Influence of Discretization in Dynamically Embedded Model Predictive Control ⋆

Yaashia Gautam * **Marco M. Nicotra** *

* *University of Colorado Boulder, ECEE Department, Boulder, CO 80309 USA (e-mail: Yaashia.Gautam@colorado.edu)*

Abstract: The paper analyzes the closed-loop stability of Dynamically Embedded Model Predictive Control for input-constrained continuous-time nonlinear systems. Given a stabilizing continuous-time optimal control problem, the proposed method performs a discrete approximation to obtain a finite number of optimization variables. The resulting optimization problem is then embedded into a dynamic feedback law that evolves in parallel to the system. Using Input-to-State Stability, it is shown that the dynamic interconnection between the ideal continuous-time model predictive controller and the dynamically embedded solver is asymptotically stable for a sufficiently small discretization step and sufficiently fast solver dynamics. Numerical results, however, highlight a counter-intuitive behavior: as the discretization step decreases, the stability of the closed-loop system tends to deteriorate. This suggests that, although the discretization should be sufficiently accurate to correctly capture the behavior of the system, oversampling the system dynamics may be just as harmful as undersampling.

Keywords: Nonlinear predictive control, Numerical methods for optimal control, Real-time optimal control, Stability of Nonlinear Systems

1. INTRODUCTION

Model Predictive Control (MPC) is a popular control strategy for stabilizing multi-input nonlinear systems subject to constraints. As detailed in Mayne et al. (2000) and Goodwin et al. (2006), the idea behind MPC is to solve an Optimal Control Problem (OCP) at each time instant and apply the first step of the optimal control sequence as an input to the system. Due to the challenges of solving continuous-time OCPs, MPC is typically formulated in discrete time. For continuous-time systems, Magni and Scattolini (2004) introduced a sampled-data MPC formulation that relies on piecewise-constant control inputs to provide a sufficient amount of time for solving the OCP. Since solving OCPs to completion can be computationally prohibitive, the literature features a number of discrete-time MPC strategies that rely on a running estimate of the OCP solution, updated at each timestep. Notable examples include Suboptimal MPC by Scokaert et al. (1999), CGMRES by Ohtsuka (2004), RTI by Diehl et al. (2005), IPA-SQP by Ghaemi et al. (2009), and TDO by Liao-McPherson et al. (2020).

Interestingly enough, the same idea can be implemented in continuous time by replacing the discrete-time OCP solver with a continuous-time equivalent. The result, hereafter referred to as *Dynamically Embedded MPC* (DE-MPC), is a continuous-time MPC scheme that tracks the solution to the OCP with a bounded error. Several variations of DE-MPC have been proposed for the case of linear systems by Feller and Ebenbauer (2014), Nicotra et al. (2019), and Yoshida et al. (2019). Although the underlying ideas are similar, none of these papers study how discretizing the OCP affects the overall stability of the closed-loop system.

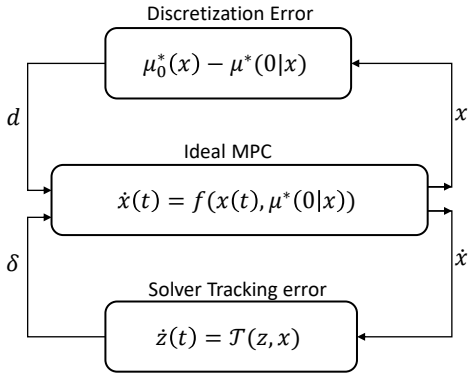

Fig. 1. DE-MPC can be interpreted as an ideal continuous-time MPC subject to two disturbances: the first is introduced by discretizing the OCP, the second is a result of not solving the OCP to completion.

This paper introduces a general DE-MPC formulation for nonlinear systems subject to input constraints. Given the traditional continuous-time MPC formulation as a starting point, the first step is to discretize the OCP so that there are a finite number of optimization variables. The discretized OCP is then solved in continuous-time using a dynamic feedback law that runs parallel to the controlled system. Using Input-to-State Stability (ISS), it is shown that, if the OCP discretization is sufficiently accurate and the OCP solver is sufficiently fast, the resulting closed-loop system is asymptotically stable. Numerical simulations validate the proposed method but highlight a counter-intuitive behavior: as the OCP discretization accuracy

⋆ This work is supported by the CMMI division of the National Science Foundation through the award number 2046212.

increases, the system becomes more difficult to stabilize. This behavior is studied in detail for a very simple linear model and is then shown to hold for a more general nonlinear problem.

The paper is organized as follows: Section 2 introduces the standard (ideal) continuous-time MPC formulation and details its main implementation challenges. Section 3 shows how to discretize a continuous-time OCP and studies the effects of discretization on the closed-loop system. Section 4 shows how the solution to the discretized OCP can be tracked with a bounded error using a dynamic feedback law and provides sufficient conditions for asymptotic stability of the closed-loop system in Figure 1. Section 5 features numerical validations of the proposed method while highlighting how the discretization step influences the stability of the closed-loop system.

2. PROBLEM STATEMENT

Consider a continuous-time system

$$\dot{x}(t) = f(x(t), u(t)) \tag{1}$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}^m$ is the control input, and $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is the system dynamics. The system is subject to input constraints $u \in \mathcal{U}$, where $\mathcal{U}$ is a closed, convex set.

To control the system using continuous-time MPC, we ideally need the solution to the optimal control problem

$$\begin{aligned} \min \quad & m(\xi(T)) + \int_0^T l(\xi(\tau), \mu(\tau))d\tau && \text{(2a)} \\ \text{s.t.} \quad & \dot{\xi}(\tau) = f(\xi(\tau), \mu(\tau)), \quad \xi(0) = x, && \text{(2b)} \\ & \mu(\tau) \in \mathcal{U}, \qquad\qquad \forall \tau \in [0, T], && \text{(2c)} \end{aligned}$$

where $T > 0$ is the prediction horizon, $m : \mathbb{R}^n \to \mathbb{R}$ is the terminal cost and $l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is the stage cost. Then, if we denote the solution mapping of (2) with $\mu^*(\tau|x)$, the continuous-time MPC can be implemented by assigning the feedback law $u(x) = \mu^*(0|x)$. As detailed in Feller and Ebenbauer (2014), state constraints can also be taken into account using barrier functions on the stage cost.

Assumption 1. The dynamics $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ are Lipschitz continuous, stabilizable, and admit the origin as a feasible equilibrium point. The functions $m : \mathbb{R}^n \to \mathbb{R}$ and $l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ are convex, twice continuously differentiable, zero at the origin and lower bounded by quadratic functions. Notably, there exists $\eta > 0$, such that $l(x, u) \geq \eta\|x\|^2$. In addition, there exists a scalar $\chi_0 > 0$ and a terminal control law $\kappa : \mathbb{R}^n \to \mathbb{R}^m$ such that, $\kappa(x) \in \mathcal{U},\ \forall \|x\| \leq \chi_0$ and

$$\nabla m(x) + l(x, \kappa(x)) \leq 0, \qquad \forall \|x\| \leq \chi_0.$$

Assumption 2. The parametrized optimal control problem (2) admits a solution mapping $\mu^*(t|x)$ that is strongly regular in x. Specific conditions for this to hold are detailed in Dontchev et al. (2019).

The following theorem states a well known result for continuous-time MPC.

Theorem 3. Under Assumptions 1 and 2, there exist positive scalars $\chi_1 > 0$ and $\Delta_1 > 0$ such that the origin of the closed loop system

$$\dot{x} = f(x, \mu^*(0|x) + d) \tag{3}$$

is ISS with state and input restrictions $\|x(0)\| \leq \chi_1$ and $\|d\|_\infty \leq \Delta_1$.

Proof. Given $d = 0$, it follows from Mayne et al. (2000) that the optimal control policy $\mu^*(\tau|x),\ t \in [0, T]$ and the resulting optimal trajectory

$$\dot{\xi}^*(\tau|x) = f(\xi^*(\tau|x), \mu^*(\tau|x)), \qquad \xi^*(0|x) = x, \tag{4}$$

are such that the optimal cost

$$V(x) = m(\xi^*(T|x)) + \int_0^T l(\xi^*(\tau|x), \mu^*(\tau|x))d\tau \tag{5}$$

satisfies $\dot{V}(x) \leq -\eta\|x\|^2,\ \forall \|x\| \leq \chi_0$. Since the unforced system is locally exponentially stable, it is also locally ISS when $d \neq 0$.

Unfortunately, the proposed continuous-time MPC formulation presents a few challenges. First, we need to solve the continuous-time optimal control problem (2). Second, we need to solve it "instantaneously" to have a continuous-time control law. To address the first challenge, we consider a discrete approximation of (2) that features a finite number of optimization variables. The second challenge will then be addressed by having a continuous-time solver run in parallel with the system, thereby tracking the OCP solution with a bounded error.

3. OPTIMAL CONTROL PROBLEM DISCRETIZATION

To discretize the OCP (2), we first define the discretization step $t_d = T/N$, where $N \in \mathbb{N}$ ensures that the prediction horizon T is an integer multiple of t_d. The system dynamics can then be discretized using the forward Euler approximation.

$$f_d(x_k, u_k) = x_k + t_d f(x_k, u_k). \tag{6}$$

Likewise, the integral portion of the cost function can be discretized using the rectangular approximation

$$\int_0^T l(\xi(\tau), \mu(\tau))dt \approx \sum_{k=0}^{N-1} t_d l(\xi_k, \mu_k), \tag{7}$$

thereby obtaining the discretized OCP

$$\begin{aligned} \min \quad & m(\xi(N)) + \sum_{k=0}^{N-1} t_d l(\xi_k, \mu_k) && \text{(8a)} \\ \text{s.t.} \quad & \xi_{k+1} = f_d(\xi_k, \mu_k), \quad \xi_0 = x, && \text{(8b)} \\ & \mu_k \in \mathcal{U}, \qquad\qquad \forall k \in [0, N-1], && \text{(8c)} \end{aligned}$$

Let $\mu_k^*(x)$ denote the solution mapping of (8). In general, the OCPs (2) and (8) are such that $\mu_k^*(x) \neq \mu^*(kt_d|x)$. Therefore, there is no guarantee that $u(x) = \mu_0^*(x)$ would be a stabilizing feedback law for system (1). To address this issue, the following lemma notes that the discretization error is proportional to t_d.

Lemma 4. There exist positive scalars $L > 0$ and $\chi_L > 0$ such that the discretization error $\Delta u_d = \mu_0^*(x) - \mu^*(0|x)$ satisfies

$$\|\Delta u_d\| \leq L t_d \|x\|, \qquad \forall\ \|x\| \leq \chi_L. \tag{9}$$

Proof. Following Assumption 2 and the definition of strong regularity, we have that, for a fixed t_d, there exists $c_1 > 0$

$$\|\mu^*(0|x)\| \leq c_1 \|x\|, \tag{10}$$

for x in a neighbourhood of the origin.
As detailed in Dontchev et al. (2019), there exists $c_2 > 0$ such that

$$\|\Delta u_d\| \leq c_2 t_d, \tag{11}$$

for a fixed x. Dontchev et al. (2019) also proves that the discretized OCP (8) is strongly regular and satisfies

$$\|\mu_0^*(x)\| \leq c_3 \|x\|, \tag{12}$$

with $c_3 > 0$. Using the triangular inequality, we then show

$$\|\mu_0^*(x) - \mu^*(0|x)\| \leq (c_2 + c_3)\|x\|. \tag{13}$$

Equation (9) follows from the fact that (11) holds for fixed t_d, (13) holds for fixed x, and both must hold simultaneously in a neighbourhood of the origin.

Remark 5. In principle, the results obtained in this paper are independent on the discretization strategy, meaning that it is possible to replace the forward Euler and rectangular approximations with any other numerical method as long as Lemma 4 remains satisfied.

Having established that the discrete approximation $\mu_0^*(x)$ belongs to a bounded neighbourhood of the ideal MPC law $\mu^*(0|x)$, the following theorem shows that the discretization retains similar ISS properties as Theorem 1.

Theorem 6. Given a sufficiently small discretization step t_d, there exist positive scalars $\chi_2 > 0$ and $\Delta_2 > 0$ such that the origin of the closed-loop system

$$\dot{x} = f(x, \mu_0^*(x) + \delta) \tag{14}$$

is ISS with state and input restrictions $\|x(0)\| \leq \chi_2$ and $\|\delta\|_\infty \leq \Delta_2$.

Proof. Due to Theorem 3, the system

$$\dot{x} = f(x, \mu^*(0|x) + \Delta u + \delta) \tag{15}$$

is ISS with restrictions $\|x(0)\| \leq \chi_1$ and $\|\Delta u + \delta\|_\infty \leq \Delta_1$. As such, given $\delta = 0$, there exists a class $\mathcal{K}$ function γ_1 such that

$$\overline{\lim_{t\to\infty}} \|x(t)\| \leq \gamma_1 \left(\overline{\lim_{t\to\infty}} \|\Delta u(t)\| \right), \quad \forall \|\Delta u(t)\| \leq \Delta_1, \tag{16}$$

where $\overline{\lim}_{t\to\infty}$ is limit superior. Following Lemma 4,

$$\overline{\lim_{t\to\infty}} \|\Delta u(t)\| \leq L t_d \overline{\lim_{t\to\infty}} \|x(t)\|, \quad \forall \|x(t)\| \leq \chi_L. \tag{17}$$

Under the specified conditions, it follows from the small gain theorem that, given a sufficiently small t_d satisfying

$$L t_d \gamma_1(s) < s, \qquad \forall s \leq \Delta_1, \tag{18}$$

the system is locally exponentially stable. As a result, it is locally ISS when $\delta \neq 0$.

Corollary 7. Under the same conditions as Theorem 6, there exists $\gamma_2 \in \mathcal{K}$ such that the closed-loop system (14) satisfies the input to output gain

$$\overline{\lim_{t\to\infty}} \|\dot{x}(t)\| \leq \gamma_2 \left(\overline{\lim_{t\to\infty}} \|\delta(t)\| \right). \tag{19}$$

Proof. Due to Assumption 1, $\dot{x} = f(x, \mu_0^*(x) + \delta)$ is bounded for any bounded state $\|x\|$ and input $\|\delta\|$. Thus, Input-to-State Stability is sufficient to ensure Input-to-Output Stability (IOS) when taking $\dot{x}$ as an output.

The main interest in Theorem 6 is that it shows that the continuous-time system can be stabilized by solving the discretized OCP (8) as opposed to the original OCP (2). However, continuous-time implementation would still require us to solve (8) instantaneously. This limitation is addressed in the following section.

4. DYNAMICALLY EMBEDDED SOLVER

The objective of this section is to drop the assumption of having to solve (8) instantaneously by using a continuous-time solver to track the solution as the OCP parameter $x(t)$ evolves over time. To this end, let $z^*(x) = [\mu_0^*(x)^T, \mu_1^*(x)^T, \dots, \mu_{N-1}^*(x)^T]^T$ be the exact solution to (8), and let the running estimate $z(t) = [\mu_0(t)^T, \mu_1(t)^T, \dots \mu_{N-1}(t)^T]^T$ be the internal state of a continuous-time solver in the form

$$\dot{z} = \mathcal{T}(z, x). \tag{20}$$

Rather than prescribing a specific solver, the following assumption provides sufficient conditions under which a generic solver can successfully run in parallel to the MPC. A notable class of solvers that satisfies this assumption is the projected gradient flow.

Assumption 8. Given a constant parameter x and a solution estimate z such that $\|z - z^*(x)\| \leq Z_s$, with $Z_s > 0$, the continuous-time solver (20) satisfies

$$(z - z^*(x))^T \mathcal{T}(z, x) \leq -\alpha\theta \|z - z^*(x)\|^\beta \tag{21}$$

where $\alpha > 0$ is a generic tunable parameter, whereas $\theta > 0$ and $\beta \geq 2$ are positive constants that depend on the specific choice of the solver.

Theorem 9. Under Assumptions 2 and 8, given a time-varying $x(t)$ satisfying $\|x\|_\infty \leq \chi_s$, with $\chi_s > 0$, there exists a scalar $\Delta_s > 0$ such that, the solver (20) is ISS with state and input restrictions $\|z(0) - z^*(x(0))\| \leq Z_s$ and $\|\dot{x}\|_\infty \leq \Delta_s$.

Proof. Consider the ISS-Lyapunov candidate function $W(z, x) = \frac{1}{2}\|z - z^*(x)\|^2$. Following Assumption 8, its time derivative satisfies

$$\dot{W} \leq -\alpha\theta \|z - z^*(x)\|^\beta + |\partial_x z^*(x)| \|z - z^*(x)\| \|\dot{x}\|, \tag{22}$$

where $\partial_x z^*(x)$ is the Clarke generalised Jacobian of the solution mapping $z^*(x)$. Due to Assumption 2, there exists a scalar $D_x^* > 0$ such that $|\partial_x z^*(x)| \leq D_x^*$, $\forall \|x\| \leq \chi_s$. Thus, the system satisfies the ISS-Lyapunov condition

$$\dot{W} < 0, \quad \forall \|z - z^*(x)\| > \alpha^{-1} \frac{D_x^*}{\theta} \|\dot{x}\|. \tag{23}$$

The input restrictions can then be obtained to ensure the condition $\|z(t) - z^*(x)\| \leq Z_s$, $\forall t \geq 0$. Namely, we have $\Delta_s = Z_s \alpha\theta / D_x^*$.

Corollary 10. Under the same conditions as Theorem 9, the continuous-time solver (20) satisfies the input to output gain

$$\overline{\lim_{t\to\infty}} \|\delta(t)\| \leq \alpha^{-1} \frac{D_x^*}{\theta} \overline{\lim_{t\to\infty}} \|\dot{x}(t)\|, \tag{24}$$

with $\delta(t) = \mu_0(t) - \mu_0^*(x(t))$.

Proof. The statement follows directly from the fact that $\delta(t)$ is the first element of the vector $z(t) - z^*(x(t))$.

Having shown that the continuous-time solver is able to track a time-varying solution with a bounded error, the following theorem addresses the stability of a dynamically embedded MPC where the solver runs in parallel to the system that is being controlled.

Theorem 11. Under Assumptions 1, 2 and 8, given a sufficiently small $t_d > 0$ satisfying Theorem 6, there exists a sufficiently large $\alpha > 0$ such that the dynamic

interconnection of (14) and (20), with $\delta(t) = \mu_0(t) - \mu_0^*(x(t))$, is Asymptotically Stable.

Proof. Theorems 6 and 9 ensure that the individual subsystems (14) and (20) are ISS with nonzero state and input restrictions. Moreover, it follows from Corollaries 7 and 10 that the two subsystems are IOS and satisfy the asymptotic gains (19) and (24). Therefore, it follows from the small gain theorem that, given a sufficiently large α such that

$$\alpha^{-1}\frac{D_x^*}{\theta}\gamma_2(s) < s, \qquad \forall s \leq \Delta_2, \tag{25}$$

the dynamic interconnection of the two subsystems is asymptotically stable with state and input restrictions $\|x(0)\| \leq \chi_f$, $\chi_f > 0$, and $\|z(0) - z^*(x(0))\| \leq Z_f$, $Z_f > 0$.

Theorem 9 states that, once the OCP has been discretized, it is possible to implement a dynamically embedded MPC by designing a sufficiently fast solver. However, the theorem does not provide insight on the relationship between t_d and α. Intuitively, one would expect that the discretization step t_d should be "as small as possible" to minimize the discretization error $\Delta u_d(x) = \mu^*(0|x) - \mu_0^*(x)$. As shown in the next section, however, this intuition may be inaccurate since the choice of t_d has a non-trivial impact on both $\gamma_2(\cdot)$ and θ. As a result, the dynamically embedded MPC may display counter-intuitive behaviors where, for a fixed α, the overall stability may benefit from *increasing* the discretization step t_d.

5. NUMERICAL ANALYSIS

5.1 Double Integrator: Discretization

To highlight and interpret the effects of discretization, we begin by studying the unconstrained linear time-invariant system

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 0.1 \end{bmatrix} u. \tag{26}$$

Given the initial condition $x(0) = [10, 0]^T$ and the control horizon $T = 30s$, consider the OCP

$$\min \quad \|\xi(T)\|_P^2 + \int_0^T \|\xi(t)\|_Q^2 + \|\mu(t)\|_R^2 \, dt \tag{27a}$$

$$s.t. \quad \dot{\xi} = A\xi + B\mu, \qquad \xi(0) = x(t), \tag{27b}$$

where $Q = I_2$, $R = 0.04$, and P is the solution to the Algebraic Riccati Equation $A^TP + PA - PBR^{-1}B^TP + Q = 0$.

As well known, the closed-form solution to this problem is the continuous Linear Quadratic Regulator (LQR) $\mu^*(0|x) = -Kx$, with $K = R^{-1}B^TP$. To study the effect of discretization, however, we now consider the discrete OCP approximation

$$\min \quad \|x_N\|_P^2 + \sum_{k=0}^{N-1} \|x_k\|_{Q_d}^2 + \|u_k\|_{R_d}^2 \tag{28a}$$

$$s.t. \quad x_{k+1} = A_d x_k + B_d u_k, \qquad x_0 = x, \tag{28b}$$

where $A_d = I_2 + t_d A$, $B_d = t_d B$, $Q_d = t_d Q$, $R_d = t_d R$, and $N = T/t_d \in \mathbb{N}$. Since P is the solution to the continuous Algebraic Riccati Equation, the solution to (28) is not the discrete LQR. Nevertheless, as detailed in Rawlings et al. (2017), the closed-form solution can still be obtained by propagating the Discrete Algebraic Riccati Equation

$$P_{k-1} = Q_d + A_d^T P_k A_d - A_d^T P_k B_d (R_d + B_d^T P_k B_d) B_d^T P_k A_d,$$

with $P_N = P$, and assigning $\mu_0^*(x) = -K_0 x$, where the optimal gain is $K_0 = (R_d + B_d^T P_0 B_d)^{-1} B_d^T P_0 A_d$.

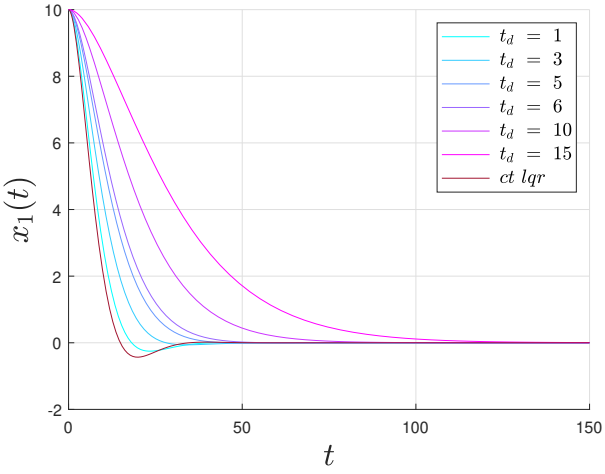

Fig. 2. *Closed-loop trajectory of the double integrator subject to the exact solution to* (28). *As t_d decreases, the system tends to the solution of the continuous-time LQR. Increasing t_d slows the closed-loop response.*

Having the closed-form expression of $\mu^*(0|x)$ and $\mu_0^*(x)$ enables us to study how the choice of t_d affects the behavior of the closed-loop system. Figure 2 shows the closed-loop response of the output $x_1(t) = [1\ 0]\, x(t)$, when subject to the control laws $u(x) = \mu^*(0|t)$ and $u(x) = \mu_0^*(x)$ for different discretization steps t_d.

As expected from Lemma 4, the error between the continuous-time optimal control policy and its discrete approximation tends to zero as $t_d \to 0$. More importantly, we note that the closed-loop response becomes slower as t_d increases. This behavior is likely due to the fact that the forward Euler approximation becomes increasingly unstable as t_d increases, thereby causing the discretized control law $\mu_0^*(x)$ to be *overdamped* compared to the continuous-time solution $\mu^*(0|x)$.

5.2 Double Integrator: Solver Dynamics

Instead of using the closed-form solutions, we now solve the discrete optimal control problem (28) using the preconditioned gradient flow solver

$$\dot{z} = -\frac{\alpha}{\text{cond}(H)} \nabla_z \left(\frac{1}{2} z^T H z + h^T z \right) \tag{29}$$

where H and h are the Hessian and the Jacobian of the cost function described in (28), defined as

$$H = 2(\hat{R} + \hat{B}^T \hat{Q} \hat{B}) \tag{30}$$

$$h = 2\hat{B}^T \hat{Q} \hat{A} x(0) \tag{31}$$

with

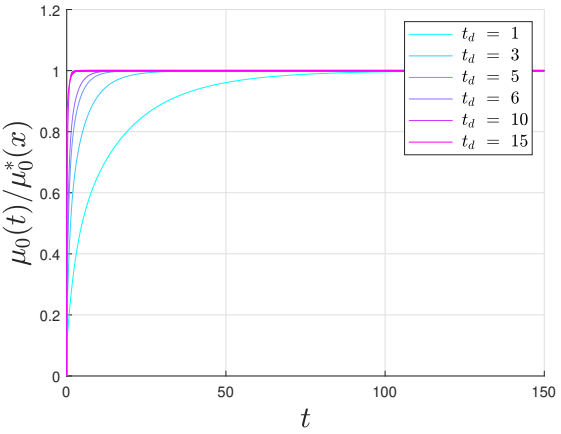

Fig. 3. *Convergence rate of* (29) *for a fixed value of* x. *Increasing* t_d *causes the solver to converge faster.*

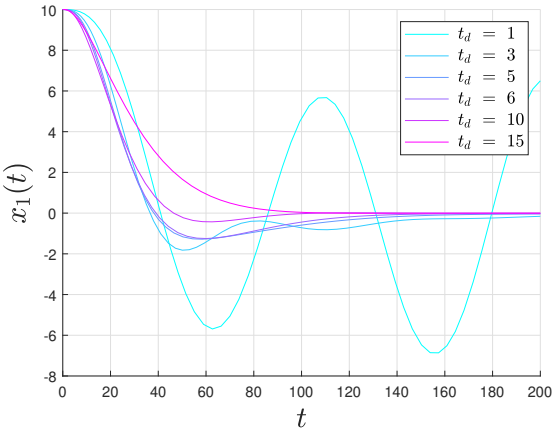

Fig. 4. *Closed-loop response of the DE-MPC for* $\alpha = 2e{-}3$. *The system is unstable for* $t_d = 1$, *but converges to the origin as* t_d *increases.*

$$\hat{B} = \begin{bmatrix} B & 0 & \dots & 0 \\ AB & B & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ A^{N-1}B & \dots & AB & B \end{bmatrix}, \qquad \hat{A} = \begin{bmatrix} A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \tag{32}$$

$$\hat{Q} = \begin{bmatrix} Q_d & & & \\ & \ddots & & \\ & & Q_d & \\ & & & P \end{bmatrix}, \qquad \hat{R} = \begin{bmatrix} R_d & & \\ & \ddots & \\ & & R_d \end{bmatrix} \tag{33}$$

Figure 3 shows the rate of convergence of (29) to the optimal input $\mu_0^*(x)$, assuming a fixed parameter x. Here, we note that higher values of t_d are associated with a higher convergence rate. This can be explained by noting that higher values of t_d entail lower $N = T/t_d$, thereby meaning that the optimization problem has fewer variables and is therefore easier to solve.

5.3 Double Integrator: Interconnection

When the solver evolves in parallel with the system, we know from Theorem 11 that, for a given discretization step t_d, there exists a sufficiently large α that stabilizes the closed-loop system. In Figure 4, we illustrate the behavior of the system for a fixed α and observe that the interconnection features better stability properties as t_d decreases.

This surprising result can be explained by considering Figures 2 and 3. Notably, we observe that higher values of t_d simultaneously slow down the dynamics of the closed-loop system and speed up the convergence rate of the solver. This combined effect makes it easier for the DE-MPC to track the solution of the discretized OCP as the discretization step increases. It is to be noted that varying the fixed parameter, α, for each t_d may help regain the stability of the closed-loop system. However, for a fixed α, the behavior is as shown in this paper.

5.4 Emergency Lane Change: Linearized Case

To further our study, we now consider the linearized dynamics of the emergency lane change maneuver found in Liao-McPherson et al. (2020), but without the dynamic extension, i.e. the control input u are the steering angles and not their derivatives. The steering angles are not a part of the state vector. The system is subject to input constraints $u_1 \in [-30°, 30°]$ and $u_2 \in [-6°, 6°]$. Given the initial conditions $x(0) = [5, 0, 0, 0]^T$ and a prediction horizon $T = .6s$, Figure 5 shows the effect of five different discretization steps on the closed-loop stability of the DE-MPC. The constrained OCP was solved using a projected primal-dual gradient flow, e.g., Bianchin et al. (2022).

Table 1. *Effect of discretization on the linearized lane change model.*

t_d	.06	.1	.12	.2	.3
N	30	18	15	9	6
$\max(\lvert eig(A_d)\rvert)$	24.5708	41.5986	50.1127	84.1691	126.7397
$\max(\lvert eig(A_d)\rvert^N)$	5e41	1e29	3e25	2e17	4e12

Once again, we verify that, for a fixed flow rate of the solver, lowering t_d can make the interconnection unstable. To provide more insight on this behavior, Table 1 reports the eigenvalues of the open-loop matrix A_d and its matrix exponential A_d^N, with $N = T/t_d$. The table shows that decreasing t_d causes A_d^N to have larger eigenvalues which, in turn, causes the optimal control problem to become ill-conditioned. The loss of stability is then explained by the fact that the dynamic solver is too slow compared to the plant, thereby causing the interconnection to become unstable. This behavior is somewhat similar to what is observed in Liao-McPherson et al. (2022) for the discrete-time case, where it was noted that, under appropriate circumstances, decreasing the prediction horizon length can be beneficial for closed-loop stability.

5.5 Emergency Lane Change: Nonlinear Case

We now consider the nonlinear model for the lane change maneuver, as described in the pervious section, where the dynamics can be found in Liao-McPherson et al. (2020). The dynamic feedback law is obtained by approximating (8) with a time-varying Quadratic Program and implementing a projected primal-dual gradient flow to track the OCP solution. Once again, we note that (given a

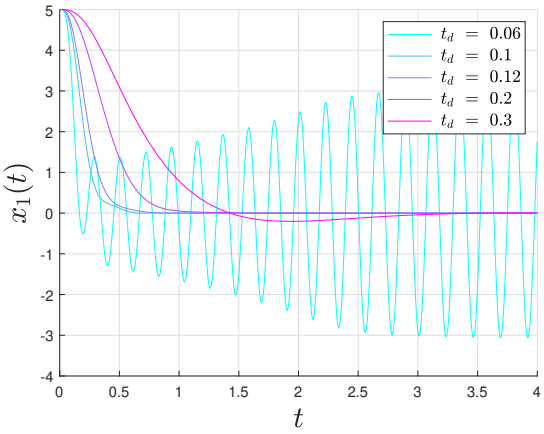

Fig. 5. *Lane change maneuver for the linearized system. Given a fixed solver, the response can be destabilized by selecting an excessively small discretization step.*

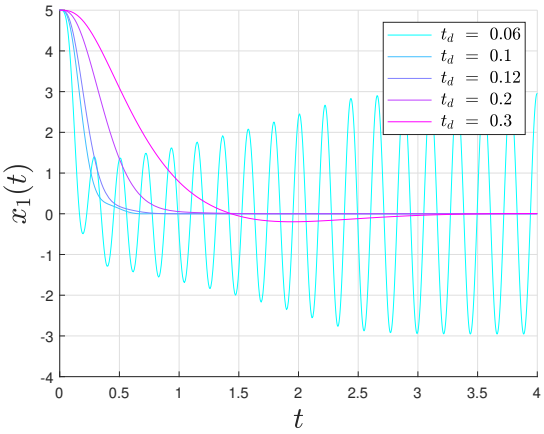

Fig. 6. *Lane change maneuver for the nonlinear model. Given a fixed solver, the response can be destabilized by selecting an excessively small discretization step.*

fixed solver) the closed-loop system becomes unstable as t_d decreases.

6. CONCLUSION

This paper analyzes the stability of Dynamically Embedded MPC applied to input-constrained continuous-time nonlinear systems. Rigorous proofs show that it is possible to mimic the behavior of continuous-time MPC by having a sufficiently accurate discrete approximation and tracking it with a sufficiently fast continuous-time solver. Numerical studies, however, highlight a counter-intuitive behavior: for a given solver, the natural instinct of choosing the discretization step "as small as possible" can actually destabilize the closed-loop system. Notwithstanding, the paper should not be interpreted as an invitation to choose an "arbitrarily large" discretization step (which can also cause instability). Rather, the message is that once the discretization is accurate enough to capture the system dynamics, any further reductions may be counter-productive to the stability of DE-MPC when taking into account the solver dynamics.

REFERENCES

Bianchin, G., Cortés, J., Poveda, J.I., and Dall'Anese, E. (2022). Time-varying optimization of LTI systems via projected primal-dual gradient flows. *IEEE Transactions on Control of Network Systems*, 9, 474–486.

Diehl, M., Findeisen, R., Allgöwer, F., Bock, H.G., and Schlöder, J.P. (2005). Nominal stability of real-time iteration scheme for nonlinear model predictive control. *IEE Proceedings-Control Theory and Applications*, 152(3), 296–308.

Dontchev, A.L., Kolmanovsky, I.V., Krastanov, M.I., Nicotra, M.M., and Veliov, V.M. (2019). Lipschitz stability in discretized optimal control with application to SQP. *SIAM Journal on Control and Optimization*, 57(1), 468–489.

Feller, C. and Ebenbauer, C. (2014). Continuous-time linear MPC algorithms based on relaxed logarithmic barrier functions. *IFAC Proceedings Volumes*, 47(3), 2481–2488.

Ghaemi, R., Sun, J., and Kolmanovsky, I.V. (2009). An integrated perturbation analysis and sequential quadratic programming approach for model predictive control. *Automatica*, 45(10), 2412–2418.

Goodwin, G., Seron, M.M., and De Doná, J.A. (2006). *Constrained control and estimation: an optimisation approach*. Springer Science & Business Media.

Liao-McPherson, D., Nicotra, M.M., and Kolmanovsky, I.V. (2020). Time-distributed optimization for real-time model predictive control: Stability, robustness, and constraint satisfaction. *Automatica*, 117, 108973.

Liao-McPherson, D., Skibik, T., Leung, J., Kolmanovsky, I.V., and Nicotra, M.M. (2022). An analysis of closed-loop stability for linear model predictive control based on time-distributed optimization. *IEEE Transactions on Automatic Control*, 67, 2618–2625.

Magni, L. and Scattolini, R. (2004). Model predictive control of continuous-time nonlinear systems with piecewise constant control. *IEEE Transactions on Automatic Control*, 49, 900–906.

Mayne, D.Q., Rawlings, J.B., Rao, C.V., and Scokaert, P.O.M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 789–814.

Nicotra, M.M., Liao-McPherson, D., and Kolmanovsky, I.V. (2019). Embedding constrained model predictive control in a continuous-time dynamic feedback. *IEEE Transactions on Automatic Control*, 64(5), 1932–1946. doi:10.1109/TAC.2018.2867359.

Ohtsuka, T. (2004). A continuation/GMRES method for fast computation of nonlinear receding horizon control. *Automatica*, 40(4), 563–574.

Rawlings, J.B., Mayne, D.Q., and Diehl, M.M. (2017). *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing.

Scokaert, P.O.M., Mayne, D.Q., and Rawlings, J.B. (1999). Suboptimal model predictive control (feasibility implies stability). *IEEE Transactions on Automatic Control*, 44(3), 648–654.

Yoshida, K., Inoue, M., and Hatanaka, T. (2019). Instant mpc for linear systems and dissipativity-based stability analysis. *IEEE Control Systems Letters*, 3(4), 811–816.