathcal{U}(t), \qquad t \in [t_k, t_k+T]\,, \tag{2d}$$

where $\hat{x}_{k|k}$ denotes the filtered state estimates (from the filter) and p is the parameter estimates (from the online- or offline estimation method). $T = T_c = T_p$ is the control and prediction horizon. Several indirect and direct methods exists for the numerical solution of this optimal control problem (Binder et al., 2001). In this paper we use a direct local collocation method (von Stryk, 1993).

We use a stochastic extension of the van der Pol oscillator model to illustrate the components of the NMPC system. The numerical methods for the simulation study are implemented in `python` and the source code is available via `GitHub`[1].

1.2 Paper organization

The paper is organised as follows. Section 2 presents the CDEKF, while Section 3 derives the use of the CDEKF for online and offline parameter estimation. Section 4 presents the local collocation method for numerical optimal control. Section 5 illustrates these components of the NMPC using the stochastic van der Pol oscillator model. Finally, Section 6 contains a short summary.

2. THE EXTENDED KALMAN FILTER

We present the CDEKF used in the NMPC as well as for the offline system identification. $\hat{x}_{k|k}$ and $\hat{P}_{k|k}$ denote the filtered state- and covariance estimates. $\hat{x}_{k|k-1}$ and $\hat{P}_{k|k-1}$ denote the predicted (one-step predictions) state- and covariance values.

2.1 The prediction scheme

Given the initial conditions

$$\hat{x}_{k-1}(t_{k-1}) = \hat{x}_{k-1|k-1}, \quad \hat{P}_{k-1}(t_{k-1}) = \hat{P}_{k-1|k-1}, \tag{3}$$

the state- and covariance are predicted by solving the system of ordinary differential equations (ODEs) given by

$$\dot{\hat{x}}_{k-1}(t) = f(\hat{x}_{k-1}(t), u(t); p), \tag{4a}$$
$$\dot{\hat{P}}_{k-1}(t) = A(t)\hat{P}_{k-1}(t) + \hat{P}_{k-1}(t)A(t)' + G(t)G(t)', \tag{4b}$$

where

$$A(t) = \frac{\partial f}{\partial x}(\hat{x}_{k-1}(t), u(t); p), \quad G(t) = g(\hat{x}_{k-1}(t), u(t); p).$$

The one-step predictions of the mean and covariance of the states are obtained as the solution of (3)-(4) at the new sample point, t_k. Consequently, the predictions of the state- and covariance are

$$\hat{x}_{k|k-1} = \hat{x}_{k-1}(t_k), \quad \hat{P}_{k|k-1} = \hat{P}_{k-1}(t_k). \tag{5}$$

2.2 The updating scheme

The literature contains many methods for the updating scheme of extended Kalman filter algorithms. They all compute the innovation by

$$e_k = y_k - h(\hat{x}_{k|k-1}), \tag{6}$$

the Kalman filter gain, K_k, by

$$C_k = \frac{\partial h}{\partial x}(\hat{x}_{k|k-1}), \tag{7a}$$
$$R_{k|k-1} = C_k \hat{P}_{k|k-1} C_k' + R_k, \tag{7b}$$
$$K_k = \hat{P}_{k|k-1} C_k' R_{k|k-1}^{-1}, \tag{7c}$$

and the filtered state estimate, $\hat{x}_{k|k}$, by

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k e_k. \tag{8}$$

The key difference is how they compute the filtered covariance, $P_{k|k}$. Two standard updating schemes for the covariance are

$$\hat{P}_{k|k} = (I - K_k C_k)\,\hat{P}_{k|k-1} \tag{9a}$$
$$= \hat{P}_{k|k-1} - K_k R_{k|k-1} K_k'. \tag{9b}$$

[1] `https://github.com/niclasbrok/nmpc_vdp.git`

Numerical implementations based on either (9a) or (9b) may give rise to bad performance and even divergence, as the *numerically* computed values are not guaranteed to be both positive (semi-)definite and symmetric. The Joseph stabilization form

$$\hat{P}_{k|k} = (I - K_k C_k)\, \hat{P}_{k|k-1}\, (I - K_k C_k)' + K_k R_k K_k'. \quad (10)$$

for updating the filtered covariance estimate guarantees that the numerical value of $P_{k|k}$ is symmetric positive (semi-)definite. The CDEKF is implemented using (10) rather than (9) for reasons of numerical stability and robustness (Schneider and Georgakis, 2013). Numerically stable alternatives based on array- and squareroot-algorithms also exist (Boiroux et al., 2016c), but are less straightforward to implemenet compared to (10).

3. PARAMETER ESTIMATION

In this section, we outline the application of the CDEKF for online and offline parameter estimation.

3.1 Online identification using the CDEKF

Using a disturbance model, the CDEKF can be used for parameter estimation in addition to state estimation. One way to do this is by augmenting the SDE by as many states as parameters undergoing the estimation

$$\mathrm{d}\boldsymbol{p} = \Sigma\, \mathrm{d}\boldsymbol{\omega}_p, \quad (11)$$

where $\boldsymbol{p} = (\boldsymbol{p}_1, ..., \boldsymbol{p}_{N_p})^\top$ are the parameters to be estimated and $\Sigma = \mathrm{diag}(\sigma_1, ..., \sigma_{N_p})$. Success of this approach depends on detectability of the augmented system. Defining $\boldsymbol{z} = (\boldsymbol{x}, \boldsymbol{p})$, the augmented SDE has the form

$$\mathrm{d}\boldsymbol{z}(t) = f_z(\boldsymbol{z}(t), u(t); \tilde{p})\, \mathrm{d}t + g_z(\boldsymbol{x}(t), u(t); \tilde{p})\, \mathrm{d}\boldsymbol{\omega}, \quad (12)$$

where

$$f_z(\boldsymbol{z}, u; \tilde{p}) = \begin{pmatrix} f(\boldsymbol{x}, u; p, \tilde{p}) \\ 0 \end{pmatrix}, g_z(\boldsymbol{x}, u; \tilde{p}) = \begin{pmatrix} g(\boldsymbol{x}, u; \tilde{p}) & 0 \\ 0 & \Sigma \end{pmatrix}.$$

$\tilde{p}$ denotes the remaining parameters that are not estimated. Using this approach, the parameters are represented as disturbance states of the system since the observation equation is still given by

$$\boldsymbol{y}(t_k) = h_z(\boldsymbol{z}(t_k)) + \boldsymbol{v}_k = h(\boldsymbol{x}(t_k)) + \boldsymbol{v}_k. \quad (14)$$

The online parameter estimates are given by the filtered values from the augmented state vector, i.e. $\hat{z}_{k|k} = (\hat{x}_{k|k}; \hat{p}_{k|k})$. This augmentation method may be used for parameter estimation as well as disturbance estimation in offset free control (Morari and Maeder, 2012).

3.2 Offline identification using an ML formulation

Another use of the CDEKF is to estimate the parameters for a batch of data (Kristensen et al., 2004) in an offline optimization. The parameter estimates are the parameter set that maximizes the likelihood of the one-step prediction errors.

Let $\{y_j\}_{j=1}^N$ denote N observations relating to the sample points $\{t_j\}_{j=1}^N$ in (1b). Define the information accumulated up until the k-th sample point as $\mathcal{Y}_k = \{y_j\}_{j=1}^k$. Then the likelihood function, $\mathcal{L}$, can be defined as

$$\mathcal{L}(p \mid \mathcal{Y}_N) \propto \phi(\mathcal{Y}_N \mid p), \quad (15)$$

where ϕ is the joint density function of the observations, $\mathcal{Y}_N$. Using the definition of conditional probabilities, the right hand side can be rewritten into

$$\phi(\mathcal{Y}_N \mid p) = \prod_{k=1}^N \phi(y_k \mid \mathcal{Y}_{k-1}, p), \quad (16)$$

such that the log-likelihood function can be expressed by

$$\log(\mathcal{L}(p \mid \mathcal{Y}_N)) = \sum_{k=1}^N \log(\phi(y_k \mid \mathcal{Y}_{k-1}, p)). \quad (17)$$

Consequently, the ML parameter estimates, p_{ML}, is given by

$$p_{\mathrm{ML}} = \arg\max_{p \in \mathbb{R}^{N_p}} \log(\mathcal{L}(p \mid \mathcal{Y}_N)) \quad (18a)$$

$$= \arg\max_{p \in \mathbb{R}^{N_p}} \sum_{k=1}^N \log(\phi(y_k \mid \mathcal{Y}_{k-1}, p)). \quad (18b)$$

Since the SDE in (1a) is driven by a Brownian motion and since the increments of a Brownian motion are Gaussian it is reasonable to assume, under some regularity conditions, that the conditional densities in (16) can be well approximated by Gaussian densities

$$\phi(y_k \mid \mathcal{Y}_{k-1}, p) = \frac{\exp\left(-\frac{1}{2} e_k' R_{k|k-1}^{-1} e_k\right)}{\sqrt{\det(R_{k|k-1})(2\pi)^{n_y}}}, \quad (19)$$

where n_y is the number of output variables.

4. NUMERICAL OPTIMAL CONTROL

In this section, we briefly present the algorithm for optimization of (2) used by the NMPC considered in this tutorial paper. The algorithm is based on a direct local collocation method presented by von Stryk (1993).

For simplicity, we let the Mayer term of (2) be zero, i.e. $l_f(x(t_k + T)) = 0$, and only consider optimal control problems with a Lagrange term. We consider only bounds on the inputs. This implies that $\mathcal{U}(t)$ denotes these bound constraints, i.e.

$$\mathcal{U}(t) = [u_{\min}, u_{\max}]. \quad (20)$$

4.1 A local collocation method

When formulating a direct solution method, the first step is to introduce a parametrization of the manipulated variable, $u(t)$. The simplest parametrization is to approximate $u(t)$ as a piecewise constant function. For reasons of simlicity, we adopt this parametrization method in the following. Hence, $u(t)$ is parametrized via the values $\{q_k\}_{k=1}^N$ and time points $\{\tau_k\}_{k=0}^N$ such that

$$u = \sum_{k=1}^N q_k \chi_{[\tau_{k-1}, \tau_k[}, \quad (21)$$

where χ_I denotes the characteristic function associated with the set I. The time points that define the subintervals of $u(t)$ also constitute the global collocation points (GCPs) of the collocation method. Fig. 2 provides a schematic overview of the GCPs in relation to the local collocation points (LCPs), $\{\gamma_j\}_{j=0}^M$. The collocation points satisfy the relations

$$t_k = \tau_0 < \tau_1 < \cdots < \tau_N = t_k + T, \quad (22a)$$

$$0 = \gamma_0 < \gamma_1 < \cdots < \gamma_M = 1. \quad (22b)$$

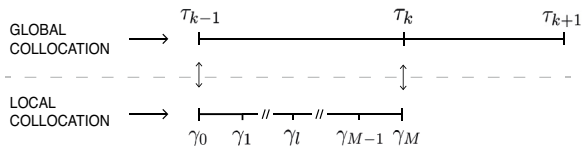

Fig. 2. Schematic overview of the collocation points.

A local collocation method approximates the ODEs on the smaller subintervals and apply a quadrature rule to impose a finite dimensional approximation. The ODEs on the interval $[\tau_{k-1}, \tau_k]$ can be formulated as an integral equation of the form

$$\boldsymbol{x}(\tau_k) - x(\tau_{k-1}) = \int_{\tau_{k-1}}^{\tau_k} f(\boldsymbol{x}(t), q_k; p)\,\mathrm{d}t. \tag{23}$$

Using the forward Euler scheme, the right hand side can be approximated via the LCPs by

$$\int_{\tau_{k-1}}^{\tau_k} f(\boldsymbol{x}(t), q_k; p)\,\mathrm{d}t \tag{24a}$$

$$= \sum_{j=1}^{M} \int_{\tau_{k-1}+\gamma_{j-1}\Delta\tau_k}^{\tau_{k-1}+\gamma_j\Delta\tau_k} f(\boldsymbol{x}(t), q_k; p)\,\mathrm{d}t \tag{24b}$$

$$\approx \Delta\tau_k \sum_{j=1}^{M} \Delta\gamma_j f(\boldsymbol{x}(\tau_{k-1} + \gamma_{j-1}\Delta\tau_k), q_k; p), \tag{24c}$$

where $\Delta\tau_k = \tau_k - \tau_{k-1}$ and $\Delta\gamma_j = \gamma_j - \gamma_{j-1}$. Next, introduce the discrete state vector as

$$s_{k,j} = \boldsymbol{x}(\tau_{k-1} + \gamma_j\Delta\tau_k). \tag{25}$$

Using this notation, the local collocation scheme can be formulated as

$$s_{k-1,M} - s_{k-1,0} = \sum_{j=1}^{M} \Delta t_{k,j} f(s_{k-1,j-1}, q_k; p), \tag{26}$$

where $\Delta t_{k,j} = \Delta\tau_k\Delta\gamma_j$. (26) solves the ODEs on the local intervals - to obtain a meaningful ODE solution, a continuity condition must be imposed, together with an initial value constraint

$$s_{0,0} = x_0 \quad \text{and} \quad s_{k-1,M} = s_{k,0}. \tag{27}$$

Finally, the objective function is also approximated using the forward Euler method with the collocation points defined in (22)

$$\int_{t_k}^{t_k+T} l(\boldsymbol{x}(t), u(t))\,\mathrm{d}t \approx \sum_{k=1}^{N}\sum_{j=1}^{M} \Delta t_{k,j} l(s_{k-1,j-1}, q_k). \tag{28}$$

Thus, the finite dimensional NLP for numerical solution of the optimal control problem (2) can be defined as

$$\min_{s,q} \quad \sum_{k=1}^{N}\sum_{j=1}^{M} \Delta t_{k,j} l(s_{k-1,j-1}, q_k) \tag{29a}$$

$$s.t. \quad s_{0,0} = x_0 \tag{29b}$$

$$s_{k-1,M} = s_{k,0} \tag{29c}$$

$$s_{k-1,M} - s_{k-1,0} = \sum_{j=1}^{M} \Delta t_{k,j} f(s_{k-1,j-1}, q_k) \tag{29d}$$

$$q_k \in \mathcal{U}(\tau_k). \tag{29e}$$

We solve (29) using `ipopt` and `python`. The numerical implementation in `python` uses the `pyipopt` package to interface to `ipopt` (Wächter and Biegler, 2006). Using `ipopt` it is possible to exploit the sparse structure that appears in the Jacobian of the constraint function of (29).

5. NUMERICAL CASE STUDY

To illustrate the methodology presented in this paper, we use the stochastic van der Pol oscillator model that is defined as

$$\mathrm{d}\boldsymbol{x}(t) = f(\boldsymbol{x}(t), u(t); \lambda)\,\mathrm{d}t + g(\boldsymbol{x}(t), u(t); \sigma)\,\mathrm{d}\boldsymbol{\omega}(t), \tag{30a}$$

where

$$f(\boldsymbol{x}, u; \lambda) = \begin{pmatrix} x_2 \\ -x_1 + \lambda(1 - x_1^2)x_2 + u \end{pmatrix}, \tag{30b}$$

$$g(\boldsymbol{x}, u; \sigma) = \begin{pmatrix} 0 & 0 \\ 0 & \sigma \end{pmatrix}. \tag{30c}$$

$\lambda > 0$ is a parameter governing the stiffness of the system and $\sigma > 0$ is a parameter related to model deficiency.The SDE model (30) is used as a model for simulating the plant as well as in the CDEKF of the NMPC. Fig. 3 provides a comparison between three realizations of the SDE (30) and the corresponding ODE. The effect of the added noise to x_2 is clearly visible.

5.1 Set-point tracking using online parameter estimation

The set-point trajectory, $\overline{x}_1(t)$, for $x_1(t)$ is defined as the step function

$$\overline{x}_1(t) = \begin{cases} 0, & t \leq 15 \\ 1, & 15 \leq t \leq 30 \\ 0, & 30 \leq t \end{cases}, \tag{31}$$

and the corresponding integrand of the control objective, l, is defined as

$$l(x(t), u(t)) = (1-\alpha)(x_1(t) - \overline{x}_1(t))^2 + \alpha u(t)^2, \tag{32}$$

where $\alpha = 1/1000$ is regularization parameter. The noise parameters are defined as

$$R_k = \sigma_\varepsilon^2 I, \quad \sigma_\varepsilon = 1/100, \quad \sigma = 1/5. \tag{33}$$

It is assumed that both states are directly observable, i.e. $h(x(t)) = x(t)$. The NLP has been constructed with an equidistant mesh such that

$$\Delta\tau_k = 1/(N-1), \quad \Delta\gamma_k = 1/(M-1), \tag{34}$$

where $N = 51$ and $M = 21$. The control and prediction horizon of the optimal control problem (2) is $T = T_c = T_p = 20$. It is assumed that observations occur equidistantly with $T_s = t_k - t_{k-1} = 0.4$. The control signals are constrained by the sets, $\mathcal{U}(\tau_k) = [-1, 1]$.

Fig. 4 shows a closed-loop simulation. For this simulation, the controller has to estimate λ. Fig. 5 shows the true value of λ as well as the online parameter estimate, $\hat{\lambda}_{k|k}$. Fig. 5 also shows how the online estimation method performs when an unmodelled disturbance is introduced. The unmodelled disturbance is introduced at $t = 45$ where the true value of λ is changed from $\lambda = 1$ to $\lambda = 3$.

Fig. 6 shows the offline parameter estimates and the corresponding log-likelihood functions based on 100 observations of the plant. The offline estimation is tested in two cases; a case where $\lambda = 1$ and a case where $\lambda = 10$. The parameters are estimated to be $\hat{\lambda} = 1.017$ and $\hat{\lambda} = 9.953$, respectively.

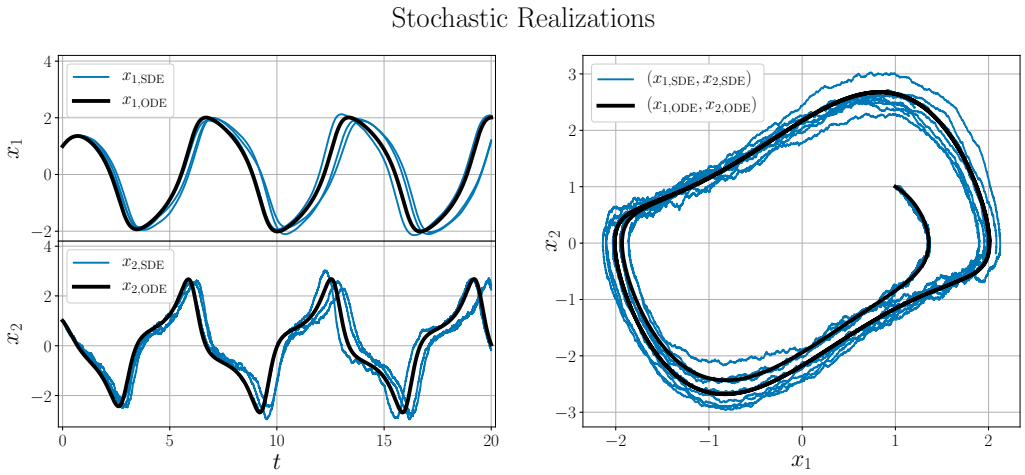

Fig. 3. The difference between an ODE realization (deterministic) and three different realizations of the SDE (stohcastic).

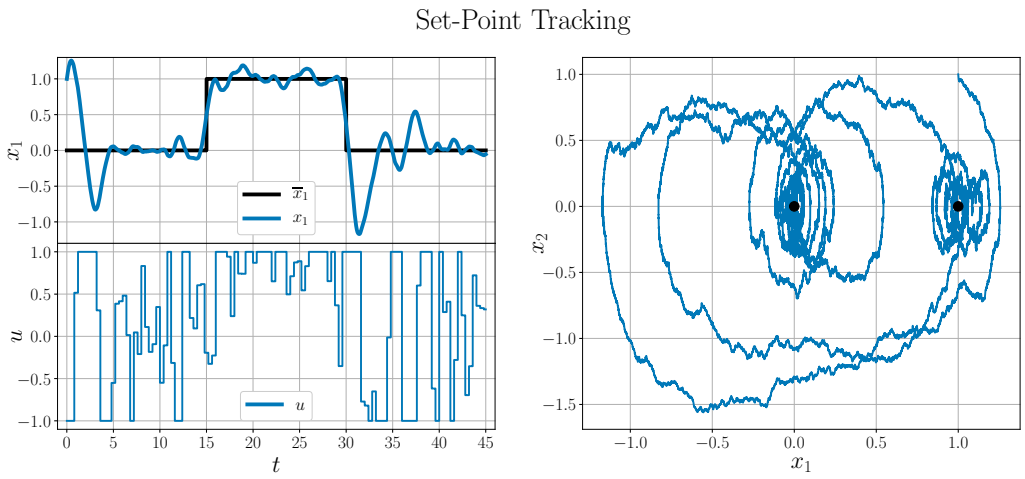

Fig. 4. A closed-loop simulation where x_1 has to track the set-point trajetory given in (31). The top-left plot shows the true state values of x_1 together with the set-point trajectory. The bottom-left plot shows the manipulated variables, $u(t)$, computed by the NMPC. The phase plot to the right illustrates the path the NMPC chooses when a set-point change is imposed.

The results presented in Fig. 4 and Fig. 5 show that the NMPC is able to simultaneously control the system to the desired set-point trajectory (31) and accurately estimate the unknown parameter, λ. However, from Fig. 5 it is seen that the parameter estimate is sensitive to the set-point change. For the estimates around $t \in \{15, 30\}$ in Fig. 5, the EKF estimates significant parameter changes despite of the fact that the true value is kept constant at $\lambda = 1$. The results from Fig. 4 also show that the resulting control signal, $u(t)$, is active around (and on) the upper bound when $\overline{x}_1(t) = 1$. This is a result of the fact that $(1, 0)$ is not an equilibrium for (30). Hence, the controller has to actively change x_2 to keep x_1 close to $\overline{x}_1$.

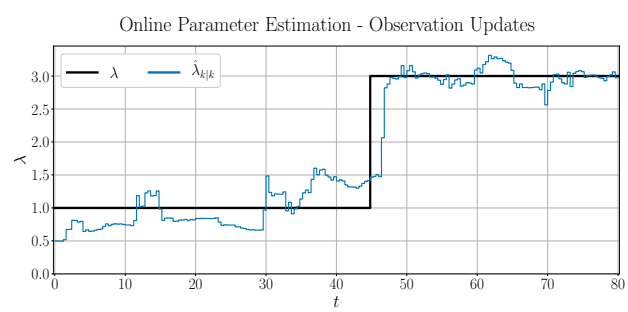

Fig. 5. The online parameter estimates of λ, where at $t = 45$ an un-modelled disturbance is introduced (λ shifts from 1 to 3).

6. SUMMARY

We provide a tutorial overview of how to construct an NMPC to regulate a stochastic system governed by SDEs. Based on the CDEKF, an online and an offline method for parameter estimation are presented. The dynamic optimization module is based on a local collocation scheme, where the forward Euler method has been used as discretization method for the dynamical equations. The performance of the closed-loop controller is investigated for a stochastic extension to the van der Pol oscillator model. The source code for the tutorial is available via `GitHub`. The key contribution and insight is to use the continuous-discrete model (1) and the same associated CDEKF in all components of the NMPC.

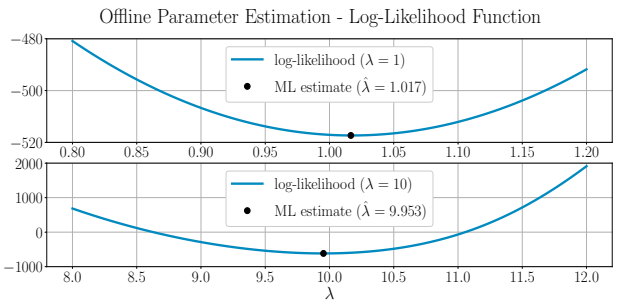

Fig. 6. Examples of offline parameter estimation of λ. Two cases are shown; the top plot shows an offline estimation problem where the true value is $\lambda = 1$ and the bottom plot shows an offline estimation problem where the true value is $\lambda = 10$. Both problems are based on 100 observations with the noise parameters given in (33).

REFERENCES

Allgöwer, F., Badgwell, T.A., Qin, J., Rawlings, J., and Wright, S.J. (1999). Nonlinear predictive control and moving horizon estimation - an introductory overview. In *Advances in Control. Highlights of ECC'99*. Springer, London, GB.

Binder, T., Blank, L., Bock, H.G., Bulirsch, R., Dahmen, W., Diehl, M., Kronseder, T., Marquardt, W., Schlöder, J.P., and von Stryk, O. (2001). Introduction to model based optimization of chemical processes on moving horizons. In M. Grötschel, S.O. Krumke, and J. Rambau (eds.), *Online Optimization of Large Scale Systems*. Springer.

Boiroux, D., Hagdrup, M., Mahmoudi, Z., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016a). An ensemble nonlinear model predictive control algorithm in an artificial pancreas for people with type 1 diabetes. In *2016 European Control Conference (ECC)*, 2115–2120.

Boiroux, D., Hagdrup, M., Mahmoudi, Z., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016b). Model identification using continuous glucose monitoring data for type 1 diabetes. *IFAC-PapersOnline*, 49(7), 759–764.

Boiroux, D., Juhl, R., Madsen, H., and Jørgensen, J.B. (2016c). An efficient square root algorithm for the computation of maximum likelihood sensitivity of continuous-discrete systems. In *55th IEEE Conference on Decision and Control (CDC)*, 3048–3053. IEEE.

Boiroux, D., Finan, D.A., Jørgensen, J.B., Poulsen, N.K., and Madsen, H. (2010). Meal estimation in nonlinear model predictive control for type 1 diabetes. *IFAC Proceedings Volumes*, 43(14), 1052–1057.

Buehler, E.A., Paulson, J.A., and Mesbah, A. (2016). Lyapunov-based stochastic nonlinear model predictive control: Shaping the state probability distribution functions. In *American Control Conference (ACC), 2016*, 5389–5394. IEEE.

Garnier, H. and Young, P.C. (2012). What does continuous-time model identification have to offer? *IFAC Proceedings Volumes*, 45(16), 810–815.

Grüne, L. and Pannek, J. (2011). *Nonlinear Model Predictive Control. Theory and Algorithms*. Springer, London, GB.

Jazwinski, A.H. (1970). *Stochastic Processes and Filtering Theory*. Academic Press, San Diego, CA, USA.

Johansen, T.A. (2011). Introduction to nonlinear model predictive control and moving horizon estimation. In M. Huba, S. Skogestad, M. Fikar, M. Hovd, T.A. Johansen, and B. Rohal-Ilkiv (eds.), *Selected Topics on Constrained and Nonlinear Control*, chapter 5, 187–240. STU Bratislava - NTNU Trondheim, Bratislava, Slovakia.

Jørgensen, J.B. and Jørgensen, S.B. (2007a). Comparison of prediction-error-modelling criteria. *Proceedings of the 2007 American Control Conference*, 140–146.

Jørgensen, J.B. and Jørgensen, S.B. (2007b). MPC-relevant prediction-error identification. *Proceedings of the 2007 American Control Conference*, 128–133.

Kristensen, N.R., Madsen, H., and Jørgensen, S.B. (2004). Parameter estimation in stochastic grey-box models. *Automatica*, 40(2), 225–237.

Mahmood, M. and Mhaskar, P. (2012). Lyapunov-based model predictive control of stochastic nonlinear systems. *Automatica*, 48(9), 2271–2276.

Mahmoudi, Z., Boiroux, D., Hagdrup, M., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016a). Application of the continuous-discrete extended Kalman filter for fault detection in continuous glucose monitors for type 1 diabetes. In *2016 European Control Conference (ECC)*, 714–719.

Mahmoudi, Z., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2017). Fault and meal detection by redundant continuous glucose monitors and the unscented kalman filter. *Biomedical Signal Processing and Control*, 38, 86–99.

Mahmoudi, Z., Wendt, S.L., Boiroux, D., Hagdrup, M., Nørgaard, K., Poulsen, N.K., Madsen, H., and Jørgensen, J.B. (2016b). Comparison of three nonlinear filters for fault detection in continuous glucose monitors. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 3507–3510.

Morari, M. and Maeder, U. (2012). Nonlinear offset-free model predictive control. *Automatica*, 48(9), 2059–2067.

Rao, G.P. and Unbehauen, H. (2006). Identification of continuous-time systems. *IEE Proceedings - Control Theory and Applications*, 153(2), 185–220.

Rawlings, J.B., Mayne, D.Q., and Diehl, M.M. (2017). *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, Madison, WI, USA.

Schneider, R. and Georgakis, C. (2013). How to NOT make the extended Kalman filter fail. *Industrial & Engineering Chemistry Research*, 52(9), 3354–3362.

Simon, D. (2006). *Optimal state estimation: Kalman, H_∞, and nonlinear approaches*. John Wiley & Sons.

von Stryk, O. (1993). Numerical solution of optimal control problems by direct collocation. In R. Bulirsch, J. Stoer, and K. Well (eds.), *Optimal Control: Calculus of Variations, Optimal Control Theory and Numerical Methods*, 129–143. Birkhäuser Basel, Basel.

Wächter, A. and Biegler, L.T. (2006). On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1), 25–57.