Model Predictive Control for the Fokker-Planck equation: analysis and structural insight

Extended Abstract

Arthur Fleig[1] and Lars Grüne[1]

Abstract— **For the control of ensembles governed by controlled stochastic differential equations we follow the approach to control the corresponding probability density function. To this end, we propose to use Model Predictive Control (MPC) for the Fokker-Planck equation. In this talk we start by describing the basic setup and illustrating the approach by numerical examples. Then, we provide first results on the analysis of the stability and performance of the MPC approach. Finally, we discuss the structure of the controller resulting from the MPC approach, particularly its dependence on space, time and on the probability density function of the ensemble under consideration.**

I. Introduction

In this talk we consider a Model Predictive Control (MPC) approach to the control of an ensemble, with the dynamics of each element of the ensemble governed by the controlled Itô stochastic differential equation (SDE)

$$dX_t = b(X_t,t,u)dt + \sigma(X_t,t)dW_t \tag{1}$$

with initial condition $X_0 \in \mathbb{R}^d$. The distribution of a large ensemble is statistically determined by its time dependent probability density function (PDF) $y : \mathbb{R}\times\mathbb{R}^d \to \mathbb{R}_0^+$. The control task thus consists of controlling the PDF of the ensemble towards a desired reference density function $y_{ref} : \mathbb{R}^d \to \mathbb{R}_0^+$. Under suitable regularity conditions, the PDF is determined by the Fokker-Planck partial differential equation

$$\partial_t y(t,x) = \sum_{i,j=1}^{d} \partial_{ij}^2 \left(a_{ij}(t,x)y(t,x)\right) + \sum_{i=1}^{d} \partial_i \left(b_i(t,x,u)y(t,x)\right)$$
$$y(0,x) = y_0(x)$$

for $(t,x)\in\mathbb{R}_{>0}\times\mathbb{R}^d$ and with $a_{ij} = \sum_k \sigma_{ik}\sigma_{jk}/2$, for details see, e.g., [8, p. 227], [9, p. 297] or [10]. Here, u can be a function of time t and/or state x.

In order to apply MPC to the problem, it is convenient to rewrite the sampled-data version of the Fokker-Planck equation as a discrete time system. To this end we fix a sampling time $T_s > 0$, sampling instants $t_n := nT_s$ for $n \in \mathbb{N}_0$ and the discrete time state

$$z(n) := y(t_n,\cdot),$$

which is now an element of an appropriate function space $\mathbb{X}$. Denoting the piece of the control function u acting from t_n to t_{n+1} shifted to $[0,T_s]$ by $u(n)$ (i.e., $u(n)(t,x) = u(t+t_n,x)$) and denoting by f the solution operator of the Fokker-Planck equation on the interval $[0,T_s]$, we can then write the discrete time dynamics as

$$z(n+1) = f(z(n),u(n)), \quad z(0) = z_0 = y_0. \tag{2}$$

Note that $u(n)$ can be either time varying or constant in time on $[0,T_s]$; the latter setting leads to a sampled data system with zero order hold. Similarly, $u(n)$ can be varying or constant in the state variable x, depending on the considered application. We denote the space of admissible control inputs for f by $\mathbb{U}$.

MPC now consists of iteratively minimizing a finite horizon functional of the form

$$J_N(z_0,u) := \sum_{k=0}^{N-1} \ell(z_u(k;z_0),u(k)) \tag{3}$$

with respect to u, where $z_u(k;z_0)$ denotes the solution of (2) for discrete time control $u = u(\cdot) \in \mathbb{U}^N$. We assume that the desired reference PDF y_{ref} is an equilibrium, i.e., that there exists an admissible control $u_{ref} \in \mathbb{U}$ such that $f(y_{ref},u_{ref}) = y_{ref}$, and define the stage cost ℓ in (3) as

$$\ell(z,u) := \frac{1}{2}\|z-y_{ref}\|^2_{L^2(\mathbb{R}^d)} + \frac{\lambda}{2}\|u-u_{ref}\|^2_2 \tag{4}$$

for a parameter $\lambda > 0$.

A feedback law μ is then obtained by the usual moving horizon iteration:

1. Given an initial value $z_\mu(0) \in \mathbb{X}$, fix the length of the optimization horizon N and set $n = 0$.
2. Initialize the state $z_0 = z_\mu(n)$ and minimize (3) with respect to $u \in \mathbb{U}^N$. 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 $z_\mu(n+1) = f(z_\mu(n),\mu(z_\mu(n)))$, set $n := n+1$ and go to step 2.

Clearly, in order to apply MPC in a meaningful way, the well-posedness and solvability of the optimal control problem in Step 2 must be ensured. This will not be a focus of this talk, but we mention that in passing related results from [1], [2], [3], [5], this is the main reason for using an L^2-functional in (4) although in the literature other

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

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

types of distances like the Wasserstein metric are sometimes preferred, cf. [7].

The two results presented in detail in this talk are outlined in the following two sections.

II. Stability analysis for spatially constant control

The use of MPC for the Fokker-Planck control problem introduced above was first proposed in [2], [3]. In these references, the particular choice $u_{ref} = 0$ was made and the class of control functions was limited to functions being constant in space, i.e., each element of the ensemble applies the same control input. While in general the optimization horizon N needs to be sufficiently large for ensuring asymptotic stability, the numerical results in [2], [3] indicate that for the setting investigated in these references the MPC closed loop is in fact asymptotically stable even for the shortest meaningful horizon $N = 2$.

A formal analytic explanation why this is the case has recently been provided in [4] and will be explained in this section. The analysis relies on the following exponential controllability property.

Definition 1: The system (2) is called exponentially controllable with respect to the stage costs ℓ if there exist constants $C \geq 1$ and $\rho \in (0,1)$ such that for all $z_0 \in \mathbb{X}$ there exists an admissible control $u_{z_0} \in \mathbb{U}^\infty$ with

$$\ell(z_{u_{z_0}}(n;z_0), u_{z_0}(n)) \leq C\rho^n \min_{u \in U} \ell(z_0, u) \quad (5)$$

for all $n \in \mathbb{N}_0$, where $z_{u_{z_0}}(k;z_0)$ denotes the solution of (2) with $u = u_{z_0}$.

If this exponential controllability property holds for ℓ from (4), then the equilibrium y_{ref} is globally asymptotically stable for the MPC closed loop provided the optimization horizon N is sufficiently large [6, Theorem 6.18]. If, moreover, exponential controllability holds with $C = 1$, then this assertion even holds for $N = 2$ [6, Section 6.6]. For proving asymptotic stability with $N = 2$, it is thus sufficient to check Definition 1 with $C = 1$.

This can be accomplished in the case where the dynamics is governed by the d-dimensional Ornstein-Uhlenbeck process, which is obtained by choosing the diffusion as

$$a_{ij} := \delta_{ij}\sigma_i^2/2, \quad (6)$$

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

$$b_i(t,x,u) := -\mu_i x + u_i \quad (7)$$

for $\mu_i > 0$ and $u_i \in \mathbb{R}$.

Clearly, for controls constant in space the possibility to control the PDF is rather limited. Indeed, for zero order hold control, the only equilibria y_{ref} of the corresponding discrete time systems dynamics are normal distributions with variance σ independent of u_{ref} and mean determined by u_{ref}. For initial conditions y_0 that are normal distributions, exponential controllability w.r.t. ℓ indeed holds for $C = 1$. However, depending on the parameters of y_0, the verification of Definition 1 with $C = 1$ may not always be possible for ℓ from (4). In the talk we will explain how to circumvent this problem by constructing a cost function equivalent to (4), i.e., a cost function that yields identical optimal trajectories, and for which Definition 1 holds with $C = 1$.

III. Structural insight

In a more general setting than that of Section II, i.e., when u becomes state dependent or when other types of SDEs are considered, estimates on the minimal stabilizing optimization horizon are not yet available. However, whenever the exponential controllability condition from Definition 1 is satisfied, we know that y_{ref} will be asymptotically stable for the MPC closed loop for sufficiently large optimization horizon N, see [6, Theorem 6.18].

In this case, the MPC approach reveals interesting structural insight about the type of the control needed to achieve asymptotic stability of a desired PDF. Indeed, due to the space dependence of the control, the control action applied on each element of the ensemble depends on the state $x = X_t$ of the individual element. As such, from the point of view of the ensemble elements, the control takes the form of a time dependent (sampled data) feedback law. However, from the point of view of the Fokker-Planck equation, the time dependence of the control is entirely induced by the state of the Fokker-Planck equation, i.e., by the evolution of the PDF. Hence, the time dependence of the control is actually not exogenous, but triggered by a space dependence on a higher, “statistical” level.

This aspect will be illustrated in the talk by numerical simulations, which will also investigate the robustness of the approach against estimation errors for the PDF of the ensemble.

References

[1] A. Addou and A. Benbrik, “Existence and uniqueness of optimal control for a distributed-parameter bilinear system,” *J. Dynam. Control Systems*, vol. 8, no. 2, pp. 141–152, 2002.

[2] M. Annunziato and A. Borzì, “Optimal control of probability density functions of stochastic processes,” *Math. Model. Anal.*, vol. 15, no. 4, pp. 393–407, 2010.

[3] ——, “A Fokker-Planck control framework for multidimensional stochastic processes,” *J. Comput. Appl. Math.*, vol. 237, no. 1, pp. 487–507, 2013.

[4] A. Fleig and L. Grüne, “Estimates on the minimal stabilizing horizon length in Model Predictive Control for the Fokker-Planck equation,” *submitted*, 2016.

[5] A. Fleig and R. Guglielmi, “Bilinear optimal control of the Fokker-Planck equation,” *submitted*, 2016.

[6] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control. Theory and Algorithms*. London: Springer-Verlag, 2011.

[7] R. Jordan, D. Kinderlehrer, and F. Otto, “The variational formulation of the Fokker-Planck equation,” *SIAM Journal on Mathematical Analysis*, vol. 29, no. 1, pp. 1–17, 1998.

[8] S. Primak, V. Kontorovich, and V. Lyandres, *Stochastic methods and their applications to communications*. John Wiley & Sons, Inc., Hoboken, NJ, 2004.

[9] P. E. Protter, *Stochastic integration and differential equations*, ser. Stochastic Modelling and Applied Probability. Springer-Verlag, Berlin, 2005, vol. 21.

[10] H. Risken, *The Fokker-Planck equation*, 2nd ed., ser. Springer Series in Synergetics. Springer-Verlag, Berlin, 1989, vol. 18.