Exponential sensitivity analysis for Model Predictive Control of PDEs ⋆

Lars Grüne * **Manuel Schaller** * **Anton Schiela** *

* *Chair of Applied Mathematics, Department of Mathematics, University of Bayreuth, 95447 Bayreuth, Germany (e-mail: author@ uni-bayreuth.de).*

Abstract: Model Predictive Control (MPC) is a control method in which the solution of optimal control problems on infinite or indefinitely long horizons is split up into the successive solution of optimal control problems on relatively short finite time horizons. Only a first part of this solution with given length is implemented as a control for the longer, possibly infinite horizon. Motivated by this application, we analyze the propagation of discretization errors in the context of optimal control of abstract evolution equations in infinite dimensional spaces. Using a particular stability property, one can show that indeed the error decays exponentially in time, leading to very efficient time and space discretization schemes for MPC. In particular, one can rigorously explain the behavior of goal oriented error estimation algorithms used in this context. Furthermore, an exponential turnpike theorem will be derived. We give particular applications of this abstract theory to admissible control of hyperbolic equations, nonautonomous and semilinear parabolic equations. Eventually, we present several numerical examples illustrating the theoretical findings.

Keywords: Control of Distributed Parameter Systems; Optimal Control; Stability.
AMS MSC 2010: 65M50, 49K40, 35Q93.

1. INTRODUCTION

Model Predictive Control (MPC) is a feedback control technique used in many industrial applications such as e.g. chemical process engineering, electrical engineering, aerospace or automotive engineering, cf. Camacho and Bordons (2007); Qin and Badgwell (2003). The goal of MPC is to solve an in general nonlinear optimal control problem (OCP) on an indefinite or infinite horizon in a feedback manner. To this end, the task is split up into the successive solution of OCPs on a finite horizon T, where only an initial part of the optimal control of length τ is then implemented in the plant. Then, after some time, the state of the plant is measured (or observed) and the process is repeated. Thus, on the one hand, a nonlinear feedback law is obtained, and on the other hand, every evaluation of the feedback law involves the solution of an OCP. While the former leads to robustness, the latter raises hope that the generated trajectory is in some way optimal on the indefinite or infinite horizon. This property can rigorously be proven, if the dynamics satisfy a so-called turnpike property, depicted in Figure 1, meaning that the solution to the dynamic OCP is close to the solution of the corresponding steady state property for the majority of the time, cf. Grüne (2016). In-depth mathematical analysis of MPC-schemes can be found in e.g. Grüne and Pannek (2016); Rawlings et al. (2017). In view of MPC, we will study two important properties of infinite dimensional optimal control problems in a very general manner.

⋆ This work was supported by the DFG Grants GR 1569/17-1 and SCHI 1379/5-1.

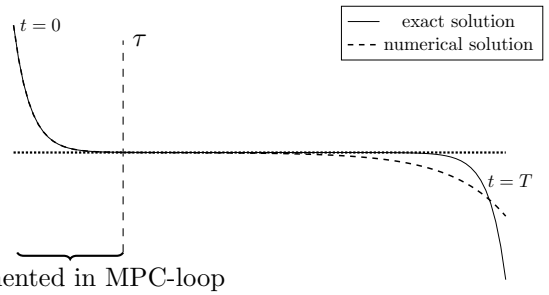

Fig. 1. Turnpike behavior of the OCP's solution, MPC-feedback and depiction of a numerical solution being (only) accurate on the initial part $[0, \tau]$.

One the one hand, we will give sufficient conditions for the turnpike property, being important to establish (quasi)optimality of the generated closed-loop trajectory. On the other hand, we aim to show that the OCP occurring in every MPC-loop can be solved efficiently by goal oriented adaptive gridding methods. To this end, we will show that the initial part of the optimal trajectory, i.e. the MPC-feedback, is only slightly influenced by perturbations happening in the far future. This property is sketched in Figure 1 and allows to use grids which are fine on a small initial part $[0, \tau]$ and coarsen up towards T, leading to a very efficient discretization for MPC, if $\tau \ll T$. This rigorously explains the behavior of goal oriented error estimation techniques which mostly only refine the time and space grid on the initial part $[0, \tau]$ when aiming to reduce the MPC-feedback error. We note that this property is not clear a priori, as the optimization involves a backwards-in-

time adjoint equation, which could propagate perturbations from close to the end time towards the initial part.

In this extended abstract we will recall some of the results obtained in Grüne et al. (2019a) and Grüne et al. (2019b) considering the linear quadratic setting and accompany them with novel results concerning nonlinear dynamics and numerical examples.

2. LINEAR ABSTRACT SCALING THEOREMS

We will first consider the case of a linear quadratic optimal control problem:

$$\begin{aligned} \min \frac{1}{2}\int_0^T \|C(x-x_d)\|_Y^2 + \|R(u-u_d)\|_U^2 \\ \text{s.t.} \quad x' = Ax + Bu + f, \quad x(0) = x_0. \end{aligned} \tag{1}$$

We assume that A generates a C0-semigroup on a Hilbert space X with norm denoted by $\|\cdot\|$, $x_0 \in X$, $f \in L_1(0,T;X)$ and $C \in L(X,Y)$, $B \in L(U,X)$, $R \in L(U,U)$ with $\|Ru\|_U \geq \gamma\|u\|^2$ for $\gamma > 0$, where U and Y are Hilbert spaces. We will discuss possible generalizations of this setting at the end of this section in Remark 1.

Formally, after elimination of the control via $u = (R^*R)^{-1}B^*\lambda + u_d$, the first order optimality conditions read

$$\underbrace{\begin{pmatrix} C^*C & -\frac{d}{dt} - A^* \\ 0 & E_T \\ \frac{d}{dt} - A & -BQ^{-1}B^* \\ E_0 & 0 \end{pmatrix}}_{=:M} \begin{pmatrix} x \\ \lambda \end{pmatrix} = \underbrace{\begin{pmatrix} C^*Cx_d \\ 0 \\ Bu_d + f \\ x_0 \end{pmatrix}}_{=:l},$$

where $E_t x = x(t)$ is a time evaluation operator. We will base our analysis on the couple (x,λ) solving this equation.

Perturbations decay exponentially. Consider a perturbed solution, solving

$$M\begin{pmatrix} \tilde{x} \\ \tilde{\lambda} \end{pmatrix} = l + \varepsilon,$$

$u = (R^*R)^{-1}B^*\tilde{\lambda} + u_d$ where ε is a perturbation bounded by $\|e^{-\mu t}\varepsilon(t)\|_{L_1(0,T;X)} \leq \rho$ with $\rho > 0$. This perturbation resembles the residual in function space of a discrete solution, i.e. it reflects the local discretization error; if ε is large at a time instance, this means that the grid is coarse and vice versa. The exponential scaling in the $L_1(0,T;X)$ norm allows for exponentially increasing perturbations ε which model grid coarsening exponentially towards T.

Theorem 1. Let (A,B) be exp. stabilizable and (A,C) be exp. detectable. Then there is $c, \mu > 0$ independent of T such that

$$\|e^{-\mu t}\begin{pmatrix} x(t) - \tilde{x}(t) \\ u(t) - \tilde{u}(t) \\ \lambda(t) - \tilde{\lambda}(t) \end{pmatrix}\| \leq c\rho \quad \text{for all } t \in [0,T].$$

The interpretation of this theorem is as follows: Consider perturbations which e.g. increase exponentially in time with rate less than μ, i.e. the estimate $\|e^{-\mu t}\varepsilon(t)\|_{L_2(0,T;X)} \leq \rho$ is satisfied. Then, the effect of this exponentially increasing perturbation on the optimal triple is of local nature, i.e. for t small, $x(t) - \tilde{x}(t)$ is small as the scaling term $e^{-\mu t}$ approaches one. This means, that perturbations happening in the far future have a negligible effect on the MPC-feedback if $\tau \ll T$.

An exponential turnpike property. Under the very same stabilizability assumptions, an exponential turnpike property can be derived. To this end, to ensure the existence of optimal steady states, we assume x_d, u_d and f to be independent of time. We denote by $(\bar{x}, \bar{u})$ the optimal steady state solving

$$\begin{pmatrix} C^*C & -A^* \\ -A & -BQ^{-1}B^* \end{pmatrix}\begin{pmatrix} \bar{x} \\ \bar{\lambda} \end{pmatrix} = \begin{pmatrix} C^*Cx_d \\ Bu_d + f \end{pmatrix},$$

and $\bar{u} = (R^*R)^{-1}B^*\bar{\lambda} + u_d$.

Theorem 2. Let (A,B) be exp. stabilizable and (A,C) be exp. detectable. Then there is $c, \mu > 0$ independently of T such that

$$\|\begin{pmatrix} x(t) - \bar{x} \\ u(t) - \bar{u} \\ \lambda(t) - \bar{\lambda} \end{pmatrix}\| \leq c(e^{\mu t} + e^{-\mu(T-t)}) \quad \text{for all } t \in [0,T].$$

Remark 1. The results from Theorem 1 and 2 can be refined by assuming more structure, i.e. A generating an analytic semigroup, cf. Grüne et al. (2019a). In this case, maximal parabolic regularity can be used to reduce the spatial regularity assumptions on the perturbations and to obtain estimates in Sobolev norms. Second, using a slightly stronger notion of stability, Theorem 1 can be extended to non-autonomous equations, where $A = A(t)$. Third, boundedness of control and observation operators B and C can be replaced by a weaker property, namely admissibility, cf. (Grüne et al., 2019b, Section 5). Finally, the initial condition on x can be accompanied by a terminal condition $x(T) = x_T$ for $x_T \in X$, if (A,B) is exactly controllable, cf. (Grüne et al., 2019b, Section 7).

3. NONLINEAR DYNAMICS

First, we give an alternative interpretation of Theorem 1 and 2: Under stabilizability and detectability assumptions, the operator M is invertible in scaled spaces, and the constants involved are independent of the time horizon. We will use this property to generalize the theory to nonlinear parabolic equations. In order to keep the presentation simple, we will introduce nonlinearity to the OCP via adding a monotone decreasing nonlinearity $f(x)$, e.g. $f(x) = -x^3$, leading to the following dynamics

$$x' = Ax + f(x) + Bu \qquad x(0) = x_0.$$

We will study (1) with these nonlinear dynamics. Deriving first order optimality conditions leads, again after elimination of the control to the nonlinear system of equations at the optimal pair $z := (x,\lambda)$

$$\underbrace{\begin{pmatrix} C^*Cx - \frac{d}{dt}\lambda - A^*\lambda - f'(x)\lambda \\ E_T\lambda \\ \frac{d}{dt}x - Ax - f(x) - BQ^{-1}B^*\lambda \\ E_0 x \end{pmatrix}}_{=:M(z)} = \underbrace{\begin{pmatrix} C^*Cx_d \\ 0 \\ Bu_d \\ x_0 \end{pmatrix}}_{=:l}. \tag{2}$$

We will now apply an implicit function theorem to the nonlinear function $G(x,\varepsilon) := M(z) - l - \varepsilon$. First, we observe that z solving (2) satisfies $G(z,0) = 0$. Analogously to the linear case, we define a perturbed variable $\tilde{z} := (\tilde{x}, \tilde{\lambda})$ solving $G(\tilde{z},\varepsilon) = 0$. The question on how the difference

$\tilde{z}-z$ depends on ε is now answered by the implicit function theorem, cf. (Zeidler, 1986, Theorem 4.B) in exponentially scaled spaces with weight $e^{-\mu t}$ or $\frac{1}{e^{-\mu t}+e^{-\mu(T-t)}}$ assuming stabilizability and detectability of the linearized equations. In the case of a turnpike property, the linearization point is the optimal steady state. However, to bound the influence of exponential perturbations, we would like to linearize around the exact optimal solution. As a consequence, the equation is intrinsically non-autonomous as $A'(x(t))$ implicitly depends on time. To this end, we assume a different stability property which also covers non-autonomous systems, introduced in Grüne et al. (2019a) under the name V-exponential stability. In a nutshell, the stability assumption for the semigroup generated by $(A+BK)$ with feedback operator K occurring in the definition of exponential stabilizability of (A, B) is replaced by a coercivity assumption on $-(A+BK)$ in a Sobolev space. All in all, Theorem 1 can be generalized to the nonlinear case as follows.

Theorem 3. Let $(A+f'(x), B)$, $(A+f'(x), C)$ be stabilizable in the sense of (Grüne et al., 2019a, Definition 3.6). Moreover, assume $\|e^{-\mu t}\varepsilon\|_{L_2(0,T;H^{-1}(\Omega))} \leq \rho$ with $\rho > 0$ sufficiently small, where Ω is the spatial domain. Then there is $\mu, c > 0$ independent of T such that

$$\|e^{-\mu t}\begin{pmatrix} x(t)-\tilde{x}(t) \\ u(t)-\tilde{u}(t) \\ \lambda(t)-\tilde{\lambda}(t)\end{pmatrix}\| \leq c\rho \qquad \forall t \in [0,T].$$

The analogous generalization of the turnpike property stated in Theorem 2 reads:

Theorem 4. Let $(A'(x)+f'(x), B)$, $(A'(x)+f'(x), C)$ be stabilizable in the sense of (Grüne et al., 2019a, Definition 3.6). Moreover, assume $\|\left(\bar{\lambda}, x_0-\bar{x}\right)\| \leq \rho$ with $\rho > 0$ sufficiently small. Then there is $\mu, c > 0$ independent of T such that

$$\|\begin{pmatrix} x(t)-\bar{x} \\ u(t)-\bar{u} \\ \lambda(t)-\bar{\lambda}\end{pmatrix}\| \leq c\rho\left(e^{-\mu t}+e^{-\mu(T-t)}\right) \qquad \forall t \in [0,T].$$

4. NUMERICAL RESULTS

The theoretical results of Theorem 1 and 3 justify the use of fine grids on the implementation horizon $[0, \tau]$ which coarsen up towards T for the solution of the underlying optimal control problems for the evaluation of the MPC-feedback law. This could be done a priori, with the downside that only qualitative results are at hand which yield no quantitative information needed for an implementation. For this reason we will use goal oriented a posteriori error estimation, cf. Meidner and Vexler (2007). This not only allows to estimate the error quantitatively, but also to estimate the error in an arbitrary measure, the so called quantity of interest (QOI). The natural quantity of interest for error estimation in an MPC-context is a truncated version of the cost function, i.e. to define the QOI via

$$I(x,u) := \frac{1}{2}\int_0^\tau \|C(x-x_d)\|_Y^2 + \|R(u-u_d)\|_U^2. \qquad (3)$$

The aim of goal oriented error estimation is now to adaptively construct time and space grids, such that the numerical solution of the OCP on these grids has a small error in the QOI, i.e. $|I(x,u)-I(\tilde{x},\tilde{u})|$ being small. The question we are interested in is now the following: How do the grids, being specialized for this QOI look like? Theorem 1 and 3 suggest, that in order to have a high-quality MPC-feedback, i.e. to be accurate on the beginning of the horizon, the grids generated by the goal oriented error estimation are fine close to the beginning of the horizon, but not towards T. We will illustrate this property at the example of optimal control with a 2d linear heat equation on the unit square and time interval $[0, 9]$.

$$\begin{aligned} \min\ & \frac{1}{2}\int_0^T \|x(t)-x_{\text{ref}}\|^2_{L_2(\Omega)} + \alpha\|u(t)\|^2_{L_2(\Omega)}\,dt, \\ & \dot{x} = d\Delta x + sx + u, \qquad x(0) = 0 \end{aligned} \qquad (4)$$

with homogeneous Dirichlet boundary conditions, x_{ref} is a time independent reference, $d > 0$ is the heat conductivity and $s \in \mathbb{R}$ an instability parameter.

Time refinement. We first consider adaptive time refinement for the optimal control problem (4), where we set $\alpha = 10^{-1}$, $s = 5$ and $\tau = 0.5$.

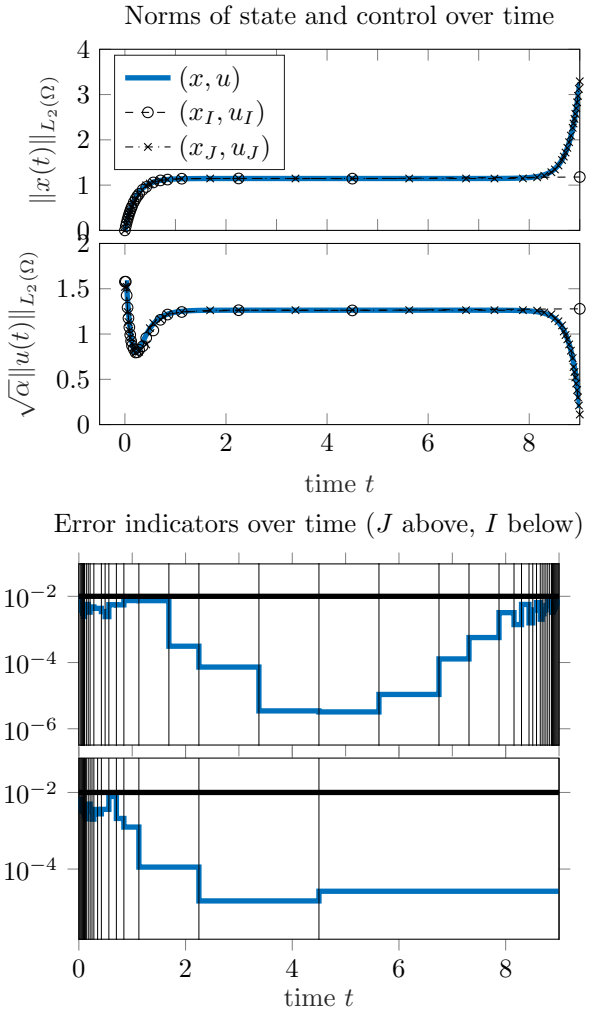

Fig. 2. Above: Norm of optimal state and control over time. The reference trajectory was computed on a very fine time grid. Below: Error indicators (blue) and grid points on time grid (vertical black lines) over time. Indicators for the cost function $J(x,u)$ in the upper and indicators for truncated cost function $I(x,u)$ in the lower plot respectively.

In the upper picture of Figure 2, the optimal state and control on a very fine grid denoted by (x, u), and the optimal pair on the two adaptively refined grids is depicted, where one refinement is w.r.t. $J(x, u)$ with the solution denoted by (x_J, u_J) and the other for $I(x, u)$, with the solution denoted by (x_I, u_I). We refined the grid until the error indicators on every time grid interval for the respective objective are below 10^{-2}. We observe the turnpike property stated in Theorem 2 for all depicted solutions. The pairs (x, u) and (x_J, u_J) coincide on the whole horizon, whereas (x_I, u_I) does not exhibit the so called leaving arc at the end of the horizon. The reason for this becomes clear in the lower picture of Figure 2: As $I(x, u)$ only involves the interval $[0, \tau]$, the grid for (x_I, u_I) is very coarse towards T, rendering it impossible to reflect the leaving arc. On the other hand, the grid for (x_J, u_J) is fine at the beginning and at the end of the interval, allowing the solution to be accurate on the whole interval. However, this shows, that in order to be accurate on the whole horizon, a remarkably higher amount of time grid point is needed as opposed to being accurate on $[0, \tau]$. Figure 3 depicts the performance of using a truncated cost function instead of the full cost function for adaptivity in an MPC-loop. To this end, we fixed the maximal number of time points for the refinement procedure. Then, in every MPC iteration, the optimal control problem was solved and adaptive grid refinement was performed until the maximal number of grid points was reached. It can be seen in Figure 3 that the closed-loop cost is significantly lower when refining the grid for the objective $I(x, u)$ as opposed to $J(x, u)$. The total number of MPC iterations performed in this example was four.

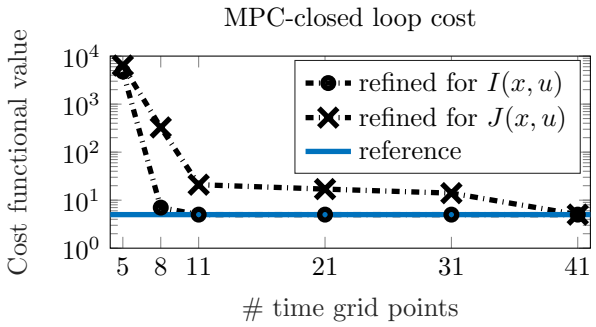

Fig. 3. Comparison of MPC-closed loop cost for refinement via $I(x, u)$ and $J(x, u)$.

Space refinement. Secondly, we will consider spatial refinement. In this case, we set $\alpha = 10^{-4}$ and $s = 0$. The grids on every time grid point are refined independently of each other. We again first compare the number of degrees of freedom needed to render the indicators for the objectives $I(x, u)$ or $J(x, u)$ on every cell below 10^{-2}. In the upper plot of Figure 4 we observe, that the required spatial degrees of freedom in every time step are significantly higher when refining for $J(x, u)$ as opposed to refining for $I(x, u)$. This again reflects the theoretical findings of Theorem 1: In order to have an accurate solution on $[0, \tau]$, the grid only needs to be fine close to $[0, \tau]$. In the lower part of Figure 4, we show how again the cost of the MPC-closed loop trajectory is lower when using $I(x, u)$ as a QOI for error estimation. We again fix the number of maximal total spatial degrees of freedom and in every MPC-loop solve the OCP with adaptive refinement until the maximal number of spatial degrees of freedom is reached.

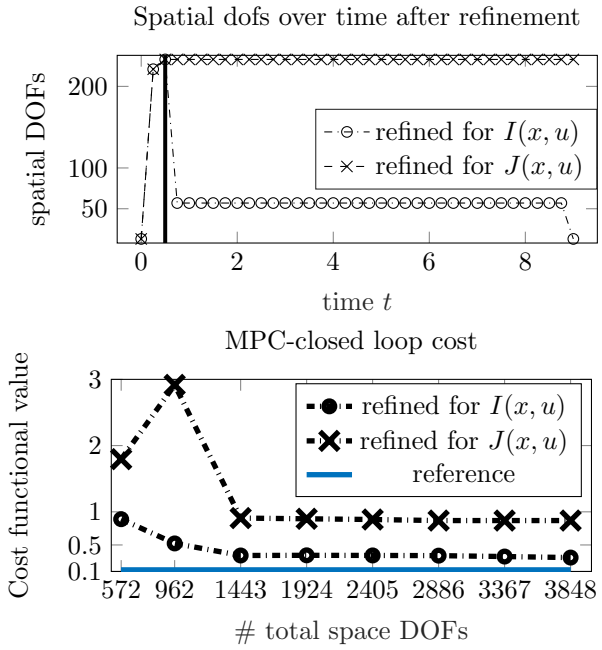

Fig. 4. Above: Spatial degrees of freedom over time after refinement. The vertical line indicates $\tau = 0.5$. Below: Comparison of MPC-closed loop cost for refinement via $I(x, u)$ and $J(x, u)$.

REFERENCES

Camacho, E.F. and Bordons, C. (2007). *Model Predictive Control*. Springer-Verlag London, 2nd edition.

Grüne, L., Schaller, M., and Schiela, A. (2019a). Sensitivity analysis of optimal control for a class of parabolic PDEs motivated by model predictive control. *SIAM Journal on Control and Optimization*, 57(4), 2753–2774. doi:10.1137/18M1223083.

Grüne, L. (2016). Approximation properties of receding horizon optimal control. *Jahresbericht Deutsche Mathematiker-Vereinigung*, 118(1), 3–37.

Grüne, L. and Pannek, J. (2016). *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer.

Grüne, L., Schaller, M., and Schiela, A. (2019b). Exponential sensitivity and turnpike analysis for linear quadratic optimal control of general evolution equations. *Journal of Differential Equations*. Available online.

Meidner, D. and Vexler, B. (2007). Adaptive space-time finite element methods for parabolic optimization problems. *SIAM Journal on Control and Optimization*, 46(1), 116–142.

Qin, S.J. and Badgwell, T.A. (2003). A survey of industrial model predictive control technology. *Control engineering practice*, 11(7), 733–764.

Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI.

Zeidler, E. (1986). *Nonlinear Functional Analysis and its Applications- I: Fixed-Point Theorems*. Springer-Verlag New York.