Available online at www.sciencedirect.com

ScienceDirect

IFAC-PapersOnLine 49-8 (2016) 260–265

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 that 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 ways to control stochastic processes. In this paper, we analyze an indirect approach introduced in Annunziato and Borzì (2010, 2013), in which the control influences the probability density function (PDF) associated with the stochastic process rather than the state of the process. 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). This approach allows to tackle 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, which is usually done by considering the average thereof, cf. Annunziato and Borzì (2010, 2013). In this setting, different classes of control functions can be used, e.g., those that do not depend on space, i.e., control inputs that 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.

The approach has similarities to solving stochastic optimal control problems via the Hamilton-Jacobi-Bellman approach, see Annunziato et al. (2014), the difference being 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 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.

A common task is to steer the state, in this case the PDF, asymptotically to a desired equilibrium. This 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 required for solving these problems: the shorter the horizon, the faster the numerical solution. On the other hand, long horizons may be needed in order to obtain stability of the resulting MPC closed loop, cf. Grüne and Pannek (2011); Altmüller and Grüne (2012). 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 that are constant in space but may be time-dependent. We show that for normally distributed PDFs, stability can always be achieved, even when looking only one time step into the future, thus resulting in the simplest possible optimal control problem with a constant control function in each MPC iteration. 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,

2405-8963 © 2016, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2016.07.451

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 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 in general computationally hard. The idea behind MPC is to circumvent this issue by iteratively solving optimal control problems on a shorter, finite time horizon and use the resulting optimal control values to construct a feedback law $\mu\colon \mathbb{X} \to \mathbb{U}$ 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 and 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 MPC scheme:

0. Given an initial value $z_\mu(0) \in \mathbb{X}$, fix the length of receding horizon $N \geq 2$ and set $n = 0$.
1. 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 $\mu(z_\mu(n)) := u^*(0)$.
2. Evaluate $z_\mu(n+1)$ according to relation (2), set $n := n+1$ and go to step 1.

By truncating the infinite horizon, an important question 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, cf. Grüne and Pannek (2011). MPC schemes that do not rely on these methods are much easier to set up and implement and are therefore often preferred in practice. In this case, the choice of the horizon length N in step 0 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. Finding it is the main task of this paper.

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 $\ell :\Leftrightarrow \exists C \geq 1, \rho \in (0,1)$ such that for each state $\mathring{z} \in Z$ there exists a control $u_{\mathring{z}} \in U$ satisfying

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

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

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, 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 meaningful horizon $N = 2$. Note that condition (5) depends on the running costs ℓ, which, in this paper, are given by

$$\ell(z(k), u(k)) = \frac{1}{2}\|z(k) - \bar{z}\|^2 + \frac{\lambda}{2}\|u(k) - \bar{u}\|^2 \tag{6}$$

for some norm $\|\cdot\|$, where $(\bar{z}, \bar{u})$ constitutes an equilibrium of (1), i.e., $f(\bar{z}, \bar{u}) = \bar{z}$. In particular, we have $\ell(\bar{z}, \bar{u}) = 0$ and $l(z, u) > 0$ for $(z, u) \neq (\bar{z}, \bar{u})$, which are necessary conditions for the following theorem resulting from (Grüne and Pannek, 2011, Theorem 6.18 and Section 6.6) to hold.

Theorem 2. Consider the MPC scheme with stage costs (6) satisfying the exponential controllability property from Definition 1 with $C = 1$ and $\rho \in (0, 1)$. Then the equilibrium $\bar{z}$ 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 $C = 1$.