

Towards a solution of mean-field control problems using model predictive control^{*}

Alfio Borzi^{*} Lars Grüne^{**}

^{*} *Chair of Scientific Computing, Institute for Mathematics
University of Würzburg, Emil-Fischer-Straße 30, 97074 Würzburg,
Germany (e-mail: alfio.borzi@mathematik.uni-wuerzburg.de)*

^{**} *Chair of Applied Mathematics, Mathematical Institute
University of Bayreuth, 95447 Bayreuth, Germany
(e-mail: lars.gruene@uni-bayreuth.de)*

Abstract: In this paper we propose a model predictive control (MPC) scheme for solving mean-field control problems. To this end, the MPC scheme is applied to a controlled Fokker-Planck equation. We test our algorithm by means of a numerical example, both with and without nonlinear coupling. We also provide numerical evidence that MPC indeed delivers approximately optimal trajectories for this example.

Keywords: model predictive control, mean-field control problem, controlled Fokker-Planck equation

INTRODUCTION

Mean-Field optimal control problems appear in stochastic control systems, in which the overall behaviour of a population of agents affects the dynamics of the single agent as well as the optimization objective, see Bensoussan et al. (2013). They have various applications in control engineering, e.g., for consensus problems, cf. Nourian et al. (2013), and for the management of large populations of flexible electric loads, see Grammatico et al. (2015).

It is well known that mean-field optimal control problems can be solved in the framework of Fokker-Planck equations, see Bensoussan et al. (2013) or Annunziato and Borzi (2018). One of the benefits of this approach is that the original stochastic optimal control problem is turned into a deterministic PDE optimization problem, for which powerful numerical solution techniques are available. However, when the time horizon of the problem is variably or infinitely long, as in typical regulation problems in control engineering, this approach is subject to severe numerical difficulties.

In this situation, Model Predictive Control (MPC) has turned out to be a valid alternative. In MPC, an optimal control problem on a long, possibly infinite time horizon is split up into the consecutive solution of problems on relatively short finite time horizons. Due to the repeated optimization, MPC yields a feedback control, which is why we refer to the solution trajectories generated by MPC as the MPC closed-loop solutions.

There is by now an established theory that explains when MPC gives approximately optimal closed-loop solution, see, e.g., Grüne (2016). The main structural property that is needed for this approximation is the so called

turnpike property, which in turn is closely connected to a dissipativity property, see Grüne and Müller (2016).

In this paper, we propose to use MPC for obtaining approximately optimal solutions for mean-field optimal control problems. After describing the MPC scheme as well as some of the necessary theoretical background, we provide a numerical study by which we explore the capability of MPC for solving mean-field control problems. Besides simulation results, we provide numerical evidence that the MPC solution indeed gives approximately optimal solutions. This is done, on the one hand, by numerically exploring the objective value and, on the other hand, by numerically verifying the occurrence of the turnpike property.

1. MEAN-FIELD CONTROL PROBLEMS

In this section we describe the class of stochastic optimal control problems that we want to address. In order to keep the presentation technically simple, we leave out several technical details. Interested readers may consult the monograph by Bensoussan et al. (2013) for the necessary mathematical background. We also refer to the recent survey by Annunziato and Borzi (2018), which discusses mean-field control problems in the context of control of the Fokker-Planck equation.

In mean-field problems, one considers an infinite number of agents with identical dynamics, which is governed by the n -dimensional Itô-stochastic differential equation (SDE)

$$dx(t) = g(x(t), m(t), v(t, x(t)))dt + \sigma(x(t))dW(t) \quad (1)$$

with initial condition $x(0) = x_0$. Here $x(t)$ is the state of the agent at time t and $v(t, \cdot)$ is the feedback control, that is assumed identical for all agents. The vector valued

^{*} This research was supported by DFG-Grant GR 1569/13-1.

function g describes the drift and the matrix valued function σ the diffusion of the SDE, and $W(t)$ is an n -dimensional standard Wiener process. The set of all agents is assumed to have the initial distribution m_0 and the quantity $m(t)$ with $m(0) = m_0$ describes the evolution of this distribution over time. Note that the evolution of $m(t)$ depends on the choice of the control $v(t, \cdot)$.

Given an initial distribution m_0 , the goal of the mean-field control problem is to find a control v and an evolution of the associated distribution of the agents m such that the functional

$$J(v, m) = \mathbb{E} \left[\int_0^T h(x(t), m(t), v(t, x(t))) dt \right]$$

is minimized for a given cost function h , where \mathbb{E} denotes the expectation operator. Observe that through this optimization criterion the optimal control $v^*(t, \cdot)$ depends on $m(t)$. Hence, these two functions are coupled and this coupling may induce nonlinearities into the problem even if g and σ in (1) are linear functions.

We note that this stochastic functional involving the expectation operator can be rewritten in a purely deterministic form as

$$J(v, m) = \int_0^T \int_{\Omega} h(x, m(t), v(t, x)) m(t, x) dx dt, \quad (2)$$

where $m(t, x)$ denotes the evaluation of the distribution function $m(t)$ in $x \in \Omega$.

Given the control v , the evolution of the probability distribution m is characterized by the Fokker-Planck equation

$$\begin{aligned} \partial_t m(t, x) - \frac{1}{2} \sum_{i,j=1}^n \partial_{x_i x_j}^2 (a_{ij}(x) m(t, x)) \\ + \sum_{i=1}^n \partial_{x_i} (g_i(x, m(t, x), v(t, x)) m(t, x)) = 0, \end{aligned} \quad (3)$$

with initial condition $m(0, x) = m_0(x)$ and suitable boundary conditions,

$$a_{ij} = \sum_{k=1}^m \sigma_{ik} \sigma_{jk},$$

and g and σ from (1). Using the Fokker-Planck equation as dynamics and (2) as cost functional, we have rewritten the mean-field control problem as a purely deterministic problem.

The mean-field formulation describes the probability density of N interacting agents when N tends to infinity (Annunziato and Borzì, 2018, Section 3). It can thus be seen as a model reduction technique for a problem which otherwise becomes numerically intractable already for moderate size of N . The interaction between the agents induces the m -dependence of g in (1), which makes the solution of the problem analytically and numerically challenging, because the Fokker-Planck equation (3) is nonlinear. However, this dependence is important in order to obtain a realistic

modelling in many applications. We will illustrate this in our numerical example at the end of this paper.

We note that mean-field type control problems are different from mean-field games. While in mean-field games a Nash equilibrium for the objectives of the agents is studied, in a mean-field control problem the overall optimum (i.e., the optimum over v and m) is sought. Hence, mean-field games try to explain the interaction between the members of a large population of non-cooperative agents, while mean-field control problems aim at finding the best possible solution depending on the current distribution of the agents. This enables us to find the best way of action for a large population of agents (which are approximated by the infinite population in the mean-field formulation) and is fundamental for problems, e.g., in crowd dynamics or traffic control, see Roy et al. (2016); Herty and Pareschi (2010). Our proposed method thus gives new algorithmic approaches for problems in these application areas.

The usual way to solve the mean-field control problem is to couple the Fokker-Planck equation (3) with a Hamilton-Jacobi-Bellman equation for computing the optimal feedback control v . The solution of this McKean-Vlasov type coupled system of nonlinear PDEs (for details see, e.g., (Bensoussan et al., 2013, Eq. (4.12)), Annunziato et al. (2014) or (Annunziato and Borzì, 2018, Section 4)) then characterizes the optimal solutions and can be used in order to compute the optimal feedback control. However, the solution becomes numerically very challenging in case of very long or even infinite time horizons T , which we want to consider in this paper.

2. MODEL PREDICTIVE CONTROL

A remedy for these difficulties is the use of Model Predictive Control (also called receding horizon control). In this approach, the problem on a long or infinite horizon is split up into the consecutive solution of problems on shorter time horizons, which are thus much easier to solve. In practical applications, MPC is often used as an online optimization method for computing a feedback law and the analysis of MPC in this context is typically focused on stability and feasibility questions. Another, equally useful aspect of MPC is that it provides approximately optimal solutions to the original problem. It can thus be regarded as a model reduction technique in time for solving optimal control problems on long or infinite horizons. We mention again that we have rewritten the stochastic mean-field problem as a purely deterministic problem. This means that we can use standard deterministic MPC schemes for its solution, which is advantageous in two aspects: on the one hand, deterministic schemes are much easier to implement as there is no need to resort to stochastic optimization. On the other hand, while some results on approximate optimality of MPC are available in the stochastic setting, e.g., in Chatterjee and Lygeros (2015), there are a much more rich and general results guaranteeing approximate optimality for deterministic MPC schemes.

MPC can be formulated in continuous time or in discrete time. Since more general and powerful analysis results for MPC are available in discrete time, we will present it in this form. To this end, we sample the solution of the Fokker-Planck equation in time by introducing a discrete

time scale $t_k := kh$, $h > 0$, $k \in \mathbb{N}$, and defining $z(k) = m(t_k, \cdot)$ as the state of a discrete time model

$$z(k+1) = f(z(k), u(k)) \quad (4)$$

with initial condition $z(0) = z_0$. Equation (4) constitutes an infinite-dimensional discrete-time control system on a suitable Banach space Z with f being the solution operator of the controlled Fokker-Planck equation. The discrete time control $u(k)$ is given by the continuous time control $v(\cdot, t)|_{[t_k, t_{k+1}]}$.

The functional $J(v, m)$ can then be rewritten as

$$J_N(z(0), u) = \sum_{k=0}^{N-1} \ell(z(k), u(k)), \quad (5)$$

where $Nh = T$ and

$$\ell(z(k), u(k)) = \int_{t_k}^{t_{k+1}} \int_{\Omega} f(x, m(t), v(t, x)) m(t, x) dx dt,$$

In this way we obtain an exact representation of the continuous time PDE optimization problem. Note that $m|_{[t_k, t_{k+1}]}$ is determined by $m(t_k) = z(k)$ and $v(\cdot, t)|_{[t_k, t_{k+1}]} = u(k)$, thus ℓ is well defined as a function of $z(k)$ and $u(k)$.

In MPC, (5) is minimized with optimization horizon $N \in \mathbb{N}$ satisfying $Nh \ll T$. Moreover, typically state constraints $z \in \mathbb{Z}$ and control constraints $u \in \mathbb{U}$ are imposed. In detail, MPC computes a feedback controller $\mu : \mathbb{Z} \rightarrow \mathbb{U}$ for the closed loop system

$$z_{\mu}(n+1) = f(z_{\mu}(n), \mu(z_{\mu}(n))) \quad (6)$$

in the following way.

0. Given an initial value $z_{\mu}(0) \in \mathbb{Z}$, fix the optimization horizon $N \in \mathbb{N}$ and set $n := 0$.
1. Measure the current state $z_{\mu}(n)$ and minimize (5) with respect to $u(\cdot) \in \mathbb{U}^N$ subject to the constraints $z(0) = z_{\mu}(n)$, $z(k) \in \mathbb{Z}$ for all $k = 1, \dots, N$ and (4). Denote the resulting optimal control sequence by $u^* \in \mathbb{U}^N$ and set $\mu(z_{\mu}(n)) := u^*(0)$.
2. Compute $z_{\mu}(n+1)$ according to (6) and set $n := n+1$. If $(n+N)h < T$, then go to 1, else stop.

For general optimal control problems, it is not clear whether this algorithm gives an approximately optimal trajectory. In what follows, we explain an approach for guaranteeing this property for the infinite horizon case. Clearly, some redundancy is needed in the optimal solutions on large or infinite time horizons, which allows to solve problems on shorter and finite horizons and still obtain approximate optimality. It turns out that the so-called turnpike property provides this kind of redundancy.

Definition 1. Let $(z^e, u^e) \in \mathbb{Z} \times \mathbb{U}$ be an equilibrium of (4), i.e., $f(z^e, u^e) = z^e$.

- (i) The infinite horizon optimal control problem for stage cost $\ell(z, u) - \ell(z^e, u^e)$ has the turnpike property at x^e if the following holds: there exists $\rho \in \mathcal{L}$ such that for each optimal trajectory z^* and all $P \in \mathbb{N}$ there is a

set $\mathcal{Q}(z^*(0), P, \infty) \subseteq \mathbb{N}_0$ with $\#\mathcal{Q}(z^*(0), P, \infty) \leq P$ and

$$\|z^*(k) - z^e\| \leq \rho(P)$$

for all $k \in \mathbb{N}_0$ with $k \notin \mathcal{Q}(z^*(0), P, \infty)$.

- (ii) The finite horizon optimal control problems have the turnpike property at z^e if the following holds: there exists $\sigma \in \mathcal{L}$ such that for each optimal trajectory z^* , $x \in \mathbb{Z}$ and all $N, P \in \mathbb{N}$ there is a set $\mathcal{Q}(z^*(0), P, N) \subseteq \{0, \dots, N\}$ with $\#\mathcal{Q}(z^*(0), P, N) \leq P$ and

$$\|z^*(k) - z^e\| \leq \sigma(P)$$

for all $k \in \{0, \dots, N\}$ with $k \notin \mathcal{Q}(z^*(0), P, N)$.

Summarizing, the turnpike property says that any optimal solution “most of the time” stays close to the equilibrium z^e . This has two consequences: first, the initial pieces for solutions for different optimization horizons are similar, because they all approach the optimal equilibrium x^e in an optimal way. Second, after a suitable period of time the effect of the initial condition becomes negligible, i.e., all optimal solutions z^* look roughly the same independent of $z^*(0)$. These properties are both easily recognized in numerical simulations of optimal trajectory, which is why it is easy to find numerical evidence for the turnpike property. This is also the way we pursue in Section 3 of this paper. A rigorous verification of the turnpike property is most easily done via dissipativity arguments. For a class of (linear) Fokker-Planck equations, such an analysis was performed in Fleig and Grüne (2019).

Many optimal control problems exhibit the turnpike property and we refer to Dorfman et al. (1958); McKenzie (1986) for classical and to Faulwasser et al. (2017); Grüne and Müller (2016); Trélat et al. (2018) for recent results in this field. If the turnpike property is satisfied, we can make an approximate optimality statement. Let

$$J_M^{cl}(z(0), \mu) = \sum_{k=0}^{M-1} \ell(z_{\mu}(k), \mu(z_{\mu}(k))) \quad (7)$$

denote the cost of the trajectory generated by MPC and

$$V_{\infty}(z(0)) := \inf_{u \in \mathbb{U}^{\infty}, z(k) \in \mathbb{Z}} \sum_{k=0}^{\infty} \ell(z(k), u(k)) - \ell(z^e, u^e)$$

denote the infinite horizon optimal value function of the problem with stage cost $\ell - \ell(x^e, u^e)$, i.e., the best value that can be achieved on the infinite horizon. We remark that $V_{\infty}(z(0))$ need not be finite, however, in many optimal control problems it is, at least for initial values $z(0)$ near z^e . In this case, the following theorem holds. Its proof is based on the similarity of the initial pieces of the optimal trajectories and can be found in Grüne (2016), Theorem 4.4.

Theorem 2. If the optimal control problem has the turnpike property for finite and infinite horizon and the optimal value functions satisfy suitable continuity and boundedness conditions, then there is a function $\delta \in \mathcal{L}$ such that the inequality

$$J_M^{cl}(z(0), \mu) + V_{\infty}(z_{\mu}(M)) \leq V_{\infty}(z(0)) + M\delta(N) \quad (8)$$

holds for all $M \in \mathbb{N}$ and all sufficiently large $N \in \mathbb{N}$.

The interpretation of (8) is as follows: the left-hand side can be seen as the value obtained if we follow the MPC trajectory for M steps and then continue in an infinite

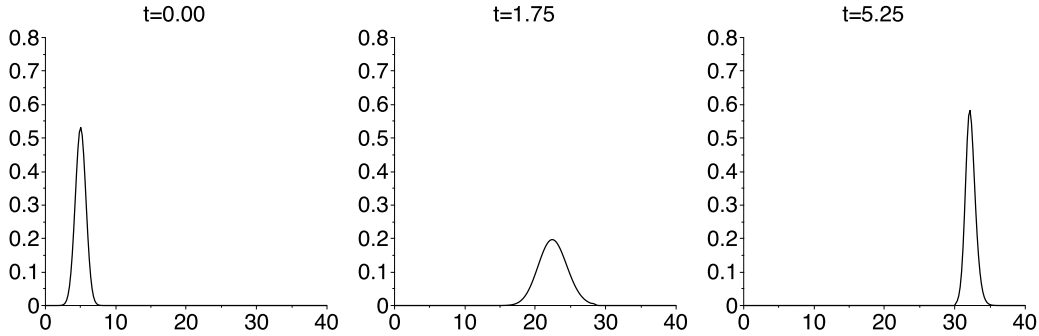


Fig. 1. Transport of the initial density into $[30, 40]$ using MPC without density dependent speed constraint

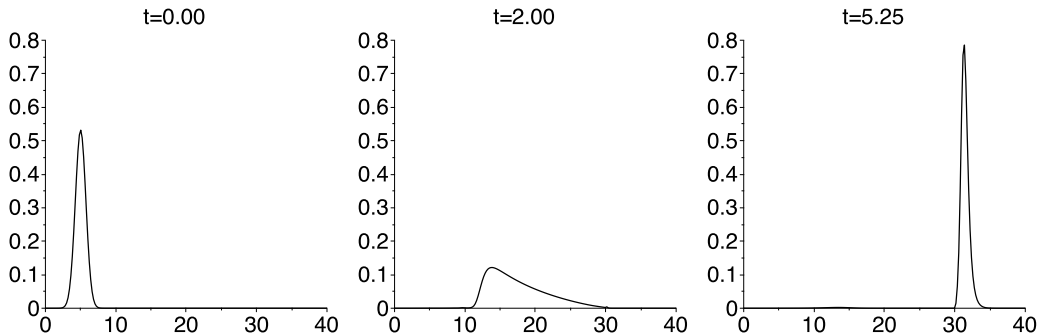


Fig. 2. Transport of the initial density into $[30, 40]$ using MPC with density dependent speed constraint

horizon optimal fashion. The inequality hence says that the cost of this trajectory is at most by $M\delta(N)$ larger than the value of the exact optimal trajectory. At the first glance, the fact that the error term $M\delta(N)$ grows with M may seem to indicate that the estimate deteriorates as $M \rightarrow \infty$. However, for typical problems, the modulus of the optimal finite horizon cost also grows with the length of the horizon M . Thus, if we consider the error relative to the optimal finite horizon cost, then it is constant in M .

3. NUMERICAL EXAMPLE

With the help of a numerical example we will explore whether the proposed approach to use MPC for mean-field control problems is valid. To this end, we study a prototypic problem, in which the task is to steer a distribution of agents following the 1d SDE

$$dx(t) = v(t)dt + 0.75dW(t)$$

from the interval $[0, 30]$ into the interval $[30, 40]$, which is achieved by using the stage cost $f(x, m, v) = m(x) + \lambda v^2$ for $x \in [0, 30]$ and $f(x, m, v) = \lambda v^2$ for $x \in [30, 40]$ with a very small $\lambda = 10^{-4}$. The control input is limited to $v \in [-2.1, 10]$.

For the numerical simulations, we first discretized the Fokker-Planck equation using the established combination of the conservative and positivity-preserving Chang-Cooper scheme in space (Chang and Cooper (1970)) and the BDF2 scheme in time, see Mohammadi and Borzi (2015). Our spatial domain was the interval $[0, 40]$, which we approximated using 201 equidistant grid points. The temporal discretization step size in the numerical scheme was $\Delta t = 0.005$ while the sampling time in the MPC

scheme was set to $h = 0.25$. Unless otherwise indicated, the optimization horizon in MPC was chosen as $N = 15$ sampling intervals. For the numerical optimization we used the Projected Gradient method to take into account control constraints. To compute the gradient, we formally derived the corresponding adjoint equation, which we discretized similarly to Annunziato and Borzi (2013).

If we think of the agents as cars on a road, then the speed of the overall population will depend on the density of the cars. Without taking this fact into account, the optimal control will only shift and rescale the initial distribution, resulting in an unrealistic solution that does not take into account the need for flattening the distribution in order to enable the cars to drive at a higher speed, see Figure 1.

In order to incorporate the limited speed at high density into the model, we scale the speed $v(t)$ in the equation with the factor $1/(1 + 5m(x))$, resulting in a reduced speed at those places x where the density $m(x)$ is high. We note that this scaling induces a nonlinearity in the Fokker-Planck equation, which also poses a challenge for the numerical solver. Here, however, the nonlinearity is of a particular form that can be implemented as a density dependent input constraint, which the discretization method described above can handle. This constraint leads to a significantly different and in particular flatter and skewed shape of the distribution during the transition, see Figure 2.

In order to verify the occurrence of turnpike behavior in our example, we perform two experiments. The turnpike property implies that the initial part of the open-loop optimal solution is similar for all horizons N that are sufficiently large. Figure (3) shows that this is what happens: for $N \geq 15$, the optimal solutions are virtually indistinguishable at times $t = 1.00$ and $t = 2.00$. The

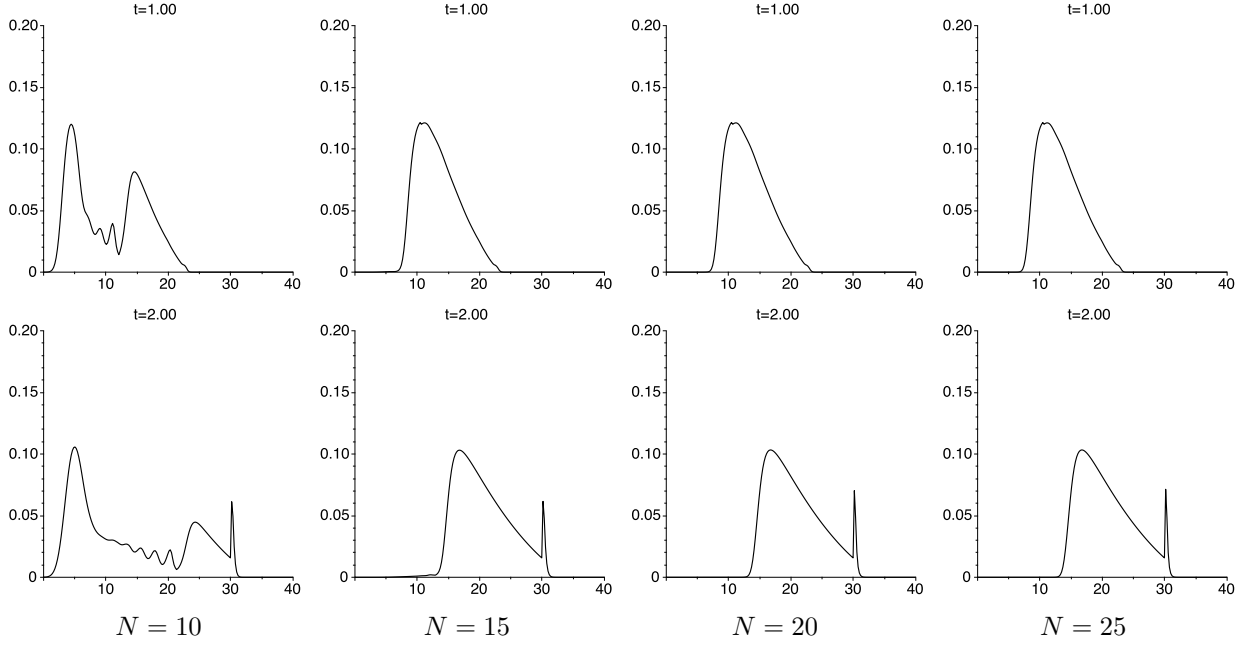


Fig. 3. Similarity of open-loop optimal solutions for different horizons $N = 10, 15, 20, 25$ (left to right) after a couple of time steps, indicating the occurrence of turnpike behavior

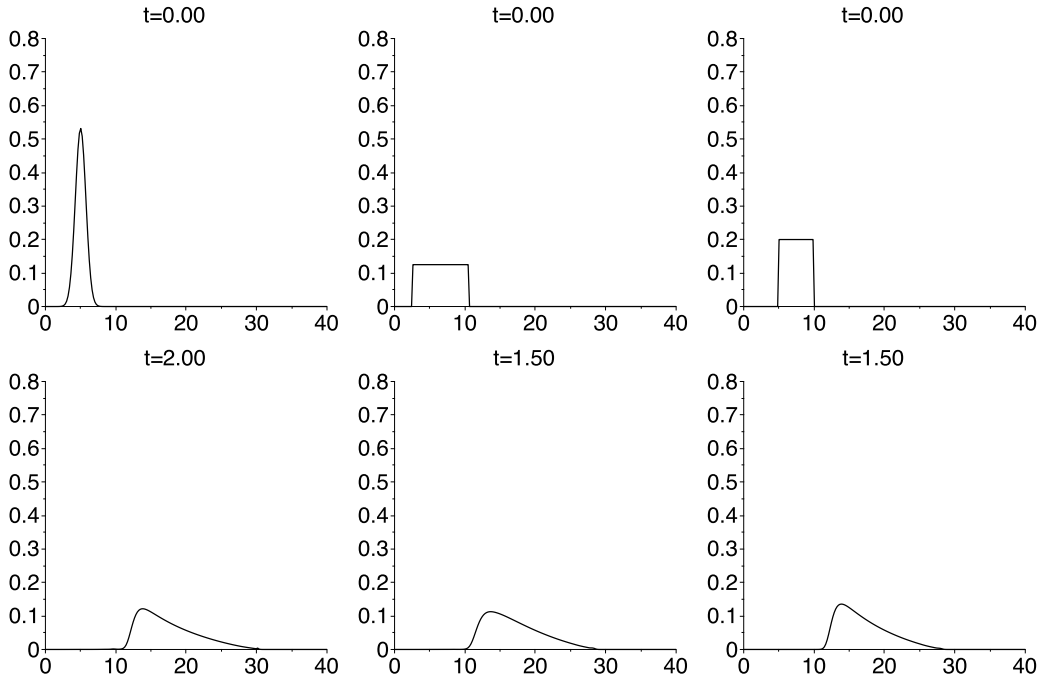


Fig. 4. Similarity of MPC solutions for different initial conditions (left to right) after a couple of time steps, indicating the occurrence of turnpike behavior

initial condition here was chosen as the uniform distribution depicted in the middle of the top row of Figure 4. This figure shows the second aspect of the turnpike property, namely that after a certain amount of time the solutions become similar independent of the initial condition. This property is true both for the open-loop and for the resulting MPC closed-loop trajectories and Figure 4 illustrates that this behaviour occurs for the MPC closed-loop trajectories for $N = 15$. Both experiments together provide a strong indication for turnpike behavior.

The occurrence of the turnpike implies via Theorem 2 that the closed-loop cost of the MPC trajectory (7) converges to the optimal value for $N \rightarrow \infty$. Notice that we evaluate (7) for $M = 30$, because for all optimization horizons $N \geq 8$ the support of the PDF $z_\mu(k)$ was contained in the interval $[30, 40]$ for all $k \geq M = 30$, hence for all such k the cost $\ell(z_\mu(k), \mu(z(k)))$ is negligible, although not exactly 0, because some control action is needed in order to prevent the density from moving back below $x = 30$. Figure (5) shows the typical behaviour of the closed-loop cost: once

the optimization horizon N has reached a threshold, the cost starts decreasing and converges to a value close to the optimal one. The threshold is related to the time that is needed to steer (parts of) the density into the desired region. If the optimization horizon is too short, the algorithm cannot figure out a strategy for reducing the stage cost of the problem.

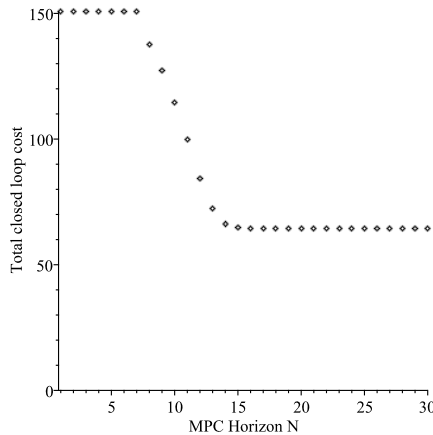


Fig. 5. Closed-loop cost for $K = 30$ for different optimization horizons $N = 1, \dots, 30$

CONCLUSION

Our numerical results show that MPC is a very promising solution method for mean-field optimal control problems. Particularly, the method is perfectly suited to handle the nonlinear coupling occurring in our model problem. Moreover, the numerical results confirm convergence of the closed-loop cost and occurrence of the turnpike property that is needed for proving this convergence. The results strongly motivate further research on MPC for mean-field optimal control problems, regarding both the derivation of rigorous theoretical approximation results and the development of efficient numerical schemes.

Acknowledgement We are indebted to Arthur Fleig for his support with the numerical simulations.

REFERENCES

- Annunziato, M. and Borzì, A. (2013). A Fokker-Planck control framework for multidimensional stochastic processes. *J. Comput. Appl. Math.*, 237(1), 487–507.
- Annunziato, M. and Borzì, A. (2018). A Fokker-Planck control framework for stochastic systems. *EMS Surveys in Math. Sciences*, 5(1/2), 65–98.
- Annunziato, M., Borzì, A., Nobile, F., and Tempone, R. (2014). On the connection between the Hamilton-Jacobi-Bellman and the Fokker-Planck control frameworks. *Appl. Math.*, 5(16), Article ID 49428, 8 pages.
- Bensoussan, A., Frehse, J., and Yam, P. (2013). *Mean field games and mean field type control theory*. SpringerBriefs in Mathematics. Springer, New York.
- Chang, J. and Cooper, G. (1970). A practical difference scheme for Fokker-Planck equations. *Journal of Computational Physics*, 6(1), 1–16.
- Chatterjee, D. and Lygeros, J. (2015). On stability and performance of stochastic predictive control techniques. *IEEE Trans. Automat. Control*, 60(2), 509–514.
- Dorfman, R., Samuelson, P.A., and Solow, R.M. (1958). *Linear programming and economic analysis*. A Rand Corporation Research Study. McGraw-Hill, New York-Toronto-London.
- Faulwasser, T., Korda, M., Jones, C.N., and Bonvin, D. (2017). On turnpike and dissipativity properties of continuous-time optimal control problems. *Automatica*, 81, 297–304.
- Fleig, A. and Grüne, L. (2019). On dissipativity of the Fokker-Planck equation for the Ornstein-Uhlenbeck process. In *Proceedings of the 3rd IFAC Workshop on Control of Systems Governed by Partial Differential Equations — CPDE 2019*, IFAC-PapersOnLine 52(2), 13–18. Oaxaca, Mexico.
- Grammatico, S., Gentile, B., Parise, F., and Lygeros, J. (2015). A Mean Field control approach for demand side management of large populations of Thermostatically Controlled Loads. In *Proceedings of the 2015 European Control Conference — ECC 2015*, 3548–3553. Linz, Austria.
- Grüne, L. (2016). Approximation properties of receding horizon optimal control. *Jahresber. DMV*, 118(1), 3–37.
- Grüne, L. and Müller, M.A. (2016). On the relation between strict dissipativity and the turnpike property. *Syst. Contr. Lett.*, 90, 45–53.
- Herty, M. and Pareschi, L. (2010). Fokker-Planck asymptotics for traffic flow models. *Kinet. Relat. Models*, 3(1), 165–179.
- McKenzie, L.W. (1986). Optimal economic growth, turnpike theorems and comparative dynamics. In *Handbook of Mathematical Economics, Vol. III*, 1281–1355. North-Holland, Amsterdam.
- Mohammadi, M. and Borzì, A. (2015). Analysis of the Chang-Cooper discretization scheme for a class of Fokker-Planck equations. *J. Numer. Math.*, 23(3), 271–288.
- Nourian, M., Caines, P.E., Malhamé, R.P., and Huang, M. (2013). Nash, social and centralized solutions to consensus problems via mean field control theory. *IEEE Trans. Automat. Control*, 58(3), 639–653.
- Roy, S., Annunziato, M., and Borzì, A. (2016). A Fokker-Planck feedback control-constrained approach for modelling crowd motion. *J. Comput. Theor. Transp.*, 45(6), 442–458.
- Trélat, E., Zhang, C., and Zuazua, E. (2018). Steady-state and periodic exponential turnpike property for optimal control problems in Hilbert spaces. *SIAM J. Control Optim.*, 56(2), 1222–1252.