

Numerical Verification of Turnpike and Continuity Properties for Time-Varying PDEs

Lars Grüne¹, Simon Pirkelmann¹

¹*Chair of Applied Mathematics, Mathematical Institute, University of Bayreuth, Germany*

December 17, 2018

Abstract

To prove approximate closed loop optimality of economic model predictive control of time varying systems, continuity assumptions for the optimal value function and the turnpike property are sufficient conditions. It can sometimes be difficult to prove these assumptions analytically for a given example. In this paper we present a numerical approach that aims to verify the assumptions by simulations for a system involving a convection-diffusion equation with boundary control. The results show that the assumptions are realistic and can be met for practical problems.

Contents

1	Introduction	1
2	Problem statement and definitions	2
3	Model predictive control	3
3.1	Assumptions for proving MPC convergence properties	4
4	Numerical Verification of Turnpike and Continuity	6
4.1	Approximate computation of an optimal operation trajectory	7
4.2	Verifying turnpike	7
4.3	Verifying continuity	8
5	Conclusion	9

1 Introduction

Model predictive control (MPC) has emerged as a method to reach a previously known equilibrium point or to follow a given reference trajectory at which the system shows a desired behavior (e.g. stability or energy efficiency). Applications of the method to time-invariant systems are numerous and there are many results that show convergence to the specified trajectory ([5], [9]).

The tracked reference trajectory or equilibrium point often comes from additional knowledge about the system or it is determined by preceding optimization problems. However, sometimes it can be difficult to calculate a suitable reference trajectory in advance because the regime of optimal operation may be quite complex (cf. [8], [2]). A promising alternative approach is the use of economic model predictive control, where the desired system behavior is not prescribed but instead is implicitly found by the controller by means of incorporating an economic criterion into the stage cost. For time-invariant systems this approach has been studied in depth (see e.g. [3] for a comprehensive overview of existing stability results).

As we have shown in earlier studies ([6], [7]), the method also works well for time-varying systems and estimates of the control performance and convergence can be proven.

In the present paper we revisit performance results from [6]. In that work, the turnpike property and a continuity property of the optimal value function were employed as sufficient conditions in the convergence proofs. However,

until now it was not clear whether these conditions can be expected to hold for control systems governed by PDEs, i.e., whether the approach from [6] is feasible for such systems and thus worth to be pursued. In this paper, we provide numerical evidence that this is the case. To this end, simulations for a partial differential equation were carried out. The promising results thus serve as a basis for future analytical considerations. The paper is organized as follows: in the first section we introduce the problem statement and define an optimality notion and the concept of optimal operation in the time-varying context. The next section presents the MPC algorithm and introduces two key assumptions, the turnpike property and continuity of the optimal value functions, which are necessary to prove a performance of the MPC method. In the final section we present an example for which we verify numerically that the two assumptions are satisfied.

2 Problem statement and definitions

Consider a discrete-time time-varying control system

$$x(k+1) = f(k, x(k), u(k)), \quad x(0) = x, \quad (1)$$

where X and U are state and control space and $f : \mathbb{N}_0 \times X \times U \rightarrow X$. The index $k \in \mathbb{N}_0$ denotes time, $x(k) \in X$ is the state of the system at time k and $u(k) \in U$ is the control.

We denote a trajectory of the system starting at time k from initial state $x \in X$ controlled by $u \in U^N$, $N \in \mathbb{N}$ by $x_u(\cdot; k, x)$. State and control constraints are expressed by the sets of admissible states at time k denoted by $\mathbb{X}(k) \subseteq X$ and the sets of admissible control values for $x \in \mathbb{X}(k)$ by $\mathbb{U}(k, x) \subseteq U$. By $\mathbb{U}^N(k, x)$ we denote the sets of admissible control sequences for initial state $x \in \mathbb{X}(k)$ up to time $k+N$, i.e. control sequences $u \in U^N$ satisfying

$$u(j) \in \mathbb{U}(k+j, x_u(j; k, x)) \text{ and } x_u(j+1; k, x) \in \mathbb{X}(k+j+1)$$

for all $j = 0, \dots, N-1$. By the set $\mathbb{U}^\infty(k, x)$ we denote the extension of this definition to infinite control sequences.

We consider a stage cost function $\ell : \mathbb{N}_0 \times X \times U \rightarrow \mathbb{R}$ and the cost functional

$$J_\infty(k, x, u) = \sum_{j=0}^{\infty} \ell(k+j, x_u(j; k, x), u(j)) \quad (2)$$

The aim is to compute a feasible control sequence $u \in \mathbb{U}^\infty(k, x)$ which minimizes the objective value $J_\infty(k, x, u)$. For general stage cost $J_\infty(k, x, u)$ is not necessarily a finite number for any control sequence u , instead it may be unbounded. This presents a problem for properly defining an optimization problem since if $J_\infty(k, x, u)$ is infinite for all control sequences u we need to explain how we can compare two control sequences based on their objective value. A way to do this is applying the notion of overtaking optimality going back to [4]. Instead of considering the cost of two control sequences separately we look at the difference of the costs. Although both control sequences produce infinite costs, the difference between the two can still be finite.

Definition 1 (Overtaking optimality). *Let $x \in \mathbb{X}(k)$ and consider a control sequence $u^* \in \mathbb{U}^\infty(k, x)$ with corresponding state trajectory $x_{u^*}(\cdot; k, x)$. The pair (x_{u^*}, u^*) is called overtaking optimal if*

$$\liminf_{K \rightarrow \infty} \sum_{j=0}^{K-1} \ell(k+j, x_u(j, x), u(j)) - \ell(k+j, x_{u^*}(j, x), u^*(j)) \geq 0 \quad (3)$$

for all $u \in \mathbb{U}^\infty(k, x)$.

A control sequence is considered to be optimal if its cost is overtaken by the cost of any other control sequence at some point. This enables us to decide which of two infinite control sequences is better when starting at a fixed initial value x .

We can now formulate our main goal as an optimization problem in the overtaking optimal sense:

$$\underset{u \in \mathbb{U}^\infty(k, x)}{\text{minimize}} J_\infty(k, x, u) \quad (4)$$

We will assume a solution of this problem exists and will denote it by u_∞^* . In the previous definition the initial state x occurs explicitly, that is for a given initial state the definition formalizes which control sequence we consider optimal. Now the initial state is no longer fixed but treated as an additional optimization variable. For any possible initial value we consider all resulting feasible system trajectories. Out of those trajectories we will call those trajectories *optimal operation trajectories* that satisfy the inequality in the following definition.

Definition 2 (Optimal operation). Let $x \in \mathbb{X}(k)$ and consider a control sequence $u^* \in \mathbb{U}^\infty(k, x)$ with corresponding state trajectory $x^* = x_{u^*}(\cdot; k, x)$. We say the system (1) is optimally operated at (x^*, u^*) if

$$\liminf_{K \rightarrow \infty} \sum_{j=0}^{K-1} \ell(k+j, x_u(j, x'), u(j)) - \ell(k+j, x^*(j), u^*(j)) \geq 0 \quad (5)$$

for all $x' \in \mathbb{X}(k)$ and $u \in \mathbb{U}^\infty(k, x')$.

Note that there may in general exist multiple optimal operation trajectories of the system that in some sense are equivalent since they all satisfy Definition 2. For time-varying systems it is in general difficult to classify the optimal operation trajectories further. They arise from the complex interplay of stage cost, dynamics and constraints of the problem and are notoriously difficult to compute. For the remainder of this paper we will assume that an optimal operation trajectory of the system exists even if it cannot be specified explicitly. Ideally, we would like to compute an optimal trajectory in order to operate our system there. As we have shown in previous papers ([6], [7]) it is possible to achieve this at least approximately by using MPC. We will introduce the method in the following section and then state some results that provide performance guarantees under certain assumptions.

3 Model predictive control

Instead of solving problem (4) on the infinite horizon we fix $N \in \mathbb{N}$ and consider the truncated cost functional

$$J_N(k, x, u) = \sum_{j=0}^{N-1} \ell(k+j, x_u(j; k, x), u(j)). \quad (6)$$

In each step of the MPC algorithm we solve the following optimal control problem

$$\underset{u \in \mathbb{U}^N(k, x)}{\text{minimize}} \quad J_N(k, x, u). \quad (7)$$

We assume that a minimizer u_N^* of this problem always exists. Note that since we only consider finite horizons N and thus $J_N(k, x, u)$ will always be finite we do not need the concept of overtaking optimality here.

Starting at initial value $x = x(k_0) \in \mathbb{X}(k_0)$ the key idea of MPC is to repeatedly solve problem (7) and apply only the first element of the resulting optimal control sequence u_N^* in order to produce an approximate solution to problem (4). In this way a trajectory $x_{\mu_N}(\cdot, x)$ is generated which we call *closed loop trajectory*.

The model predictive control algorithm is summarized in Algorithm 1.

Algorithm 1 (MPC algorithm). For each time instant $k = k_0, k_0 + 1, \dots$:

1. Measure the current state $x = x(k)$ of the system.
2. Solve the optimal control problem (7) in order to obtain the optimal control sequence $u_{N,x}^*$.
3. Apply the first element of $u_{N,x}^*$ as a control to the system during the next sampling period, i.e. use the feedback law $\mu_N(k, x) := u_{N,x}^*(0)$.
4. Set $k := k + 1$.

An important quantity for assessing the quality of the solution is the cost of the closed loop for L time steps defined by

$$J_L^{cl}(k_0, x, \mu_N) = \sum_{j=0}^{L-1} \ell(k_0+j, x_{\mu_N}(j, x), \mu_N(k_0+j, x_{\mu_N}(j, x))).$$

In the following we present two key assumptions that are sufficient for proving a performance bound for the closed loop cost.




Figure 1: An illustration of the finite horizon turnpike property for a time-varying optimal operation trajectory x^* .

3.1 Assumptions for proving MPC convergence properties

The first crucial assumption is the so-called turnpike property. It demands that open-loop trajectories of the finite and infinite horizon optimal control problems (4) and (7) are most of the time close to an optimal operation trajectory (x^*, u^*) . This is formalized in the next definition.

Assumption 1 (Turnpike property). *Consider a trajectory pair (x^*, u^*) at which the system (1) is optimally operated. We assume that the optimal control problem (7) has the turnpike property at (x^*, u^*) , i.e. that the following holds:*

There exists $\sigma \in \mathcal{L}^1$ such that for each $k \in \mathbb{N}_0$, each optimal trajectory $x_{u_N}^(\cdot, x)$, $x \in \mathbb{X}(k)$ and all $N, P \in \mathbb{N}$ there is a set $\mathcal{Q}(k, x, P, N) \subseteq \{0, \dots, N\}$ with $\#\mathcal{Q}(k, x, P, N) \leq P$ and*

$$|(x_{u_N}^*(M, x), u_N^*(M))|_{(x^*(k+M), u^*(k+M))} \leq \sigma(P)$$

for all $M \in \{0, \dots, N\} \setminus \mathcal{Q}(k, x, P, N)$.

Furthermore, we assume that the infinite horizon optimal control problem (4) has the turnpike property, i.e. there exists $\rho \in \mathcal{L}$ such that for each $k \in \mathbb{N}_0$, each optimal trajectory $x_{u_\infty}^(\cdot, x)$, $x \in \mathbb{X}(k)$ and all $P \in \mathbb{N}$ there is a set $\mathcal{Q}(k, x, P, \infty) \subseteq \mathbb{N}_0$ with $\#\mathcal{Q}(k, x, P, \infty) \leq P$ and*

$$|(x_{u_\infty}^*(M, x), u_\infty^*(M))|_{(x^*(k+M), u^*(k+M))} \leq \rho(P)$$

for all $M \in \mathbb{N}_0 \setminus \mathcal{Q}(k, x, P, \infty)$.

The definition of the turnpike property may seem technically involved. However, it is in fact easy to visualize. What the assumption demands is that for each open-loop trajectory there can only be a finite number (at most P) of time instances when the trajectory is far from the optimal operation trajectory. In addition we can specify a bound on the distance for the points which are close to the turnpike. Figure 1 illustrates this idea graphically.

The second assumption we need is a property regarding the optimal value functions of problems (4) and (7). As noted before, the objective value $J_\infty(k, x, u)$ may become infinite for all possible control sequences u . This means also the optimal value function, classically defined as $V_\infty(k, x) := \inf_u J_\infty(k, x, u)$, would be infinite for all k and x . To prevent this we instead consider a shifted cost functional, which deducts the cost of an optimal operation trajectory (x^*, u^*) , and define the optimal value function in terms of this shifted cost. For consistency we also use the shifted cost functional for the definition of the optimal value function of the finite horizon problem even though here it would be possible to define it without.

Definition 3 (Shifted cost and optimal value functions).

We define the shifted stage cost

$$\hat{\ell}(k, x(k), u(k)) := \ell(k, x(k), u(k)) - \ell(k, x^*(k), u^*(k))$$


and corresponding shifted cost functionals

$$\hat{J}_N(k, x, u) := \sum_{j=0}^{N-1} \hat{\ell}(k+j, x_u(j; k, x), u(j))$$

and

$$\hat{J}_\infty(k, x, u) := \sum_{j=0}^{\infty} \hat{\ell}(k+j, x_u(j; k, x), u(j)).$$

¹ $\mathcal{L} := \{\sigma : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+ \mid \sigma \text{ is continuous and strictly decreasing with } \lim_{s \rightarrow \infty} \sigma(s) = 0\}$



(a) First component of γ_V as a function in N for fixed $x \in \mathcal{B}_\varepsilon(x^*(k))$. Since $\gamma_V(\cdot, \|x - x^*\|)$ is monotonously decreasing in N this means the value $|\hat{V}_N(k, x) - \hat{V}_N(k, x^*(k))|$ must be bounded. (b) The second component of γ_V as a function in x for fixed horizon length N . $\gamma_V(N, \cdot)$ increases monotonically with growing distance $\|x - x^*\|$.

Figure 2: Graphical illustration of the modulus of continuity γ_V from Assumption 2.

The optimal value functions of problems (7) and (4) are defined as

$$\hat{V}_N(k, x) := \inf_{u \in \mathbb{U}^N(k, x)} \hat{J}_N(k, x, u).$$

and

$$\hat{V}_\infty(k, x) := \inf_{u \in \mathbb{U}^\infty(k, x)} \hat{J}_\infty(k, x, u).$$

It can be shown (cf. [6]) that the optimal value functions attain finite values for all k and x if the turnpike property holds and the following continuity property is satisfied.

Assumption 2 (Continuity property). We assume that the optimal value function \hat{V}_N is (approximately) continuous at x^* in the following uniform way: for each $k \in \mathbb{N}_0$ there is an open ball $\mathcal{B}_\varepsilon(x^*(k))$, $\varepsilon > 0$, around $x^*(k)$ and a function $\gamma_V : \mathbb{R}_0^+ \times \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ with $\gamma_V(N, r) \rightarrow 0$ if $N \rightarrow \infty$ and $r \rightarrow 0$, and $\gamma_V(\cdot, r)$, $\gamma_V(N, \cdot)$ monotonic for fixed r and N , such that for all $x \in \mathcal{B}_\varepsilon(x^*(k)) \cap \mathbb{X}(k)$ and all $N \in \mathbb{N}$ the inequality

$$|\hat{V}_N(k, x) - \hat{V}_N(k, x^*(k))| \leq \gamma_V(N, |x|_{x^*(k)}) \quad (8)$$

holds.

Moreover, we also assume approximate continuity of the optimal value function on the infinite horizon: for each $k \in \mathbb{N}_0$ there is an open ball $\mathcal{B}_\varepsilon(x^*(k))$, $\varepsilon > 0$, around $x^*(k)$ and a function $\omega_V \in \mathcal{K}_\infty^2$ such that for all $x \in \mathcal{B}_\varepsilon(x^*(k)) \cap \mathbb{X}(k)$ it holds

$$|\hat{V}_\infty(k, x) - \hat{V}_\infty(k, x^*(k))| \leq \omega_V(|x|_{x^*(k)}).$$

An alternative stronger condition for (8) can be used if the modulus of continuity γ_V can be chosen independently of the horizon length N . In this case, inequality (8) is replaced by

$$|\hat{V}_N(k, x) - \hat{V}_N(k, x^*(k))| \leq \varphi_V(|x|_{x^*(k)}) \quad (9)$$

with a function $\varphi_V \in \mathcal{K}_\infty$.

The following result from [6] shows that the cost of the MPC closed loop approximates the cost of an optimal operation trajectory on the infinite horizon.

Theorem 1 (cf. Thm. 1 in [6]). Let Assumptions 1 and 2 hold. Then for each time $k \in \mathbb{N}_0$, each $N \in \mathbb{N}$ sufficiently large, and an arbitrary number of time steps $L \in \mathbb{N}$ in the MPC algorithm the shifted closed loop cost satisfies

$$\hat{J}_L^{cl}(k, x, \mu_N) \leq \hat{V}_\infty(k, x) - \hat{V}_\infty(k + L, x_{\mu_N}(L, x)) + L\delta(N) \quad (10)$$

with a function $\delta \in \mathcal{L}$.

² $\mathcal{K}_\infty := \{\alpha : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+ \mid \alpha \text{ is continuous, strictly increasing and unbounded with } \alpha(0) = 0\}$

Inequality (10) in the theorem expresses that the cost of the trajectory generated by MPC approximates the cost of an initial piece of an optimal trajectory on the infinite horizon except for error terms depending on the horizon length. It implies that $J_L^{cl}(k, x, \mu_N)$ is the initial piece of an approximately overtaking optimal trajectory for the non-shifted cost with the same error $L\delta(N)$. Since this cost typically grows linearly in L , the relative error $L\delta(N)/J_L^{cl}(k, x, \mu_N)$ is bounded independently of L .

In the proof of Theorem 1 and its preceding lemmas Assumptions 1 and 2 are used as follows. Assumption 1 establishes a link between the finite horizon open loops of the MPC problem and the optimal operation trajectory (x^*, u^*) by demanding that there are time instances when they are sufficiently close. This is then exploited by Assumption 2 which enables us to 'jump' from the MPC open loop to the optimal operation trajectory without changing the cost too much.

Although these are convenient assumptions to make for the purpose of proving this theorem it is not clear that they are reasonable or realistic assumptions for real problems. In the following section we present a problem that is motivated by a practical application and demonstrate numerically that for this problem the two assumptions are satisfied and thus justify the use of our assumptions.

4 Numerical Verification of Turnpike and Continuity

In this section we present an example for which we want to verify the turnpike and continuity properties from the last section numerically. We consider a time-varying convection diffusion equation

$$\begin{aligned} \frac{\partial y}{\partial t} - \alpha \nabla^2 y + w \nabla y &= 0 \text{ on } Q := \Omega \times [0, \infty), \\ y(0) &= y_0 \text{ on } \Omega, \end{aligned} \quad (11)$$

on a domain Ω with the following conditions at the boundary $\Gamma = \Gamma_{out} \cup \Gamma_c$:

$$\begin{aligned} \frac{\partial y}{\partial n} + \gamma_{out} y &= \delta_{out} y_{out} \text{ on } \Sigma_{out} := \Gamma_{out} \times [0, \infty), \\ \frac{\partial y}{\partial n} + \gamma_c y &= \delta_c u \text{ on } \Sigma_c := \Gamma_c \times [0, \infty). \end{aligned} \quad (12)$$

In this example $y : Q \rightarrow \mathbb{R}$ is the state, $u : \Sigma_c \rightarrow \mathbb{R}$ and $w : \Omega \times [0, T] \rightarrow \mathbb{R}$ are control functions and $\gamma_c, \delta_c : \Sigma_c \rightarrow \mathbb{R}$, $\gamma_{out}, \delta_{out} : \Sigma_{out} \rightarrow \mathbb{R}$ are coefficient functions. The time-variance is due to the time-varying function $y_{out} : \Sigma_{out} \rightarrow \mathbb{R}$ at the boundary.

We consider the infinite horizon optimal control problem

$$\min_{y, u, w} J(y, u, w) = \frac{1}{2} \|u\|_{L^2(\Sigma_c)}^2 + \frac{1}{2} \|w\|_{L^2(Q)}^2 \quad (13)$$

subject to equations (11), (12) and the constraints

$$\underline{u} \leq u \leq \bar{u} \text{ on } \Sigma_c, \quad (14)$$

$$\underline{y} \leq y \leq \bar{y} \text{ on } \Omega \times [0, \infty), \quad (15)$$

with lower and upper bounds for state and control where $\Omega_y \subseteq \Omega$ is a subdomain.

This example is motivated by a problem from application from HVAC (heating, ventilation, air conditioning), where the temperature y of a room is subject to changing outside temperatures y_{out} and should be controlled in an energy efficient way by means of a heater (control u) and a ventilation unit (control w).


For the numerical implementation we consider as domain the unit interval, i.e. $\Omega = [0, 1]$, and as subdomain $\Omega_y = [\frac{1}{4}, \frac{3}{4}]$. Control and state constraints are chosen as $\bar{u} = -\underline{u} = \frac{1}{4}$,

$$\bar{y}(x, t) = -\underline{y}(x, t) = \begin{cases} \frac{3}{20}, & \text{for } x \in \Omega_y, \\ 10, & \text{for } x \in \Omega \setminus \Omega_y. \end{cases}$$


Further parameters are $\alpha = 1$, $\gamma_{out} = \delta_{out} = 10^6$, $\gamma_c = 0$ and $\delta_c = 10$. For y_{out} we choose the periodic function $y_{out}(t) = \frac{3}{10} \sin(10t)$.

To fit in the setting of the previous section the problem is discretized in time using the implicit Euler method with a sampling rate $h = 10^{-2}$ and in space using the finite element method with the *FEniCS* toolbox ([1]) with $n_y = 300$ degrees of freedom. The resulting finite dimensional optimization problems are solved using *Ipopt* ([10]). The source code is available online³.


³Bitbucket repository: <https://bitbucket.org/spirkelmann/simulations-cpde-2019/>, commit: 5e163d2



(a) Initial state $\tilde{y}_L^*(0)$ of the numerically computed optimal operation trajectories for decreasing sampling rate h . The results indicate a lack of regularity of the initial state of the optimal operation trajectory.



(b) L_2 norms of the approximate optimal operation trajectories \tilde{y}_L^* for different sampling rates h over a fixed time horizon of $T = 1.0$.



(c) L_2 norms of the approximate optimal operation trajectories \tilde{y}_L^* with a fixed sampling rate $h = 10^{-2}$ and different horizon lengths L .

Figure 3: Numerical evidence for the convergence of the approximately computed optimal operation trajectories \tilde{y}_L^* .

4.1 Approximate computation of an optimal operation trajectory

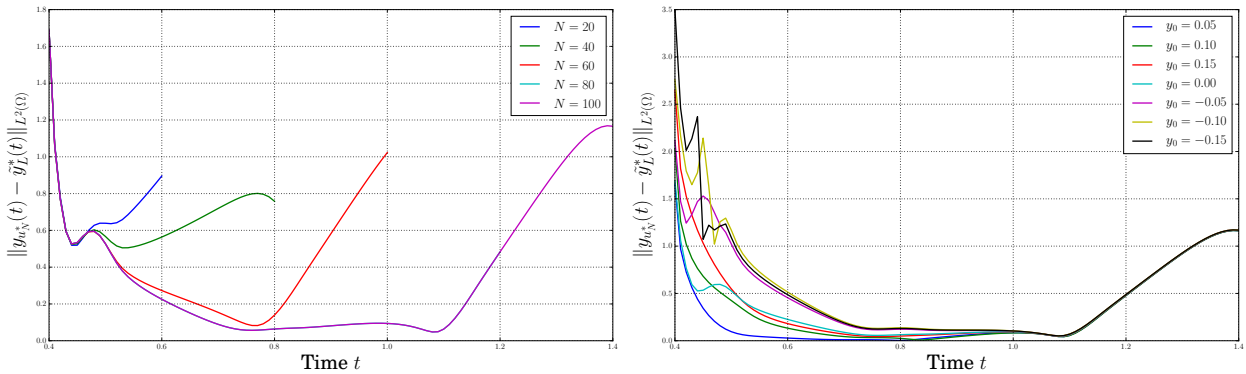
To check Assumptions 1 and 2, it is first necessary to compute an optimal operation trajectory pair (x^*, u^*) from Definition 2 called (y^*, u^*) in the notation of this section. To the best of our knowledge this cannot be done analytically. Computing it numerically is also impossible since this would involve solving an optimal control problem on an infinite horizon. Instead we compute a surrogate by choosing a large (but finite) L and solve a single open loop problem where the initial value y_0 is left as a free variable. We denote this approximation by $(\tilde{y}_L^*, \tilde{u}_L^*)$.

Numerical evidence suggests that for decreasing sampling rate $h \rightarrow 0$ the initial state of the optimal operation trajectory y^* is not a regular function in space but rather a distribution (see Figure 3a). This implies that the computed approximate optimal operation trajectory \tilde{y}_L^* with an initial state that can only approximate this distribution may not be close to the optimal operation trajectory. In practice this is not an issue because the smoothing property of the convection diffusion equation causes solutions to be sufficiently regular for each $t > 0$. In fact it can be observed in simulations that for decreasing sampling rate h the approximate optimal operation trajectories quickly converge to what we presume is the *true* optimal operation trajectory if the time horizon is sufficiently large (see Figure 3b). Moreover, for fixed sampling rate h and varying L the initial pieces of open loop solutions \tilde{y}_L^* are close which also suggests convergence to the optimal operation trajectory y^* (see Figure 3c).

For these reasons it seems justified to choose the sampling rate $h = 10^{-2}$ and the horizon of $L = 500$ to obtain an approximation of the optimal operation trajectory (y^*, u^*) for the purpose of the following simulations.

4.2 Verifying turnpike

For the numerical verification of the turnpike property from Assumption 1 we demonstrate that open loop trajectories $y_{u_N^*}$ of problem (7) are most of the time in a neighborhood of the optimal operation trajectory y^* ,



(a) L_2 distance between the optimal operation trajectory \tilde{y}_L^* and open loop trajectories $y_{u_N^*}$ of the MPC algorithm starting at time $t_0 = 0.4$ and initial state $y_0 \equiv 0$ for different horizon lengths N . (b) L_2 distance between the optimal operation trajectory \tilde{y}_L^* and open loop trajectories $y_{u_N^*}$ of the MPC algorithm with fixed horizon $N = 100$ and different constant initial states $y(x, t_0) \equiv y_0$ at time $t_0 = 0.4$.

Figure 4: Simulations show that the turnpike property holds.

which for our purposes is replaced by \tilde{y}_L^* . In addition, the assumption demands that as the horizon increases the size of the neighborhood shrinks, i.e. the open loop solutions get closer to the optimal operation trajectory. It should be noted that numerically we can only verify the finite horizon turnpike property in this way since for the infinite horizon turnpike we would need access to solutions of the problem on the infinite horizon.

Figure 4a illustrates that open loop trajectories approach the optimal operation trajectory. The trajectories exhibit an approaching and a leaving arc that is typical for the turnpike property. In Figure 4b open loop trajectories for a fixed horizon N but different initial states y_0 are shown. The plots demonstrate that all trajectories approach the same optimal operation trajectory. These results indicate that the turnpike property is likely to be satisfied for the example.

4.3 Verifying continuity

In order to verify the continuity property from Assumption 2 we need to check that the optimal value function for an initial state on the optimal operation trajectory does not change too much when we disturb this initial state. As in the previous section, we can only check the continuity assumption of the finite horizon problem numerically.

Formally, for each time point k and optimal state $y^*(k)$ we consider the quantity

$$\delta_k(N, \varepsilon) := |\hat{V}_N(k, y^*(k)) - \hat{V}_N(k, y_\varepsilon)| \quad (16)$$

for different horizon lengths N and disturbed states $y_\varepsilon \in \mathcal{B}_\varepsilon(y^*(k))$. Since the shifted optimal value function for finite horizons satisfies

$$\hat{V}_N(k, y) = V_N(k, y) - \sum_{j=0}^{N-1} \ell(k+j, y^*(k+j), u^*(k+j)) \quad (17)$$

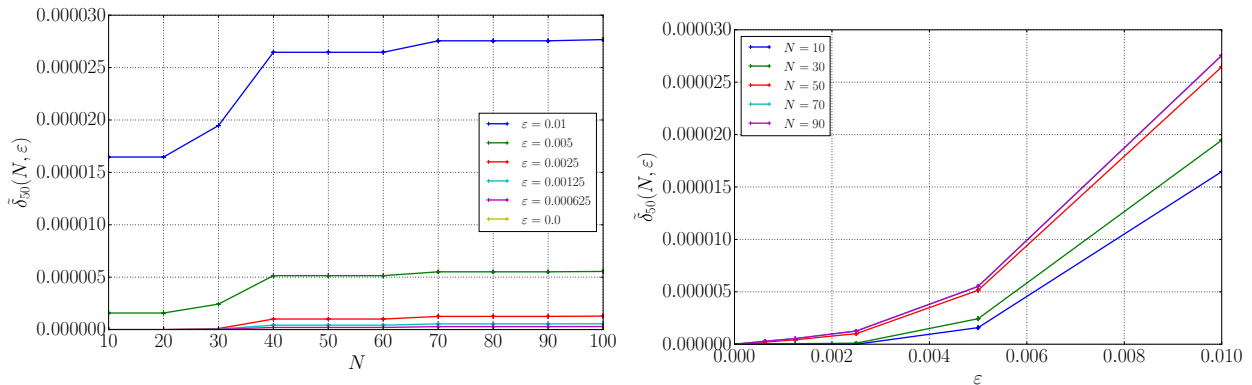
it holds that

$$\begin{aligned} \delta_k(N, \varepsilon) &= |\hat{V}_N(k, y^*(k)) - \hat{V}_N(k, y_\varepsilon)| \\ &= |V_N(k, y^*(k)) - V_N(k, y_\varepsilon)|. \end{aligned} \quad (18)$$

where V_N is the optimal value function of problem (7) given by

$$V_N(k, y) := \inf_u J_N(k, y, u). \quad (19)$$

As it turned out, in this example the stronger continuity condition (9) can be verified numerically, i.e., the function resulting from (16) can be bounded by a \mathcal{K}_∞ function that is independent of N . To check this numerically we fix a time point $t = kh$ and consider the state $\tilde{y}_L^*(t)$ on the optimal operation trajectory at that time point. Then, for decreasing $\varepsilon_i := \varepsilon_0 \frac{1}{2^i}$, $i \in \{0, \dots, n\}$, we generate a number random disturbances $y_{\varepsilon_i}^j$, $j \in \{1, \dots, m\}$ of the optimal state such that $\varepsilon_i = \|y_{\varepsilon_i}^j - \tilde{y}_L^*(t)\|_{L^2(\Omega)}$. For each of the initial conditions $y_{\varepsilon_i}^j$ generated in this way we solve the optimal control problem (7) for different horizon lengths $N \in \{N_1, \dots, N_l\}$. Thus we obtain samples of optimal value functions V_N for varying N in a neighborhood of the optimal operation



(a) Plot the difference between the optimal value function V_N at an initial value on the optimal operation trajectory \tilde{y}_L^* and the optimal value function at the disturbed states y_ε as a function in the horizon length N and for different magnitude of the disturbance ε . It can be observed that for each ε the function is bounded by a constant for increasing N .

(b) Difference between the optimal value function V_N at an initial value on the optimal operation trajectory \tilde{y}_L^* and the optimal value function at the disturbed states y_ε as a function in the magnitude of the disturbance ε for a selection different horizon lengths N . Obviously, the function can be upper bounded by a \mathcal{K}_∞ function.

Figure 5: Simulation results that show that the continuity property holds.

trajectory. Out of all samples we choose the ones with maximum deviation from the optimal value function at the optimal operation trajectory, i.e.

$$\tilde{\delta}_k(N, \varepsilon) := \max_j |V_N(k, \tilde{y}_L^*(t)) - V_N(k, y_{\varepsilon_i}^j)|. \quad (20)$$

For a sufficiently large number of samples this gives a good approximation of δ_k in a neighborhood of the optimal operation trajectory. Finally, we remark that of course we would have to check the conditions on δ_k for all time instances k . Because the optimal operation trajectory of the example exhibits periodic behavior (cf. Figure 3c) we could restrict ourselves to checking one period.

In the following we show exemplary results for a single time point $k = 50$ corresponding to the time $t = 0.5$. The chosen results are representative for all time points. The parameters from the above discussion were chosen as $\varepsilon_0 = 10^{-2}$, $n = 5$, $m = 10$, $N_i = 10i$, $i \in \{1, \dots, 10\}$.

Figure 5 shows the computed function $\tilde{\delta}_k$ as a function in its first and second component. The top figure 5a shows that $\tilde{\delta}_k(N, \varepsilon)$ is indeed bounded in N and thus it is possible to find a modulus of continuity φ_V that satisfies the required assumptions. Similarly, the bottom figure 5b demonstrates that the required upper bound for $\tilde{\delta}_k(N, \cdot)$ exists and satisfies the monotonicity assumptions. Thus we can conclude that, at least numerically, the continuity property holds for this example.

5 Conclusion

In this paper we have considered assumptions that are useful for proving convergence properties of MPC. In order to check whether the assumptions are fulfilled for a practical example, several simulations were carried out. The numerical results indicate that our assumptions are satisfied for a PDE example motivated by an application. This gives us confidence that they are the 'right' type of assumptions. The results thus motivate future investigations in which the assumptions will be shown analytically for simple systems.

References

- [1] M. S. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells. The fenics project version 1.5. *Archive of Numerical Software*, 3(100), 2015.
- [2] Z. Dong and D. Angeli. Analysis of economic model predictive control with terminal penalty functions on generalized optimal regimes of operation. *International Journal of Robust and Nonlinear Control*, 28(16):4790–4815, 2018.
- [3] T. Faulwasser, L. Grüne, and M. A. Müller. Economic nonlinear model predictive control. *Foundations and Trends in Systems and Control*, 5(1):1–98, January 2018.
- [4] D. Gale. On optimal development in a multi-sector economy. *Rev. Econ. Studies*, 34(1):1–18, 1967.

- [5] L. Grüne and J. Pannek. *Nonlinear Model Predictive Control. Theory and Algorithms*. Springer, second edition, 2017.
- [6] L. Grüne and S. Pirkelmann. Closed-loop performance analysis for economic model predictive control of time-varying systems. In , editor, *Proceedings of the 56th IEEE Conference on Decision and Control (CDC 2017)*, pages 5563–5569, Dec 2017.
- [7] L. Grüne and S. Pirkelmann. Economic model predictive control for time-varying system: Performance and stability results, July 2018.
- [8] M. A. Müller and L. Grüne. Economic model predictive control without terminal constraints for optimal periodic behavior. *Automatica*, 70:128–139, August 2016.
- [9] J. B. Rawlings and D. Q. Mayne. *Model predictive control: Theory and design*. 2009.
- [10] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, Mar 2006.