Estimates on the Minimal Stabilizing Horizon Length in Model Predictive Control for the Fokker-Planck Equation ⋆

A. Fleig * **L. Grüne** *

* *Department of Mathematics, University of Bayreuth, Germany (e-mail: arthur.fleig@uni-bayreuth.de, lars.gruene@uni-bayreuth.de)*

Abstract: In a series of papers by Annunziato and Borzì, Model Predictive Control of the Fokker-Planck equation has been established as a numerically feasible way for controlling stochastic processes via their probability density functions. Numerical simulations suggest that the resulting controller yields an asymptotically stable closed loop system for optimization horizons looking only one time step into the future. In this paper we provide a formal proof of this fact for the Fokker-Planck equation corresponding to the controlled Ornstein-Uhlenbeck process using an L^2 cost and control functions which are constant in space. The key step of the proof consists in the verification of an exponential controllability property with respect to the stage cost. Numerical simulations are provided to illustrate our results.

Keywords: Model predictive control, Optimal control, Partial differential equations, Stabilizing feedback, Stochastic processes

1. INTRODUCTION

There are various approaches to the control of stochastic processes. In this paper, we analyze an indirect approach introduced in Annunziato and Borzì (2010, 2013), in which the probability density function (PDF) rather than the state of an individual stochastic process is controlled in an optimal way. Since the evolution of the PDF is determined by the Fokker-Planck equation, this leads to an optimal control problem governed by this parabolic partial differential equation (PDE). The approach has similarities to the approach of solving stochastic optimal control problems via the Hamilton-Jacobi-Bellman approach, see Annunziato et al. (2014), however, it differs from this approach in that the optimal control is derived by optimizing the solution of a PDE rather than deriving the optimal control from the solution of a PDE. This allows to approach the control problem from a more global point of view, controlling the collective statistical behavior of the system as represented by the probability density function rather than optimizing the individual behavior of the systems. This global view also allows to use different classes of control functions, like functions which do not depend on space, i.e., control inputs which are independent of the current state of the stochastic process and are thus particularly easy to implement. This class of functions was used in Annunziato and Borzì (2010, 2013) and is also considered in this paper.

Controlling the state of the Fokker-Planck equation asymptotically to a desired equilibrium PDF can be formulated as an infinite horizon optimal control problem. Solving such a problem directly is, however, a very difficult and often computationally infeasible task. Model Predictive Control (MPC), also known as Receding Horizon Control, provides a way to circumvent these computational difficulties by splitting up the problem into the iterative solution of finite horizon problems, see, e.g., the monographs Grüne and Pannek (2011) or Rawlings and Mayne (2009). Obviously, the length of the needed finite optimization horizon directly influences the numerical effort needed for solving these problems: the shorter the horizon, the faster the numerical solution. On the other hand, long horizons may be needed on order to obtain stability of the resulting MPC closed loop. While numerical results in Annunziato and Borzì (2010, 2013) indicate that for the Fokker-Planck equation very short optimization horizons are sufficient for obtaining stability, a formal proof of this fact is to the best of our knowledge missing up to now.

In this paper, we close this gap for the Fokker-Planck equation corresponding to the controlled Ornstein-Uhlenbeck process using an L^2 cost and control functions which are constant in space. We show that for normally distributed PDFs, stability can always be achieved using an optimization horizon which looks only one time step into the future, thus resulting in the simplest possible optimal control problem. Our analysis relies on an exponential controllability condition for the considered stage cost, which is established for different cases depending on the ratio of the variance of the initial PDF to the drift and diffusion coefficients in the Fokker-Planck equation, employing a suitably chosen equivalent stage cost for one of the cases. The remainder of the paper is organized as follows. In Section 2 we introduce the concept of Model Predictive Control including the exponential controllability condition and the stability result we are going to use. Section 3 defines the problem setting, particularly the Fokker-Planck

⋆ This work was supported by the DFG project Model Predictive Control for the Fokker-Planck Equation, GR 1569/15-1. The paper was written while the second author was visiting the University of Newcastle, Australia.

equation we are going to control. Section 4 derives the main stability result by checking the exponential controllability condition. Our results are illustrated by a numerical example in Section 5 before we conclude the paper in Section 6.

2. MODEL PREDICTIVE CONTROL

In this section, we briefly present the concept of MPC. A more detailed introduction can be found in the monographs Grüne and Pannek (2011) and Rawlings and Mayne (2009).

We consider nonlinear discrete time control systems

$$z(k+1) = f(z(k), u(k)), \quad z(0) = z_0, \tag{1}$$

with $k \in \mathbb{N}_0$, state $z(k) \in \mathbb{X} \subset Z$ and control $u(k) \in \mathbb{U} \subset U$, where Z and U are metric spaces. State and control constraints are incorporated in $\mathbb{X}$ and $\mathbb{U}$, respectively. Continuous time models are sampled using a (constant) sampling rate $T_s > 0$. Given an initial state z_0 and a control sequence $(u(k))_{k\in\mathbb{N}_0}$, the solution trajectory is denoted by $z_u(\cdot; z_0)$.

Stabilization and tracking problems can be recast as infinite horizon optimal control problems using a tracking type cost function. However, solving infinite horizon optimal control problems governed by PDEs is computationally hard. The idea behind MPC is to circumvent this issue by iteratively solving optimal control problems on a shorter, finite time horizon, resulting in a feedback control $\mu\colon \mathbb{X} \to \mathbb{U}$ for the closed loop system

$$z_\mu(k+1) = f(z_\mu(k), \mu(z_\mu(k))). \tag{2}$$

Instead of minimizing a cost functional

$$J_\infty(z_0, u) := \sum_{k=0}^{\infty} \ell(z_u(k; z_0), u(k)), \tag{3}$$

the finite horizon cost functional

$$J_N(z_0, u) := \sum_{k=0}^{N-1} \ell(z_u(k; z_0), u(k)) \tag{4}$$

is minimized, where $N \geq 2$ is the optimization horizon length. The continuous function $\ell\colon Z \times U \to \mathbb{R}_{\geq 0}$ defines the stage costs, also called running costs. The feedback law μ is constructed through the following steps:

1. Given an initial value $z_\mu(0) \in \mathbb{X}$, fix the length of receding horizon N and set $n = 0$.
2. Initialize the state $z_0 = z_\mu(n)$ and minimize (4) subject to (1). Apply the first value of the resulting optimal control sequence denoted by $u^* \in \mathbb{U}^N$, i.e., set $\mu(z_\mu(n)) := u^*(0)$.
3. Evaluate $z_\mu(n+1)$ according to relation (2), set $n := n+1$ and go to step 2.

By truncating the infinite horizon, an important question that arises is whether the MPC closed loop system is asymptotically stable. One way to enforce stability is to add terminal conditions to (4). In the PDE setting, this approach has been investigated, e.g., by Ito and Kunisch (2002); Dubljevic et al. (2006); Dubljevic and Christofides (2006). Terminal constraints are added to the state constraints $\mathbb{X}$, terminal costs influence the cost functional J_N. However, constructing a suitable terminal region or finding an appropriate terminal cost is a challenging task. MPC schemes that do not rely on these methods are much easier to set up and implement and are therefore often preferred in practical applications. In this case, the choice of the horizon length N in step 1 of the MPC algorithm is crucial: Longer horizons make the problem computationally harder, shorter horizon lengths may lead to instability of the MPC closed loop system. Therefore, the smallest horizon that yields a stabilizing feedback is of particular interest, both from the theoretical and practical point of view.

Similar to Altmüller and Grüne (2012), the study in this work relies on a stability condition proposed in Grüne and Pannek (2011) that, together with the exponential controllability assumption below, ensures the relaxed Lyapunov inequality, cf. (Grüne and Pannek, 2011, Thm. 6.14 and Prop. 6.17). This inequality has been introduced in Lincoln and Rantzer (2006) to guarantee stability of the MPC closed loop solution.

Definition 1. The system (1) is called exponentially controllable with respect to the stage costs $\ell :\Leftrightarrow \exists C \geq 1, \rho \in (0,1)\ \forall \mathring{z} \in Z\ \exists u_{\mathring{z}} \in U$:

$$\ell(z_{u_{\mathring{z}}}(n; \mathring{z}), u_{\mathring{z}}(n)) \leq C\rho^n \min_{u\in U} \ell(\mathring{z}, u) \tag{5}$$

for all $n \in \mathbb{N}_0$.

Using the stability condition from Grüne and Pannek (2011), the minimal stabilizing horizon can be deduced from the values of the overshoot bound C and the decay rate ρ. For more details, in particular on the influence of C and ρ, see Altmüller and Grüne (2012). The most important difference in the influence of C and ρ for our study is that for fixed C, it is generally impossible to arbitrarily reduce the horizon N by reducing ρ. However, for $C = 1$, stability can be ensured even for the shortest possible horizon $N = 2$. More precisely, (Grüne and Pannek, 2011, Theorem 6.18 and Section 6.6) yield the following theorem.

Theorem 2. Consider the MPC scheme with stage costs (15) satisfying the exponential controllability property from Definition 1 with $C = 1$ and $\rho \in (0,1)$. Then the equilibrium y_{eq} from (15) is globally asymptotically stable for the MPC closed loop for each optimization horizon $N \geq 2$.

In the subsequent analysis, we will therefore try to find a (suboptimal) control $u_{\mathring{z}}$ that satisfies exponential controllability with $C = 1$.

3. PROBLEM SETTING

We consider the Cauchy problem

$$\partial_t y - \sum_{i,j=1}^{d} \partial^2_{ij}(a_{ij} y) + \sum_{i=1}^{d} \partial_i (b_i(u) y) = 0 \text{ in } Q \tag{6}$$

$$y(\cdot, 0) = y_0 \tag{7}$$

where $a_{ij}\colon Q \to \mathbb{R}$, $b_i\colon Q \times U \to \mathbb{R}$ are given functions for $Q := \mathbb{R}^d \times (0,T)$, $y_0\colon \mathbb{R}^d \to \mathbb{R}$ is the initial state and $y\colon Q \to \mathbb{R}$ is the unknown. The control u is acting on the drift term and can be a function of time and or space.

Equation (6) is called the Fokker-Planck equation, or Forward Kolmogorov equation. Under appropriate assumptions, cf. (Primak et al., 2004, p. 227) and (Protter, 2005,

p. 297), it models the evolution of probability density functions associated with continuous-time stochastic processes described by the Itô stochastic differential equation

$$dX_t = b(X_t, t; u)dt + \sigma(X_t, t)dW_t, \quad t \in (0, T) \quad (8)$$

with initial condition $X_0 \in \mathbb{R}^d$. Here, $dW_t \in \mathbb{R}^m$ is an m-dimensional Wiener process, $m \in \mathbb{N}$, and $b = (b_1, ..., b_d)$ is the vector valued drift function, and the diffusion matrix $\sigma(X_t, t) \in \mathbb{R}^{d\times m}$ has full rank. The coefficient functions a_{ij} in (6) are related to this matrix via $a_{ij} = \sum_k \sigma_{ik}\sigma_{jk}/2$. For an exhaustive theory and more details on the connection between stochastic processes and the Fokker-Planck equation, we refer to Risken (1989). Since this equation models the evolution of a PDF y, y needs to satisfy the standard properties of a PDF, i.e.,

$$y(x,t) \geq 0, (x,t) \in Q, \quad \int_{\mathbb{R}^d} y(x,t)\, dx = 1, t \in (0,T). \quad (9)$$

In the subsequent analysis, we consider the Ornstein-Uhlenbeck process in 1D and the multidimensional version thereof. In this case, the diffusion is constant and diagonal, and is given by

$$a_{ij} := \delta_{ij}\sigma_i^2/2, \quad (10)$$

where $\sigma_i > 0$, and δ_{ij} is the Kronecker delta. The drift is defined by

$$b_i(x, t; u) := -\mu_i x + u_i \quad (11)$$

for $\mu_i > 0$ and $u_i \in \mathbb{R}$. We assume that the initial probability density function is that of a d-dimensional multivariate normal distribution with mean vector $\mathring{\mu}$ and covariance matrix $\Sigma_{ij} = \delta_{ij}\mathring{\sigma}_i^2$, i.e.,

$$y_0(x) = \left((2\pi)^d \prod_{i=1}^d \mathring{\sigma}_i^2\right)^{-1/2} \exp\left(-\sum_{i=1}^d \frac{(x_i - \mathring{\mu}_i)^2}{2\mathring{\sigma}_i^2}\right) \quad (12)$$

In this case, the solution of the Fokker-Planck equation exists in closed form and is given by

$$y(x,t) = \frac{\exp\left(-\sum_{i=1}^d \frac{\left(x_i - \left[\frac{u_i}{\mu_i} + \left(\mathring{\mu}_i - \frac{u_i}{\mu_i}\right)e^{-\mu_i(t-t_0)}\right]\right)^2}{2\left(\frac{\sigma_i^2}{2\mu_i} + \left[\mathring{\sigma}_i^2 - \frac{\sigma_i^2}{2\mu_i}\right]e^{-2\mu_i(t-t_0)}\right)}\right)}{\sqrt{(2\pi)^d \prod_{i=1}^d \left(\frac{\sigma_i^2}{2\mu_i} + \left[\mathring{\sigma}_i^2 - \frac{\sigma_i^2}{2\mu_i}\right]e^{-2\mu_i(t-t_0)}\right)}}, \quad (13)$$

which converges to

$$\bar{y}(x; u) := \left((2\pi)^d \prod_{i=1}^d \frac{\sigma_i^2}{2\mu_i}\right)^{-1/2} \exp\left(-\sum_{i=1}^d \frac{\left(x_i - \frac{u_i}{\mu_i}\right)^2}{\frac{\sigma_i^2}{\mu_i}}\right) \quad (14)$$

as $t \to \infty$. In particular, $y_{eq} := \bar{y}(\cdot; u_{eq})$ is an equilibrium solution of (6) for coefficient functions (10) and (11). Note that a space-independent control u influences only the drift term in the FP equation, i.e., the mean of the distribution, not its variance.

For the stage costs we consider

$$\ell(y(n), u(n)) := \frac{1}{2}\|y(\cdot, t_n) - y_{eq}\|^2_{L^2(\mathbb{R}^d)} + \frac{\lambda}{2}|u(t_n) - u_{eq}|^2, \quad (15)$$

where $|\cdot|$ is the Euclidean norm and t_n are the sampling times $t_n = nT_s$. For small values of the MPC sampling rate T_s, (4) with (15) approximates an $L^2(Q)$ tracking of the state. For $N = 2$, minimizing (4) with ℓ from (15) is equivalent to minimizing

$$\frac{1}{2}\|y(\cdot, T_s) - y_{eq}\|^2_{L^2(\mathbb{R}^d)} + \frac{\lambda}{2}|u(t_0) - u_{eq}|^2, \quad (16)$$

which, for $u_{eq} = 0$, is the cost used in Annunziato and Borzì (2010, 2013).

The L^2 norm is omnipresent in almost all PDE-constrained optimal control problems considered in scientific literature, with many existence and uniqueness theorems and results concerning optimality conditions, see for example Tröltzsch (2010). These can either be applied directly or have been extended to fit our setting, cf. Annunziato and Borzì (2010, 2013); Addou and Benbrik (2002); Fleig and Guglielmi (2015). Other options such as the Wasserstein metric possibly fit the problem of tracking a PDF better, cf. Jordan et al. (1998). However, to the best of our knowledge, they lack a sound theory regarding existence of optimal controls.

We remark that (15) fulfills $\ell(y_{eq}, u_{eq}) = 0$ and $l(y, u) > 0$ for $(y, u) \neq (y_{eq}, u_{eq})$. In optimal control, instead of $|u - u_{eq}|^2$, $|u|^2$ is commonly used. However, ℓ would then violate the condition $\ell(y_{eq}, u_{eq}) = 0$ that is necessary for Theorem 2 to hold. Employing $|u|^2$ in this case leads to so-called economic MPC, see Grüne (2013); Grüne et al. (2015) for more details. Investigating the stability of the MPC closed loop system in the economic MPC framework is beyond the scope of this work.

Our aim is to analyze exponential controllability w.r.t. the stage costs (15) according to Definition 1 in order estimate the minimal stabilizing horizon length depending on the constants C and ρ in (5). To simplify the presentation, we focus on the one-dimensional case. However, the results are easily extended to the multi-dimensional setting.

One promising candidate for an exponentially stabilizing control u_z in (5) is the constant control $u_z(n) = u(t_n) \equiv u_{eq}$. In this case, the analysis is simplified since the term penalizing the control in the running costs (15) vanishes and the left hand side of (5) given by $\ell(z_{u_{\mathring{z}}}(n; \mathring{z}), u_{\mathring{z}}(n)) = \ell(y(n), u_{eq})$ can be calculated explicitly:

$$2\ell(y(n), u_{eq}) = \|y(\cdot, t_n) - y_{eq}\|^2_{L^2(\mathbb{R})}$$
$$= \frac{\sqrt{\mu}}{\sqrt{2\pi\sigma^2}}\left(1 + \frac{1}{\sqrt{\gamma(t_n)}} - \frac{2\sqrt{2}\exp(-\delta(t_n))}{\sqrt{\gamma(t_n)+1}}\right), \quad (17)$$

where

$$\gamma(t) := 1 + (\alpha - 1)e^{-2\mu t} > 0, \quad (18)$$

$$\delta(t) := \frac{Ze^{-2\mu t}}{\gamma(t) + 1} \geq 0, \quad (19)$$

with

$$\alpha := 2\mu\mathring{\sigma}^2/\sigma^2 > 0, \quad (20)$$

$$Z := \frac{(\mathring{\mu} - \frac{u_{eq}}{\mu})^2}{\frac{\sigma^2}{\mu}} = \frac{(\mu\mathring{\mu} - u_{eq})^2}{\mu\sigma^2} \geq 0. \quad (21)$$

Exponential controllability follows from the inequality

$$V_\alpha(t) \leq Ce^{-Kt}V_\alpha(0), \quad (22)$$

for

$$V_\alpha(t) := 1 + \frac{1}{\sqrt{\gamma(t)}} - \frac{2\sqrt{2}\exp(-\delta(t))}{\sqrt{\gamma(t)+1}} \quad (23)$$

(we can cancel out the constant factor in (17)) with constants $C, K > 0$. For sampling times $t_n = nT_s$, ρ^n in (5) is then defined by $\rho^n = e^{-KnT_s}$.

4. STABILITY OF THE MPC CLOSED LOOP SOLUTION

We first give an interpretation of the parameters α and Z introduced in the previous section. The parameter Z indicates the distance between the initial mean of the distribution $\mathring{\mu}$ and μ in the dynamics. The former parameter, α, relates the initial variance $\mathring{\sigma}^2$ to σ^2. If $\alpha = 1$, the variance of the distribution does not change in time since $\mathring{\sigma}^2 = \sigma^2/(2\mu)$ in (13). For $\alpha < 1$, the variance of the distribution is increasing in time since $\mathring{\sigma}^2 < \sigma^2/(2\mu)$. Analogously, the variance of the distribution shrinks in time if $\alpha > 1$.

In order to conclude stability of the MPC closed loop solution from the exponential controllability condition (5), an exponentially stabilizing control needs to exist for the initial state $\mathring{z} = z_\mu(n) = y(t_n, \cdot)$ in every MPC iteration. Hence, the value of α may change from one step to the next, i.e. $\alpha_{n+1} \neq \alpha_n$, where α_n denotes the value of α in the n-th MPC iteration. It is important to note, however, that for space-independent control the sign of $\alpha_n - 1$ does not change with n. This is due to the monotone convergence of α_n to 1 we get from reformulating the change in the variance in (13),

$$\mathring{\sigma}^2_{n+1} = \frac{\sigma^2}{2\mu} + \left(\mathring{\sigma}^2_n - \frac{\sigma^2}{2\mu}\right) e^{-2\mu T_s},$$

to

$$\alpha_{n+1} = 1 + (\alpha_n - 1)e^{-2\mu T_s}. \tag{24}$$

In order to prove (22) we now consider the three cases $\alpha = 1$, $\alpha < 1$ and $\alpha > 1$ separately.

The case $\alpha = 1$: In this case, the shape of the PDF stays the same since the space-independent control u_{eq} can only move the PDF as a whole. We have

$$V_1(t) = 2 - 2e^{-Ze^{-2\mu t}/2} \tag{25}$$

and we can prove the following proposition.

Proposition 3. For $V_1(t)$, inequality (22) holds with $C = 1$ and $K = 2\mu e^{-Z/2}$.

Proof. We show $V_1'(t) \leq -KV_1(t)$ to conclude our assertion. To this end, consider

$$\begin{aligned}
&V_1'(t) + KV_1(t)\\
=& -4\mu \left(\frac{Z}{2}e^{-2\mu t}e^{-Ze^{-2\mu t}/2} - e^{-Z/2} + e^{-Z/2}e^{-Ze^{-2\mu t}/2}\right)\\
=& -4\mu \left(e^{-Ze^{-2\mu t}/2}\left[\frac{Z}{2}e^{-2\mu t} + e^{-Z/2}\right] - e^{-Z/2}\right)\\
=& -4\mu \left(e^{-XY}\left[XY + e^{-X}\right] - e^{-X}\right),
\end{aligned}$$

where $X := Z/2 \geq 0$ and $Y := e^{-2\mu t} \in (0, 1]$. For arbitrary but fixed X we define the C^∞ function $f(Y) := e^{-XY}\left(XY + e^{-X}\right) - e^{-X}$. It can easily be shown that $f(0) = 0$ and $f(1) \geq 0$. By calculating $f'(Y)$, one can show that $f(Y)$ is monotonously increasing on $(0, Y^*)$, with $Y^* := (1 - e^{-X})/X$ being the unique root of $f'(Y)$, and monotonously decreasing on $(Y^*, 1]$. Therefore, $f(Y) \geq 0$, which concludes the proof.

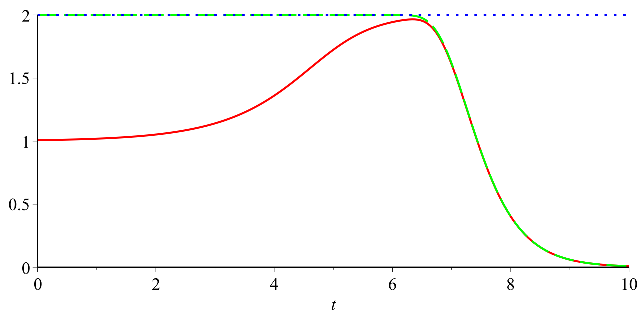

Fig. 1. $V_\alpha(t), \tilde{V}_\alpha(t)$, and $\tilde{V}_\alpha(0)e^{-Kt}$ for $\mu = \mathring{\mu} = \sigma = 1$, $\mathring{\sigma} = 100$, and $u_{eq} = 2000$

Since $C = 1$, the MPC closed loop system is asymptotically stable even for the shortest possible horizon N.

The case $\alpha < 1$: For $\alpha < 1$, since the shape of the PDF becomes wider in time, initially, the costs may be higher compared to $\alpha = 1$, i.e., $V_\alpha(0) \geq V_1(0)$. However, they also drop more quickly, i.e., $V_\alpha'(t) \leq V_1'(t)$. The idea is to prove

$$f(t) := V_1(0)V_\alpha(t) - V_\alpha(0)V_1(t) \leq 0, \tag{26}$$

since for $V_1(0) \neq 0$, w.l.o.g., we then get

$$V_\alpha(t) \leq \frac{V_\alpha(0)}{V_1(0)}V_1(t) \leq e^{-Kt}V_\alpha(0). \tag{27}$$

Obviously, $f(0) = 0$ and $\lim_{t\to\infty} f(t) = 0$. Analogously to the proof of Proposition 3, one can conclude the assertion by showing that $f(t)$ is monotonously decreasing on $[0, t^*)$ with t^* being the unique root of $f'(t)$ and monotonously increasing on (t^*, ∞) for some unique t^*. If $t^* < 0$, the argument still holds because of the monotonicity on $(t^*, \infty) \subset [0, \infty)$.

The case $\alpha > 1$: If $\alpha > 1$, because of the shrinking variance of the distribution, the running costs at the beginning may rise, i.e., $V_\alpha'(t) > 0$ for $t \in [0, t^*)$ for some $t^* > 0$. This is the case for example for $\mu = \mathring{\mu} = \sigma = 1$, $\mathring{\sigma} = 100$, and $u_{eq} = 2000$, cf. Figure 1. It is due to the L^2 norm used in the stage costs ℓ. Obviously, the exponential controllability condition (22) does not hold for $C = 1$. To circumvent this issue, we add $1 - \frac{1}{\sqrt{\gamma(t)}}$ to $V_\alpha(t)$, yielding

$$\tilde{V}_\alpha(t) := 2 - \frac{2\sqrt{2}\exp(-\delta(t))}{\sqrt{\gamma(t) + 1}}. \tag{28}$$

This new cost function fulfills all necessary requirements, e.g., nonnegativity. Furthermore, since $\gamma(t)$ does not depend on u, the cost $\tilde{\ell}$ corresponding to $\tilde{V}_\alpha$ yields the same optimal control as ℓ and Theorem 2 can be applied to $\tilde{\ell}$.

Proposition 4. For $\tilde{V}_\alpha(t)$ with $\alpha > 1$, inequality (22) holds with $C = 1$ and $K = \frac{2\mu}{\alpha+1}e^{-Z/(\alpha+1)}$.

Idea of Proof. The procedure is the same as in the proof of Proposition 3, i.e., we have

$$\tilde{V}_\alpha'(t) + K\tilde{V}_\alpha(t) = -2\sqrt{2}\mu f(Y) \tag{29}$$

where for $Y := e^{-2\mu t}$ and arbitrary, but fixed α and Z we define the C^1 function

$$f(Y) := -\frac{e^{-Z/(\alpha+1)}}{\alpha+1}\left(\sqrt{2} - \frac{2e^{-\frac{ZY}{2+(\alpha-1)Y}}}{\sqrt{2+(\alpha-1)Y}}\right)$$
$$+ \frac{Ye^{-\frac{ZY}{2+(\alpha-1)Y}}}{(2+(\alpha-1)Y)^{3/2}}\left(2Z - \frac{2ZY(\alpha-1)}{2+(\alpha-1)Y} + (\alpha-1)\right).$$

To show $f(Y) \geq 0$, we claim that $f(0), f(1) \geq 0$ and that f is monotonously increasing for $Y \in (0, Y^*)$ for some Y^*. Obviously, $f(0) = 0$. Furthermore,

$$f(1) = \underbrace{\frac{e^{-Z/(\alpha+1)}}{(\alpha+1)^{3/2}}}_{\geq 0} g(\alpha, Z) \tag{30}$$

where

$$g(\alpha, Z) := 2Z - \frac{2Z(\alpha-1)}{\alpha+1} + \alpha - 1 - \sqrt{2(\alpha+1)} + 2e^{-Z/(\alpha+1)}.$$

To show $f(1) \geq 0$, we prove $g(\alpha, Z) \geq 0$:

$$\frac{\partial g}{\partial Z}(\alpha, Z) = \frac{4 - 2e^{-Z/(\alpha+1)}}{\alpha+1} > 0, \tag{31}$$

$$g(\alpha, 0) = (\alpha+1) - \sqrt{2(\alpha+1)} > 0 \text{ for } \alpha > 1. \tag{32}$$

Next, we investigate the monotonicity of f. For $Z = 0$, we have

$$f'(Y; Z=0) = \frac{(\alpha-1)}{((\alpha-1)Y+2)^{3/2}}\left(1 - \frac{\frac{3}{2}(\alpha-1)Y}{(\alpha-1)Y+2} - \frac{1}{\alpha+1}\right), \tag{33}$$

i.e.,

$$f'(Y; Z=0) = 0 \quad \Leftrightarrow \quad Y = \frac{4\alpha}{\alpha^2 + 2\alpha - 3} =: Y^*. \tag{34}$$

Since

$$f'(0; Z) = \frac{(2Z + \alpha - 1)(\alpha + 1 - e^{-Z/(\alpha+1)})}{2\sqrt{2}(\alpha+1)} > 0, \tag{35}$$

$f(Y; Z=0)$ is monotonically increasing in $(0, Y^*)$. Since Y^* is the unique root of f', together with $f(0), f(1) \geq 0$ we conclude the assertion for $Z = 0$. Note that even if $Y^* \notin (0, 1]$ or if Y^* does not exist, the argument still holds.

In the case $Z > 0$, a simple formula for Y^* does not appear to exist. However, numerical computations indicate that a unique root Y^* of $f'(Y)$ exists also in this case, cf. Figure 2. Hence, we can conclude nonnegativity of $f(Y)$ for $Y \in [0, 1]$ also in this case, which completes the proof.

To summarize, in all three cases we can apply Theorem 2 in order to conclude asymptotic stability of the MPC closed loop solution for the shortest possible horizon $N = 2$.

5. NUMERICAL SIMULATIONS

For our numerical study, we examine the three cases presented in the previous section, i.e., $\alpha = 1$, $\alpha < 1$, and $\alpha > 1$. The Ornstein-Uhlenbeck process is considered on $Q := \Omega \times [0, 5]$ instead of $\mathbb{R}^d \times [0, 5]$, where $\Omega = (-7.5, 7.5)$. The model parameters are chosen such that the values of the PDF outside of Ω are negligible. Rather than using the explicit form (13), we solve the Fokker-Planck equation numerically, employing the Chang-Cooper scheme and the BDF2 scheme to discretize space and time, respectively, to get an approximation of second order, cf. Mohammadi and Borzì (2015).

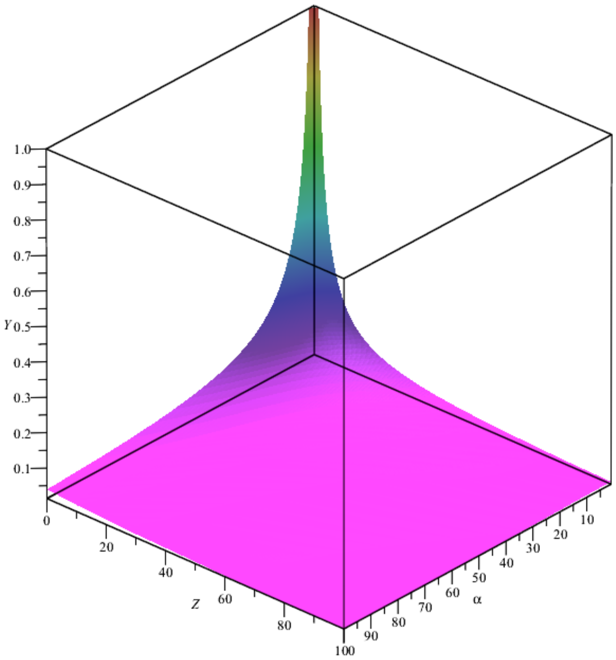

Fig. 2. Implicit plot of $f'(Y; \alpha, Z) = 0$

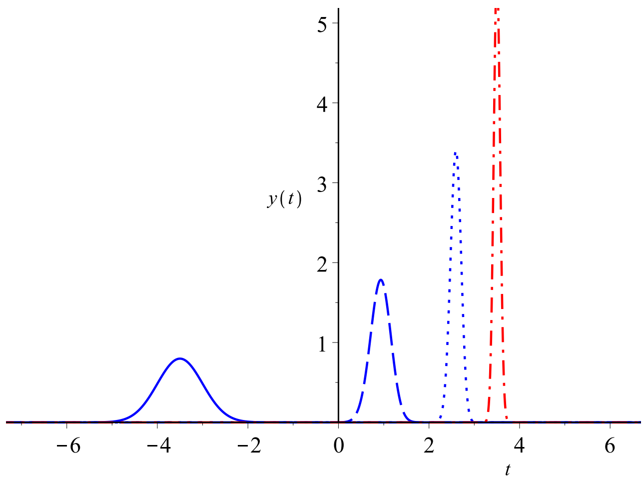

Fig. 3. $y(x, 0)$ (solid blue), $y(x, 1)$ (dashed blue), $y(x, 2)$ (dotted blue) and $y_{eq}(x)$ (dot-dashed red) for $\alpha > 1$

We fix $\mu = 1, \mathring{\mu} = -3.5$, and $u_{eq} = 3.5$. For $\alpha = 1$, we choose the remaining model parameters $(\sigma, \mathring{\sigma}) = (1/\sqrt{8}, 1/4)$. The cases $\alpha < 1$ and $\alpha > 1$ are modeled by $(\sigma, \mathring{\sigma}) = (0.5, 0.1)$ and $(\sigma, \mathring{\sigma}) = (0.1, 0.5)$, yielding $(\alpha, Z) = (0.08, 196)$ and $(\alpha, Z) = (50, 4900)$, respectively.

In the MPC algorithm, we only look one time step into the future. The sampling time T_s is 0.1. We use the costs defined by (16) with $\lambda = 0.25$. The optimal control problem is solved using a BFGS scheme. We employ necessary optimality conditions that are analogous to the ones derived in Annunziato and Borzì (2010).

Table 1. Total costs for u^* and u_{eq}

	$\alpha = 1$	$\alpha < 1$	$\alpha > 1$
u_{eq}	32.35	21.54	131.27
u^*	27.34 (-15,49%)	19.56 (-9.2%)	79.75 (-39.25%)

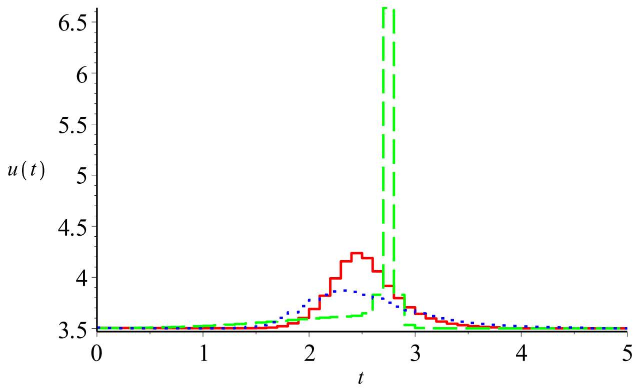

Fig. 4. Calculated optimal control u^* for $\alpha = 1$ (solid red), $\alpha < 1$ (dotted blue) and $\alpha > 1$ (dashed green)

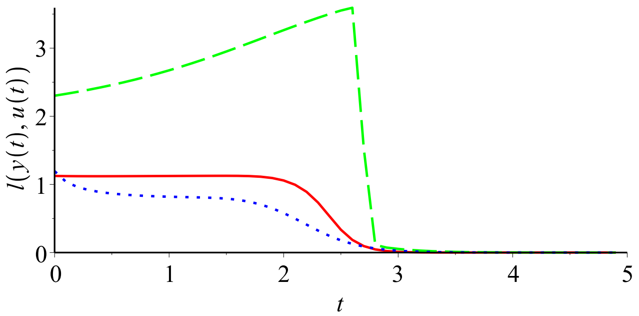

Fig. 5. Stage costs $\ell(y, u^*)$ for $\alpha = 1$ (solid red), $\alpha < 1$ (dotted blue) and $\alpha > 1$ (dashed green)

For $\alpha > 1$, Figure 3 shows the PDF $y(t)$ at various times and the equilibrium solution y_{eq}. The corresponding control and stage costs in every MPC iteration are displayed in Figures 4 and 5, respectively. In these two figures, the cases $\alpha = 1$ and $\alpha < 1$ are included. In all three cases, the optimal control stays near u_{eq} until the PDF y is close enough to y_{eq} such that a higher value of the control helps reaching the target faster at reasonable cost. The costs in Figure 5 develop as predicted. Table 1 displays the total costs for the optimal control $u = u^*$ and $u = u_{eq}$, showing the sub-optimality of u_{eq}. In conclusion, the numerical simulations coincide with our theoretical findings.

6. CONCLUSION

For the Ornstein-Uhlenbeck process and the multidimensional version thereof, we can conclude asymptotic stability of the MPC closed loop solution even for the shortest possible horizon. Our numerical simulations coincide with these findings.

It is only natural to extend these results for stochastic processes with no closed form solution, e.g., the Shiryaev process, to space-dependent controls, and to different cost functions. Regarding the latter, not only can different norms be employed, but one can replace the term penalizing the control, $|u-u_{eq}|^2$, by $|u|^2$, which, depending on the desired state and the stochastic process, leads to economic MPC.

REFERENCES

Addou, A. and Benbrik, A. (2002). Existence and uniqueness of optimal control for a distributed-parameter bilinear system. *J. Dynam. Control Systems*, 8(2), 141–152.

Altmüller, N. and Grüne, L. (2012). Distributed and boundary model predictive control for the heat equation. *GAMM-Mitt.*, 35(2), 131–145.

Annunziato, M. and Borzì, A. (2010). Optimal control of probability density functions of stochastic processes. *Math. Model. Anal.*, 15, 393–407.

Annunziato, M. and Borzì, A. (2013). A Fokker-Planck-based control framework for multidimensional stochastic processes. *J. Comput. Appl. Math.*, 237, 487–507.

Annunziato, M., Borzì, A., Nobile, F., and Tempone, R. (2014). On the connection between the hamilton-jacobi-bellman and the fokker-planck control frameworks. *Applied Mathematics*, 5(16), 2476–2484.

Dubljevic, S. and Christofides, P. (2006). Boundary predictive control of parabolic PDEs. In *American Control Conference, 2006*, 49–56.

Dubljevic, S., El-Farra, N.H., Mhaskar, P., and Christofides, P.D. (2006). Predictive control of parabolic PDEs with state and control constraints. *International Journal of robust and nonlinear control*, 16, 749–772.

Fleig, A. and Guglielmi, R. (2015). Bilinear optimal control of the Fokker-Planck equation. *preprint*.

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

Grüne, L. (2013). Economic receding horizon control without terminal constraints. *Automatica*, 49, 725–734.

Grüne, L., Semmler, W., and Stieler, M. (2015). Using nonlinear model predictive control for dynamic decision problems in economics. *Journal of Economic Dynamics & Control*, 60, 112–133.

Ito, K. and Kunisch, K. (2002). Receding horizon optimal control for infinite dimensional systems. *ESAIM: COCV*, 8, 741–760.

Jordan, R., Kinderlehrer, D., and Otto, F. (1998). The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1), 1–17.

Lincoln, B. and Rantzer, A. (2006). Relaxing dynamic programming. *IEEE Transactions on AutomaticControl*, 51(8), 1249–1260.

Mohammadi, M. and Borzì, A. (2015). Analysis of the Chang-Cooper discretization scheme for a class of Fokker-Planck equations. *Journal of Numerical Mathematics*, 23(3), 271–288.

Primak, S., Kontorovich, V., and Lyandres, V. (2004). *Stochastic methods and their applications to communications*. John Wiley & Sons, Inc., Hoboken, NJ. doi: 10.1002/0470021187.

Protter, P.E. (2005). *Stochastic integration and differential equations*, volume 21 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin.

Rawlings, J. and Mayne, D. (2009). *Model Predictive Control: Theory and Design*. Nob Hill Publishing.

Risken, H. (1989). *The Fokker-Planck equation*, volume 18 of *Springer Series in Synergetics*. Springer-Verlag, Berlin, second edition.

Tröltzsch, F. (2010). *Optimal Control of Partial Differential Equations: Theory, Methods, and Applications*. Graduate Studies in Mathematics. American Mathematical Society.