21st International Symposium on
Mathematical Theory of Networks and Systems
July 7-11, 2014. Groningen, The Netherlands

Some results on Model Predictive Control for the Fokker-Planck equation*

Extended Abstract

Arthur Fleig[1], Lars Grüne[1] and Roberto Guglielmi[1,2]

***Abstract*— A Model Predictive Control scheme is applied to track the solution of a Fokker-Planck equation over a fixed time horizon. We analyse the dependence of the total cost functional on several parameters of the algorithm, in particular on the prediction horizon, on the regularization parameter, and on the sampling time. Comparison among different numerical simulations show valuable improvements by properly tuning the scheme's parameters. Our numerical study is complemented by a theoretical controllability analysis explaining the superior performance of controls with time and space dependence.**

I. Introduction

A. Model Predictive Control

Model Predictive Control (MPC), also known as Receding Horizon Control (RHC), is a control method that computes a feedback law by iteratively solving optimal control problems on finite time horizon, instead of coping directly with an optimal control problem on an infinite time interval (see the monographs [17] and [12] for an introduction to MPC). MPC can be briefly described as follows: consider a discrete time control system of the form

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

with $k \in \mathbb{N}_0$, state $z(k) \in \mathbb{X}$ and control $u(k) \in \mathbb{U}$ for suitable state and control constraint sets $\mathbb{X} \subset Z$ and $\mathbb{U} \subset U$, where the state space Z and the control space U are metric spaces. The MPC scheme constructs a feedback law $\mu : \mathbb{X} \to \mathbb{U}$ for the closed loop system

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

through the following steps:

0. Given an initial value $z_\mu(0) \in \mathbb{X}$, fix the length of the receding horizon N and set $n = 0$.
1. Initialize the state $z_0 = z_\mu(n)$ and minimize the functional
$$J_N(z_0, u) := \sum_{k=0}^{N-1} l(k, z(k), u(k)) \tag{3}$$
subject to (1). Let $u^* \in \mathbb{U}^N$ be the resulting optimal control and 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.

B. MPC for Partial Differential Equations

The application of MPC to infinite dimensional systems governed by Partial Differential Equations (PDEs) goes back to the work [15], where terminal constraints and control Lyapunov functionals were added as terminal costs to garantee the stability of the closed loop solution to the finite horizon problem. Furthermore, MPC schemes were applied to parabolic PDEs with either distributed or boundary control in [7] and [8]. However, the construction of suitable terminal regions and costs is in general a challenging task; for this reason, in most industrial applications an MPC scheme without terminal constraints is adopted [18]. Indeed, under suitable conditions the optimization objective will force the optimal trajectories to end up in a suitable terminal region (see, for example, [11]). Several theoretical results on the MPC setting without stabilizing terminal constraints have been collected in [12]. In the case of a linear parabolic equation with either distributed or boundary control, a rigorous analysis of the dependence of the receding horizon on the cost functional and the system parameters has been developed in [1]. Furthermore, a comparison between the qualitative behaviour of an MPC scheme for a heat equation for different types of boundary control has been carried out in [2].

In a PDE context, the solution of the equation is, of course, defined in continuous time. In a continuous time setting, the discrete times k from Section I-A indicate the re-optimization times and the map g from (1) can be obtained by sampling the PDE model in time. The state $z(k)$ in the discrete time model then represents the infinite dimensional state of the PDE model at time $t = t_k$; for details see, e.g., [2, Section 3].

II. Optimal control and the Fokker-Planck equation

Our interest for studying the Fokker-Planck equation mainly stems from its connection with the optimal control of the Probability Density Function (PDF) associated to stochastic processes. For explaning this connection, let us consider the continuous time stochastic process described by the (Itô) stochastic differential equation

$$dX_t = b(X_t, t; u)dt + \sigma(X_t, t)dW_t\,, \tag{4}$$

*This work was partially supported by the EU under the 7th Framework Program, Marie Curie Initial Training Network FP7-PEOPLE-2010-ITN SADCO, GA number 264735-SADCO.

[1]Chair of Applied Mathematics, Department of Mathematics, University of Bayreuth, 95440 Bayreuth, Germany,
`arthur.fleig, lars.gruene,`
`& roberto.guglielmi@uni-bayreuth.de`

[2]Roberto Guglielmi is currently a Marie Curie fellow within SADCO

where $t \in [0, T_E]$ for a fixed terminal time $T_E > 0$ and the state variable $X_t \in \mathbb{R}$ is subject to deterministic infinitesimal increments of the drift term b and to random infinitesimal increments dW_t of a Wiener process. In the setting considered here the control function acts through the term b and may either depend on space and time or may be merely time dependent as in [3], [4].

In deterministic dynamics, the optimal control is achieved by finding the control law u that minimizes a given objective given by a cost functional $J(X,u)$.

In the non-deterministic case of (4), the state evolution X_t represents a random variable. Therefore, when dealing with stochastic optimal control, usually the average of the cost function is considered [9]. In particular, the cost functional usually is of the form

$$J(X,u) = \mathbb{E}\left[\int_0^{T_E} L(t, X_t, u(t))dt + \psi(X_{T_E})\right],$$

for suitable running cost L and terminal cost ψ.

On the other hand, the state of a stochastic process can be characterized by the shape of its statistical distribution which is represented by the Probability Density Function (PDF). Therefore, a control methodology defined via the PDF would provide an accurate and flexible control strategy that could accommodate a wide class of objectives, cf. also [6, Section 4]. For this reason, in [10], [13], [14], [20] probability density function control schemes were proposed, where the cost functional depends on the PDF of the stochastic state variable. In this way, a deterministic objective results and no average is needed.

As shown in [3], [4], the PDF associated to the stochastic process (4) satisfies a Fokker-Planck equation [19] with a control acting through the divergence term. This is a partial differential equation of parabolic type with Cauchy data given by the initial PDF distribution. It can be expressed as follows

$$\begin{aligned} &\partial_t f(x,t) - \tfrac{1}{2}\partial_x^2\left(\sigma(x,t)^2 f(x,t)\right) + \partial_x\left(b(x,t;u)f(x,t)\right) = 0, \\ &f(x,t_k) = \rho_k(x), \end{aligned} \tag{5}$$

on the domain $Q_k := \Omega \times (t_k, t_k + T)$, where $\Omega \subset \mathbb{R}^d, d \in \mathbb{N}$, $t_k := kT$ with a sampling time $T > 0$, and some given initial distribution ρ_k. Notice that in general, the space domain in (5) is $\mathbb{R}^d$ instead of Ω. However, if localized SDEs are under consideration, or if the objective is to keep the PDF within a given compact set of Ω and the probability to find X_t outside of Ω is negligible, then it is reasonable to consider bounded $\Omega \subset \mathbb{R}^d$ with homogeneous Dirichlet boundary conditions. Given a desired distribution $f_d : \Omega \times \mathbb{R} \to \mathbb{R}$, a minimization problem of type (3) subject to (5) that attempts to keep f as close as possible to f_d using a control that is only time dependent can be posed as

$$\min_u J(f,u) := \frac{1}{2}\sum_{k=0}^{N-1}\left(\|f - f_d\|^2_{L^2(Q_k)} + \lambda|u(t_k)|^2\right) \tag{6}$$

for some positive constant λ, also called Tychonov regularization. For a control depending on time and space, the objective becomes

$$\min_u J(f,u) := \frac{1}{2}\sum_{k=0}^{N-1}\left(\alpha\|f - f_d\|^2_{L^2(Q_k)} + \lambda\|u\|^2_{L^2(Q_k)}\right), \tag{7}$$

where the weight $\alpha > 0$ is introduced solely for numerical reasons.

Setting $z(k) = f(\cdot, t_k)$, (6) and (7) can be rewritten as (3) with

$$l(k,z,u) = \frac{1}{2}\|f_z - f_d\|^2_{L^2(Q_k)} + \frac{\lambda}{2}|u(t_k)|^2$$

and

$$l(k,z,u) = \frac{\alpha}{2}\|f_z - f_d\|^2_{L^2(Q_k)} + \frac{\lambda}{2}\|u\|^2_{L^2(Q_k)},$$

respectively. Here, f_z denotes the solution of (5) with initial time t_k and initial distribution $\rho_k = z$.

III. MPC FOR THE FOKKER-PLANCK EQUATION

In this talk we are going to present numerical results as well as first steps towards a theoretical analysis of MPC applied to the Fokker-Planck equation. Probably the first papers which applied MPC to the Fokker-Planck equation were [3] in the one dimensional case and [4] in higher dimensions, both with the purpose to track a (smooth) target trajectory f_d. The particular type of MPC scheme in these references uses the horizon $N = 2$ and the functional (3) with

$$J(z,u) = \frac{1}{2}\|z(1) - f_d(\cdot, t_1)\|^2_{L^2(\Omega)} + \frac{\lambda}{2}|u|^2.$$

The numerical results we are going to present in this talk extend these results in various ways: on the one hand, we use the cost functionals (6) and (7) for implementing longer prediction horizons $N > 2$ in MPC. Particularly, we investigate the interplay between N and the sampling time T and its impact on the quality of the solution.

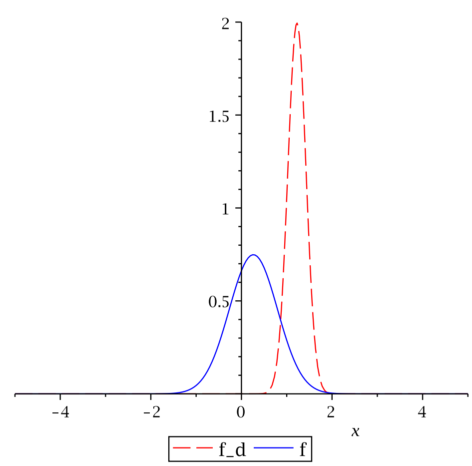

Fig. 1: Ornstein-Uhlenbeck (state)

Sampling time $T = 0.05$, $N = 2$, $u = u(t)$, $t = 1.2$

As an example, consider the two different MPC simulations for $T = 0.05$ on $\Omega =]-5, 5[$ depicted in Figures 1 to 5,

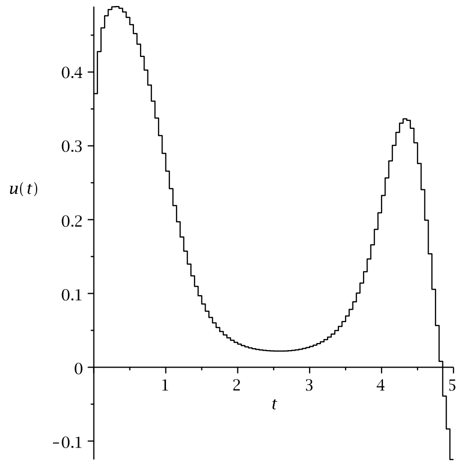

Fig. 2: Ornstein-Uhlenbeck (control)

Sampling time $T = 0.05$, $N = 2$, $u = u(t)$

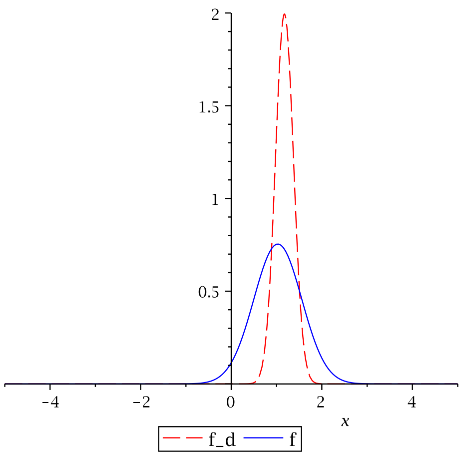

Fig. 3: Ornstein-Uhlenbeck (state)

Sampling time $T = 0.05$, $N = 11$, $u = u(t)$, $t = 1.2$

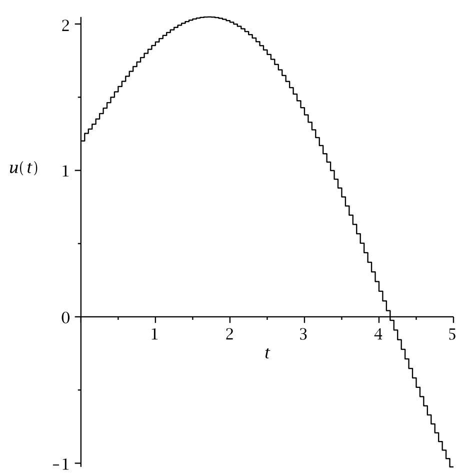

Fig. 4: Ornstein-Uhlenbeck (control)

Sampling time $T = 0.05$, $N = 11$, $u = u(t)$

for the tracking of the PDF of the one dimensional Ornstein-Uhlenbeck process with $b(x,t,u) = u - x$ and $\sigma(x,t) = 0.8$ in equation (4). The initial and target PDF used are

$$\rho_0(x) = \frac{1}{\sqrt{2\pi \cdot 0.1^2}} \exp\left(-\frac{x^2}{2 \cdot 0.1^2}\right)$$

and

$$f_d(x,t) = \frac{\exp\left(-\frac{[x-2\sin(\pi t/5)]^2}{2 \cdot 0.2^2}\right)}{\sqrt{2\pi \cdot 0.2^2}},$$

respectively. Furthermore, the parameters in the objective functionals (6) and (7) are given by $\lambda = 0.1$ and $\alpha = 100$. Note that Ω is chosen large enough such that the error made from disregarding $\mathbb{R} \setminus \Omega$ is negligible. Obviously, for this example the controlled PDF (in solid blue) tracks the desired PDF (in dashed red) much better for larger prediction horizon N.

On the other hand, we investigate the improvements which can be achieved when the control u is chosen time and space dependent. From a control point of view, this corresponds to a control structure which has both state dependent (i.e., feedback) character but may also vary with time. However, since MPC yields a (discrete time) feedback law [12, Section 3], the time dependence of u is actually induced via the dependence on the evolution of the PDF f, i.e., via a dependence on the state of the Fokker-Planck equation (5). The additional dependence on the state x of (4) allows for a significant increase of the quality of the tracking of the MPC feedback, which becomes visible by comparing Figure 5 with state and time dependent $u = u(t,x)$ with the merely time dependent $u = u(t)$ in Figure 3.

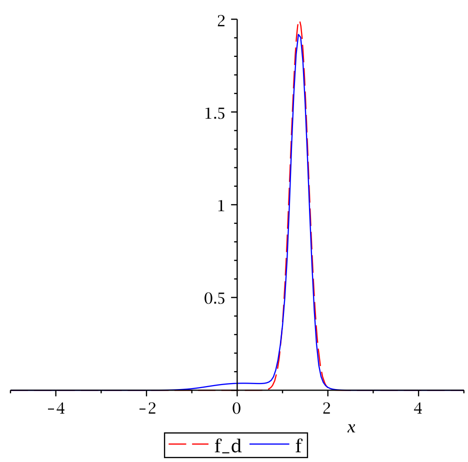

Fig. 5: Ornstein-Uhlenbeck

Sampling time $T = 0.01$, $N = 16$, $u = u(x,t)$, $t = 1.2, \alpha = 100$

The simulations shown in the figures have been obtained by using space discretization as in [3] and explicit Euler discretization in time. The optimization problem has been solved with the projected gradient method and the Newton method with BFGS Hessian approximation.

On the theoretical side, in the talk we explain the role of controllability properties of the Fokker-Planck equation for obtaining good tracking results. We explain how the theoretical controllability results for the Fokker-Planck equation [5], [16] imply that a control dependent both on time and space allows for the highly precise tracking via MPC feedback laws illustrated in Figure 5.

REFERENCES

[1] N. Altmüller and L. Grüne, Distributed and boundary model predictive control for the heat equation, GAMM-Mitt. 35, 2:131–145, 2012.
[2] N. Altmüller and L. Grüne, A comparative stabiliy analysis of Neumann and Dirichlet boundary MPC for the heat equation, Proceedings of the 1st IFAC Workshop on Control of Systems Modeled by Partial Differential Equations — CPDE 2013, 161–166, 2013.
[3] M. Annunziato and A. Borzì, Optimal control of probability density functions of stochastic processes, Math. Model. Anal. 15:393–407, 2010.
[4] M. Annunziato and A. Borzì, A Fokker-Planck-based control framework for multidimensional stochastic processes, J. Comput. Appl. Math. 237:487–507, 2013.
[5] A. Blaquiere, Controllability of a Fokker-Planck equation, the Schrödinger system, and a related stochastic optimal control. Dynam. Control 2, 3:235–253, 1992.
[6] R. Brockett, New issues in the mathematics of control, B. Engquist and W. Schmid, eds., Mathematics unlimited—2001 and beyond, Springer, Berlin, 189–219 2001.
[7] S. Dubljevic and P.D. Christofides, Boundary Predictive Control of Parabolic PDEs, Proceedings of the 2006 American Control Conference, Minneapolis, Minnesota, USA, 49–56, 2006.
[8] S. Dubljevic, N.H. El-Farra, P. Mhaskar and P.D. Christofides, Predictive control of parabolic PDEs with state and control constraints, Intern. J. Robust Nonlin. Control, 16:749–772, 2006.
[9] W.H. Fleming and R.W. Rishel, Deterministic and Stochastic Optimal Control. Springer, Berlin, 1975.
[10] M.G. Forbes, M. Guay and J.F. Forbes, Control design for first-order processes: shaping the probability density of the process state. Journal of process control, 14:399–410, 2004.
[11] L. Grüne, NMPC without terminal constraints, Proceedings of the IFAC Conference on Nonlinear Model Predictive Control 2012 — NMPC'12, 1–13, 2012.
[12] L. Grüne and J. Pannek, Nonlinear Model Predictive Control. Theory and Algorithms, Springer, London, 2011.
[13] G. Jumarie. Tracking control of nonlinear stochastic systems by using path cross-entropy and Fokker-Planck equation. Internat. J. Systems Sci., 23(7):1101–1114, 1992.
[14] M. Karny. Towards Fully Probabilistic Control Design. Automatica, 32(12):1719–1722, 1996.
[15] K. Ito and K. Kunisch, Receding horizon optimal control for infinite dimensional systems, ESAIM: COCV, 8:741–760, 2002.
[16] A. Porretta, On the planning problem for a class of mean field games. C. R. Math. Acad. Sci. Paris 351, no. 11–12:457–462, 2013.
[17] J.B. Rawlings and D.Q. Mayne, Model Predictive Control: Theory and Design, Nob Hill Publishing, 2009.
[18] S.J. Qin and T.A. Badgwell, A survey of industrial model predictive control technology, Control Engineering Practice, 11(7):733–764, 2003.
[19] R. Risken. The Fokker-Planck Equation: Methods of Solution and Applications. Springer, Berlin, 1996.
[20] H. Wang. Robust Control of the Output Probability Density Functions for Multivariable Stochastic Systems with Guaranteed Stability. IEEE Trans. Autom. Control, 44(11):2103–2107, 1999.