Towards continuous-time MPC: a novel trajectory optimization algorithm

Souvik Das[1], Siddhartha Ganguly[1], Muthyala Anjali[2], and Debasish Chatterjee[1]

Abstract— This article introduces a numerical algorithm that serves as a preliminary step toward solving continuous-time model predictive control (MPC) problems directly *without* explicit time-discretization. The chief ingredients of the underlying optimal control problem (OCP) are a linear time-invariant system, quadratic instantaneous and terminal cost functions, and convex path constraints. The thrust of the method involves finitely parameterizing the admissible space of control trajectories and solving the OCP satisfying the given constraints at every time instant in a tractable manner without *explicit* time-discretization. The ensuing OCP turns out to be a convex semi-infinite program (SIP), and some recently developed results are employed to obtain an optimal solution to this convex SIP. A numerical illustration on a benchmark model is included to show the efficacy of the algorithm.

Index Terms— Model predictive control, numerical optimal control, convex semi-infinite programs

I. INTRODUCTION

This article is concerned with an important component in the implementation of model predictive control (MPC) of continuous-time linear plants directly in continuous time. Of course, since digital implementation of controllers is normative today, there is an implicit discretization of time associated with the digital implementation, and this discretization due to digital implementation *continues to stay* in our framework. Consequent to the preceding reconciliation, one issue must be immediately addressed in the context of a continuous-time implementation of MPC — the time interval needed to compute solutions to the underlying finite-horizon constrained optimal control problem. This latency may be substantial due to the complexity of the optimal control problem, and is typically much longer than the interval of discretization associated with digital implementation; therefore, an *online* implementation of this scheme may not be suitable.

To circumvent this difficulty, we establish a novel tractable numerical algorithm for numerically solving the finite-horizon constrained optimal control problem underlying an MPC strategy, *directly in continuous-time*. This algorithm is fine-tuned to handle the most popular class of problem data

with linear dynamics, quadratic cost functions, and affine constraints.

With this numerical algorithm at hand, our goal, to be realized over our subsequent works, is to leverage the *explicit MPC* technology: the implicit feedback in the MPC scheme is computed at a set of points on the admissible set of states, and an interpolation mechanism (guaranteeing tight control of the ensuing error) constructs an *explicit* feedback by means of *offline* numerical routines along the lines of [7]. The result is a *feedback* that is 'close' to the implicit feedback map in the continuous-time MPC, and online implementation merely consists of evaluating this feedback map at each instant of time. Moreover, the extent of 'closeness' of the interpolated feedback map and the implicit feedback of the MPC scheme is a choice of the designer. In particular, as will be demonstrated in subsequent reports, the ability to pre-assign an error margin lets us embed this uncertainty directly into the optimal control problem underlying the MPC strategy; no issues related to loss of feasibility arise as a consequence.

Contributions

It is evident from the preceding discussion that numerically computing solutions to the finite-horizon constrained optimal control/trajectory optimization problem is the critical issue in the context of explicit continuous-time MPC. Indeed, since an explicit map of the underlying feedback is our eventual target (and its synthesis strategy will be reported elsewhere), this offline computation at specific points in the feasible set constitutes the key step at hand, and we shall execute this step in this article. Here are our key contributions:

1) We introduce a new algorithm to solve a class of finite horizon continuous-time optimal control problems that serve as a first step towards addressing a continuous-time MPC without *explicit* time discretization. The established algorithm instead treats the admissible space of control functions corresponding to the underlying continuous-time OCP to be finite-dimensional (see Definition 2.1 in §II and the subsequent discussions).

2) In our approach the uncountable family of convex constraints in the continuous-time OCP is left intact, which makes the problem challenging but arguably more accurate in comparison to conventional methods such as *direct multiple shooting*, *direct collocation*, etc., that rely on time-discretization. We direct the readers to Remark (3.4) for a detailed discussion.

3) There have been previous efforts to address the finite-horizon optimal control problem in MPC directly in continuous-time; we draw attention to [11], [14] and the references therein, where a discretization scheme of the control input using piecewise linear functions is adopted *on a finite time-grid* to ensure that the constraints are satisfied at the grid-points. In comparison, an internal feature of our algorithm intelligently selects a predefined number of time points so that the satisfaction of the constraints at those sample points *implies* that the constraints hold for all time; certain convexity structures are central to this step that converts an infinitary condition into a finitary one. Moreover, the value function and the optimizers of the finite horizon optimal control problem (more precisely a reformulated version of the problem, see (FOCP)) is equal to the value function of the recast semi-infinite optimization problem; see (11), Theorem 3.1, and the discussion thereafter.

Notation

We let $\mathbb{N}^* := \{1, 2, \ldots\}$ denote the set of positive integers, $\mathbb{N} := \mathbb{N}^* \cup \{0\}$. The vector space $\mathbb{R}^d$ is equipped with standard inner product, $\langle x, y \rangle := \sum_{j=1}^{d} x_j y_j$ for every $x, y \in \mathbb{R}^d$. Let X and Y be subsets of some finite dimensional Euclidean spaces, then space of continuous functions on a domain X taking values in Y is denoted by $\mathcal{C}(X; Y)$ and by $\mathbb{L}^\infty(X; Y)$ we denote the space of essentially bounded function with the essential supremum norm.

II. PROBLEM STATEMENT

Let us consider a linear time-invariant controlled dynamical system, modeled by the ordinary differential equation

$$\dot{x}(t) = Ax(t) + Bu(t) \text{ for a.e. } t \in [0, T], \quad (1)$$

where $T > 0$ is a fixed time horizon, $x(t) \in \mathbb{R}^d$ is the vector of states and $u(t) \in \mathbb{R}^m$ is the vector of control/actions at time t, and $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times m}$ are the state and the actuation matrices, respectively. Assume that the initial state $x(0) := x \in \mathbb{R}^d$ and that the nonempty final set $\mathbb{X}_F \subset \mathbb{R}^d$ are given, i.e.,

$$x(T) \in \mathbb{X}_F \subset \mathbb{R}^d, \quad (2)$$

and that the control trajectory satisfies $u(\cdot) \in \mathcal{U}$, where

$$\mathcal{U} := \{u(\cdot) \in \mathbb{L}^\infty([0, T]; \mathbb{U}) \mid u(t) \in \mathbb{U} \text{ for a.e } t \in [0, T]\}. \quad (3)$$

Here $\mathbb{U} := \prod_{i=1}^{m} \mathbb{U}_i$, and the sets $\mathbb{U}_i$, for each $i = 1, \ldots, m$, are given nonempty, connected, and compact intervals with nonempty interiors. A control $u(\cdot)$ is *feasible* if it is measurable,[1] satisfies the control constraint (3), and the corresponding solution $x(\cdot)$ of (1) satisfies (2). We define the objective function by

$$V_T(x, u(\cdot)) := c_F(x(T)) + \int_0^T c(x(t), u(t)) \, dt. \quad (4)$$

---

[1] For us the word 'measurability' always refers to Lebesgue measurability, and 'a.e.' refers to almost everywhere relative to the Lebesgue measure.

Over the feasible controls $u(\cdot) \in \mathcal{U}$, the finite horizon continuous-time optimal control problem is typically posed:

$$\inf_{u(\cdot)} \quad V_T(x, u(\cdot))$$

$$\text{sbj. to} \quad \begin{cases} \text{dynamics } (1) \text{ and its associated data,} \\ x(0) = x, \ x(T) \in \mathbb{X}_F \subset \mathbb{R}^d, \\ x(t) \in \mathbb{X} \quad \text{for each } t \in [0, T], \\ u(\cdot) \in \mathcal{U}. \end{cases} \quad \text{(OCP)}$$

Tractability of (OCP)

Notice that the optimal control problem (OCP) is an *infinite dimensional* optimization problem due to the fact that $u(\cdot) \in \mathcal{U}$ and $\mathcal{U}$ is infinite-dimensional. In contrast, (OCP) admits a finite-dimensional avatar when minimized over admissible control trajectories that are *finitely* parameterized (e.g., if $\mathcal{U}$ is the linear span of certain 'basis functions'), which leads to a narrower set of admissible controls compared to the general setting of (OCP); see [13]. This motivates the following definition:

Definition 2.1: Let $N \in \mathbb{N}^*$ and $\mathcal{D} := (\psi_i)_{i \in \mathbb{N}} \subset \mathcal{C}([0, T]; \mathbb{R})$ be a dictionary of bounded continuous functions. We define the discretized set of admissible controls $\mathcal{U}_\mathcal{D} \subset \mathcal{C}([0, T]; \mathbb{R})$ by

$$\mathcal{U}_\mathcal{D} := \text{span}\{\psi_i : [0, T] \to \mathbb{R} \mid \text{ for } i = 1, \ldots, N\},$$

which is the span of finitely many basis functions from the dictionary $\mathcal{D}$.

Assumption 2.2 (Standing Assumption): We stipulate that an N-tuple of linearly independent elements from dictionary $\mathcal{D}$ have been extracted and fixed.

For each $i \in \{1, \ldots, m\}$, the i-th component $t \mapsto u_i(t)$ of the control trajectory $t \mapsto u(t)$ is permitted to be a linear combination of basis functions $t \mapsto \psi_j(t)$ for $j = 1, \ldots, N$. Define $\Psi(t) := (\psi_1(t) \ \psi_2(t) \ \ldots \ \psi_N(t))$. Then we have the component-wise parameterization:

$$[0, T] \ni t \mapsto u_i^\mathcal{D}(t) = \sum_{j=1}^{N} \alpha_{i,j} \psi_j(t) =: \langle \alpha_i, \Psi(t) \rangle \quad (5)$$

for $i = 1, \ldots, m$, where $\alpha_{i,j} \in \mathbb{R}$ are the *control coefficients* (to be determined). Consequently, one can write the control trajectory $t \mapsto u^\mathcal{D}(t)$ in the compact form

$$u^\mathcal{D}(t) = \begin{pmatrix} \alpha_{1,1} & \alpha_{1,2} & \cdots & \alpha_{1,N} \\ \alpha_{2,1} & \alpha_{2,2} & \cdots & \alpha_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{m,1} & \alpha_{m,2} & \cdots & \alpha_{m,N} \end{pmatrix} \Psi(t) =: \alpha \Psi(t),$$

$$(6)$$

for each $t \in [0, T]$, where $\alpha \in \mathbb{R}^{m \times N}$ is the coefficient matrix. Note that with the above parameterization, the control trajectory $t \mapsto u^\mathcal{D}(t) \in \mathbb{U}$ and $u^\mathcal{D}(\cdot) \in \mathcal{U}$ — this is because each component $t \mapsto u_i^\mathcal{D}(t) \in \mathbb{U}_i$ and $u_i^\mathcal{D}(\cdot) \in \mathcal{C}([0, T]; \mathbb{R})$ where $i = 1, \ldots, m$. The choice of the generating functions $(\psi_j)_{j=1}^{N}$ depends primarily on the type of applications one has in mind and is up to the designer.

The finite-dimensional optimal control problem over the set of feasible controls taking values from $\mathcal{U}_{\mathcal{D}}$, is given by:

$$\min_{(u_i^{\mathcal{D}}(\cdot))_{i=1}^m \subset \mathcal{U}_{\mathcal{D}}} \quad V_T\big(x, u^{\mathcal{D}}(\cdot)\big)$$

$$\text{sbj. to} \quad \begin{cases} \text{dynamics (1) and its associated data,} \\ x(0) = x, \ x(T) \in \mathbb{X}_F \subset \mathbb{R}^d, \\ x(t) \in \mathbb{X} \quad \text{for each } t \in [0, T], \\ u^{\mathcal{D}}(t) \in \mathbb{U} \quad \text{for each } t \in [0, T], \end{cases}$$

(FOCP)

with the following data:

(FOCP-a) a quadratic *instantaneous cost*

$$(\xi, \mu) \mapsto c(\xi, \mu) := \langle \xi, Q\xi \rangle + \langle \mu, R\mu \rangle \in [0, +\infty[$$

and a quadratic *terminal cost*

$$\xi \mapsto c_F(\xi) := \langle \xi, P\xi \rangle \in [0, +\infty[$$

with given positive semi-definite matrices $Q = Q^\top \in \mathbb{R}^{d \times d}$ and $P = P^\top \in \mathbb{R}^{d \times d}$, and a given positive definite matrix $R = R^\top \in \mathbb{R}^{m \times m}$.

(FOCP-b) The state constraint set $\mathbb{X} \subset \mathbb{R}^d$ is a closed and convex set, and the terminal constraint set $\mathbb{X}_F \subset \mathbb{X} \subset \mathbb{R}^d$ is a compact and convex set. Each set $\mathbb{X}$ and $\mathbb{X}_F$ contains $0 \in \mathbb{R}^d$ in their respective interiors.

We denote the value of the objective in (FOCP) at initial state x and admissible control $u^{\mathcal{D}}(\cdot)$ by $V_T(x, u^{\mathcal{D}}(\cdot))$, and

$$V_T^*(x) = \text{the optimal value of (FOCP);} \qquad (7)$$

since the initial states enter (FOCP) parametrically, of course $V_T^* : \mathbb{X}_T \longrightarrow [0, +\infty[$ depends on the parameter x where $\mathbb{X}_T$ denotes the set of initial states for which the problem (FOCP) is feasible.

Assumption 2.3: We stipulate that Slater's condition holds for (FOCP).[2]

Assumption 2.4: We enforce the following assumptions:

○ In the context of the optimal control problem (FOCP) with its associated data (FOCP-a)–(FOCP-b), we assume that $\mathbb{X}_T \neq \varnothing$, and that the problem (FOCP) admits an *optimal state-action trajectory* $[0, T] \ni t \mapsto \big(x_*(t; x), u_*(t; x)\big)$;

○ the state-action trajectory $[0, T] \ni t \mapsto \big(x_*(t; x), u_*(t; x)\big)$ is a normal extremal in sense of [6, Definition 2.1].

Proposition 2.5: Consider the optimal control problem (FOCP) with its associated data (FOCP-a)–(FOCP-b). Let Assumption (2.3) and Assumption (2.4) hold, and let the constraint sets $\mathbb{X}, \mathbb{X}_F$, and $\mathbb{U}$ are polytopic. Then $[0, T] \ni t \mapsto u_*(t; x)$ is Lipschitz continuous for every $x \in \mathbb{X}_T$.

Proof: Notice that the assumptions (H1)–(H4) in [6] are satisfied. The proof follows immediately by invoking [6, Theorem 3.1]. ∎

Despite its finite-dimensional nature, the optimal control problem (FOCP) in a computation-theoretic sense is NP-hard

[2]Assumption 2.3 states that the (11) is strictly feasible and is a technical requirement for the proof of Theorem 3.1 to go through.

in general, and involves uncountably many set-membership constraints since the state and the action variable in (FOCP) are subjected to the constraints *at each instant* $t \in [0, T]$.

III. Solving (FOCP) via a convex SIP

The finite-dimensional problem (FOCP) is a convex program and admits a tractable approximation algorithm up to arbitrary precision, which is the subject of §III. To this end, we reformulate (FOCP) in the language of a convex semi-infinite program. Recall that the discretized control trajectory is given by the expression $t \mapsto u^{\mathcal{D}}(t) = \alpha\Psi(t) \in \mathbb{U}$ as given in (6). Let us define the set of admissible α's by

$$\mathcal{P} := \big\{\alpha \mid x(t) \in \mathbb{X}, \ \alpha\Psi(t) \in \mathbb{U} \text{ for each } t \in [0, T]\big\}.$$

The solution of (1) starting from an initial state x is given by the variation of constants formula:

$$t \mapsto x(t) = e^{At}x + \int_0^t e^{A(t-\tau)} Bu^{\mathcal{D}}(\tau)\, \mathrm{d}\tau$$

$$= e^{At}x + \int_0^t e^{A(t-\tau)} B\alpha\Psi(\tau)\, \mathrm{d}\tau. \qquad (8)$$

Let us fix the notation

$$\overline{N} := mN. \qquad (9)$$

Then (FOCP) can be recast as a convex semi-infinite program:

$$\min_{\alpha \in \mathcal{P}} \quad V_T\big(x, \alpha\Psi(\cdot)\big)$$

$$\text{sbj. to} \quad \begin{cases} x(0) = x, \ x(T) \in \mathbb{X}_F \subset \mathbb{R}^d, \\ x(t) \in \mathbb{X} \quad \text{for each } t \in [0, T], \\ \alpha\Psi(t) \in \mathbb{U} \quad \text{for each } t \in [0, T], \end{cases} \quad (10)$$

for each $x \in \mathbb{X}_T$. Our main focus will now be on solving the problem (10). To this end, define the function $[0, T]^{\overline{N}} \ni (t_1, \ldots, t_{\overline{N}}) =: \mathcal{T} \mapsto \mathcal{G}(\mathcal{T}; x) \in \mathbb{R}$ for every $x \in \mathbb{X}_T$, by

$$\mathcal{G}(\mathcal{T}; x) \quad :=$$

$$\min_{\alpha \in \mathcal{P}} \quad V_T\big(x, \alpha\Psi(\cdot)\big)$$

$$\text{sbj. to} \quad \begin{cases} x(0) = x, \ x(T) \in \mathbb{X}_F \subset \mathbb{R}^d, \\ x(t_i) \in \mathbb{X} \quad \text{for all } i = 1, \ldots, \overline{N}, \\ \alpha\Psi(t_i) \in \mathbb{U} \quad \text{for all } i = 1, \ldots, \overline{N}. \end{cases} \quad (11)$$

In the light of Assumption 2.4, the 'min' in (10) and (11) are well-defined, which leads to the following critical observation:

Theorem 3.1: Consider the continuous-time optimal control problem (FOCP) along with its associated data (FOCP-a)-(FOCP-b). Consider also (10) and (11) along with their associated data and notations. Suppose that Assumption 2.3 holds. If the finite sequence $\mathcal{T}^* := (t_1^*, \ldots, t_{\overline{N}}^*)$ is a solution of the maximization problem

$$\sup_{\mathcal{T} \in [0, T]^{\overline{N}}} \mathcal{G}(\mathcal{T}; x) \quad \text{for } x \in \mathbb{X}_T, \qquad (12)$$

then $V_T^*(x) = \mathcal{G}(\mathcal{T}^*; x)$, where $V_T^*(\cdot)$ is defined in (7).

Remark 3.2: Let us discuss the assertion of Theorem 3.1. The continuous-time optimal control problem (11) has been

formulated as a finite-dimensional optimization problem with a finite number of path constraints induced by the sequence $\mathcal{T}^*$. The above proposition asserts that it is sufficient to consider only (but intelligently picked) $\overline{N}$-many discrete points $\mathcal{T}$ at which the path constraints must be satisfied for the purpose of solving (10). To recover such an $\overline{N}$-tuple of optimal point, the maximization problem (12) in Theorem 3.1 must be solved globally on $[0,T]^{\overline{N}}$.

Proof: [Proof of Theorem 3.1] Observe that the set $\mathcal{P}$ is a compact and convex subset of $\mathbb{R}^{m \times N}$, in view of the fact that the sets $\mathbb{X}$ and $\mathbb{U}$ are compact and convex. This implies that feasible set

$$\mathcal{S} := \big\{ \alpha \in \mathcal{P} \mid \text{the constraints in (10) hold} \big\}$$

corresponding to (10) is closed and convex. Define the map $\mathbb{R}^{m \times N} \ni \alpha \mapsto F(\alpha) := \alpha \Psi(t) \in \mathbb{U}$. Note that $F(\cdot)$ is continuous in α and consequently any open set around the origin in $\mathbb{U}$ has a preimage in $\mathcal{P}$ which is open. Similar argument follows for the constraint $x(t) \in \mathbb{X}$, implying that $\mathcal{P}$ has a nonempty interior in the light of Assumption 2.3. Moreover, the cumulative cost function

$$\mathcal{P} \ni \alpha \mapsto \int_0^T \big(\langle x(t), \, Qx(t) \rangle + \langle \alpha \Psi(t), \, R\alpha\Psi(t) \rangle \big) \, \mathrm{d}t,$$

is both continuous and convex in α. The existence of an optimizer $\mathcal{T}^* := (t_1^*, t_2^*, \ldots, t_{\overline{N}}^*) \in [0,T]^{\overline{N}}$ solving (12) then follows from Assumption 2.3 and [1, Theorem 4.1]. Invoking [3, Theorem 1], we assert that $V_T^*(x) = \mathcal{G}(\mathcal{T}^*; x)$ for each $x \in \mathbb{X}_T$. ∎

Remark 3.3: Note that the solution to the maximization problem in Theorem 3.1 is equal to the solution to the convex SIP (10) by a technique established in [15].

Algorithm to solve (FOCP)

Here, we present an algorithm to solve the continuous-time optimal control problem directly. Recall that the quantity

$$\alpha \in \underset{\zeta \in \mathcal{P}}{\arg\min} \big\{ V_T(x; \zeta\Psi(\cdot)) \,\big|\, x(t_i) \in \mathbb{X}, \, \zeta\Psi(t_i) \in \mathbb{U},$$
$$\text{for all } i = 1, \ldots, \overline{N} \big\}, \qquad (13)$$

where $V_T\big(x; \alpha\Psi(\cdot)\big)$ is defined in (4).

Remark 3.4: Contemporary methods for direct trajectory optimization involve discretization of both time and the space of admissible control functions via employing some form of approximation at the level of control [5], [8] or both states and control [9]. These methods are broadly classifiable into two categories — *direct multiple shooting* and *direct collocation* — that involve discretization of the time axis, and discretization of the dynamics by means of suitable numerical integrator schemes, adding finitely many equality constraints to the nonlinear program, one at each point of discretization. These schemes produce a finitely parametrized family of trajectories that are then subjected to finitely many set-membership constraints, one each at the points of discretization: the uncountable family of path constraints are relaxed to finitely many constraints on the discretized time grid. The resulting finite-dimensional nonlinear program is

---

Algorithm 1: Simulated annealing based algorithm for continuous-time optimal control problem

Data : threshold number of iterations τ;
Initialize: initialize constraint points:
$t_1^0, t_2^0, \ldots, t_{\overline{N}}^0 \in [0,T]$; initial guess for maximum value $\mathcal{G}_{\max}$; initial guess for the initial solution $\overline{\alpha}$;

1 while $n \leqslant \tau$ **do**
2   *Sample* (via simulated annealing-based algorithm) constraint time indices:
     $\mathcal{T}^n = (t_1^n, t_2^n, \ldots, t_{\overline{N}}^n) \in [0,T]^{\overline{N}}$
3   *Evaluate* $\mathcal{G}_n := \mathcal{G}(\mathcal{T}^n; x)$ as defined in (11)
4   *Recover* the solution α_n as given in (13)
5   **if** $\mathcal{G}_n \geqslant \mathcal{G}_{\max}$ **then**
6   Set $\mathcal{G}_{\max} \leftarrow \mathcal{G}_n$
7   Set $\overline{\alpha} \leftarrow \alpha_n$
8   Update $n \leftarrow n + 1$
9 end

---

then solved. On the other hand, the uncountable family of constraints in (OCP) is left intact in our approach, which makes it far more challenging but consequently and arguably, more accurate compared to conventional direct methods.

IV. STABILITY

Recall from §I that we seek to leverage the explicit MPC technology along the lines of [7]. To this end we include a technical result concerning the stability of the dynamical system (1) under the feedback policy $\xi \mapsto u_*^{\mathrm{o}}(\xi) := u_*(0; \xi)$.

Assumption 4.1: We stipulate that $c_F(\cdot), \mathbb{X}_F$, and $c(\cdot, \cdot)$ must satisfy following properties [12]:

○ There exists a feedback given by $\mathbb{X}_F \ni \xi \mapsto g_F(\xi) \in \mathbb{U}$, such that the terminal set $\mathbb{X}_F$ is positively invariant for the system $\dot{\xi} = A\xi + Bg_F(\xi)$.
○ The following inequality holds for all $\xi \in \mathbb{X}_F$:

$$\langle \nabla_\xi c_F(\xi), \, A\xi + Bg_F(\xi) \rangle \leqslant -c(\xi, g_F(\xi)). \qquad (14)$$

Here is our key technical result:

Theorem 4.2: Consider the constrained optimal control problem (FOCP) along with its data (FOCP-a)–(FOCP-b), and suppose that Assumption (4.1) holds. Let the state and action constraint sets be polytopic. Assume that $\mathbb{X}_T \neq \varnothing$, and $[0,T] \ni t \mapsto \big(x_*(t; x), u_*(t; x) \big)$ be an optimal state-action trajectory. Define $\mathbb{X}_T \ni \xi \mapsto u_*^{\mathrm{o}}(\xi) := u_*(0; \xi)$. Then we have the following inequality for all $\xi \in \mathbb{X}_T$:

$$\langle \nabla_\xi V_T^*(\xi), \, A\xi + Bu_*^{\mathrm{o}}(\xi) \rangle \leqslant -c(\xi, u_*^{\mathrm{o}}(\xi)). \qquad (15)$$

Remark 4.3: Theorem (4.2) demonstrates that *under the feedback policy* $\mathbb{X}_T \ni \xi \mapsto u_*^{\mathrm{o}}(\xi) := u_*(0; \xi)$ the closed-loop system $\dot{x}(t) = Ax(t) + Bu_*^{\mathrm{o}}(x)$ is stable via the descent-like property (15) of the value function $V_T^*(\cdot)$ along the closed-loop trajectories of the system (1). The complete proof can be found in [4].

V. DISCUSSION AND NUMERICAL EXPERIMENT

In this section, with the aid of a numerical experiment, we demonstrate the ability of our algorithm to generate constrained state and control trajectories that satisfy the given constraints for *all* time. We again emphasize the fact that unlike conventional methods we directly deal with uncountably many path constraints without any time discretization. The procedure and the tractable algorithm have been discussed in detail above; see §III. We also highlight the presence of a *turnpike behavior* [10] in the example in §V-A. This phenomenon serves as a motivation to employ model predictive control strategies, which we will report in our subsequent articles. We report that all numerical experiments are conducted in a computer with Intel(R) Core(TM) $i5-5200U$ CPU clocked at 2.20GHz with 8GB RAM.

A. Stabilization of an inverted pendulum on a cart

Here we consider an inverted pendulum on a movable cart with a restricted cart length as shown in Fig. 1. This is a benchmark under-actuated system with the horizontal force acting on the cart as the only control variable; see [2, §4.3] for more details. The inverted pendulum is in an upright position pivoted to the cart having an unstable equilibrium point. Our objective is to stabilize the inverted pendulum on the cart for a given time $T > 0$. The cart-pendulum system is a nonlinear time-invariant fourth-order dynamical system and we consider a linear approximation around the vertical position of the inverted pendulum.

Linear system formulation:

We linearize the dynamics of the inverted pendulum system around the origin with the states $x_1(t) := x(t)$ as the position of the cart, $x_2(t) := \theta(t)$ as the angle made by the pendulum with the vertical, $x_3(t) := \frac{\mathrm{d}x_1}{\mathrm{d}t}(t)$ as the velocity of the cart, $x_4(t) := \frac{\mathrm{d}x_2}{\mathrm{d}t}(t)$ as the angular velocity of the pendulum and the control input $t \mapsto u(t)$ as the horizontal force acting on the cart. The system matrix A and the input matrix B are defined by

$$A := \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & \frac{-m_0^2 g l^2}{(J_0 + m_0 l^2)p} & 0 & 0 \\ 0 & \frac{m_0 g l(m_0 + M)}{(J_0 + m_0 l^2)p} & 0 & 0 \end{bmatrix}, B := \begin{bmatrix} 0 \\ 0 \\ \frac{1}{p} \\ \frac{-m_0 l}{(J_0 + m_0 l^2)p} \end{bmatrix},$$

where $p := m_0 + M - (m_0^2 l^2 / J_0 + m_0 l^2)$. Here, m_0 is the mass of the pendulum, M is the mass of the cart, $2l$ is the length of the pendulum, L is the length of the cart track available on both sides of the origin, and J_0 is the moment of inertia of the pendulum (see [2] for more details on the model). Let us consider the quadratic cost

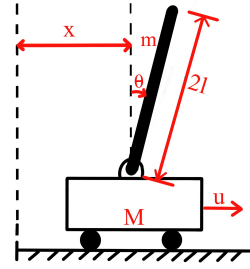$$J(u(\cdot)) := \int_0^T |u(t)|^2 \, \mathrm{d}t.$$

Fig. 1: Inverted pendulum on a cart setup.

1) Numerical simulation: We consider the sinusoidal basis functions for the discretized control trajectory in (5):

$$\psi_i^{\mathbf{s}}(t) := \begin{cases} 1 & \text{if } i = 1, \\ \sin(2\pi(i-1)t) & \text{if } i \in \{2, \ldots, \frac{N+1}{2}\}, \\ \cos\left(2\pi\left(i - \frac{N+1}{2}\right)t\right) & \text{if } i \in \{\frac{N+3}{2}, \ldots, N\}. \end{cases} \tag{16}$$

Substituting the expression (5) with $\psi_i(\cdot) = \psi_i^{\mathbf{s}}(\cdot)$ in the cost function of (16) leads to:

$$J(\alpha) = \alpha^\top P \alpha \quad \text{with } \alpha = \begin{pmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_N \end{pmatrix}^\top,$$

and the matrix P is defined by

$$P := \begin{pmatrix} \langle \psi_1, \psi_1 \rangle & \langle \psi_2, \psi_1 \rangle & \cdots & \langle \psi_N, \psi_1 \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle \psi_1, \psi_N \rangle & \langle \psi_2, \psi_N \rangle & \cdots & \langle \psi_N, \psi_N \rangle \end{pmatrix},$$

where $\langle \psi_i(\cdot), \psi_j(\cdot) \rangle$ is the standard T-horizon $\mathbb{L}^2$-inner product. Now we solve the global maximization problem:

$$\sup_{\mathcal{T} \in [0,T]^{\overline{N}}} \mathcal{G}(\mathcal{T}; x), \ x := x(0)$$

and the objective $\mathcal{G}(\cdot; x)$ for every $x \in \mathbb{X}_T$ is defined in the following fashion:

$$\mathcal{G}(\mathcal{T}; x) :=$$
$$\min_{\alpha \in \mathbb{R}^{m \times N}} \alpha^\top P \alpha$$

$$\text{sbj. to} \quad \begin{cases} \alpha \in \mathcal{P}, \\ |x(t_i)| \leqslant \varepsilon \text{ for } t_1, \ldots, t_{\overline{N}} \in [0, T], \\ \text{where } \varepsilon := (\varepsilon_1 \ \varepsilon_2 \ \varepsilon_3 \ \varepsilon_4)^\top. \end{cases}$$

Here, the inner minimization problem is constrained on a finite set of time instants $(t_1, \ldots, t_{\overline{N}})$.

2) Numerical simulation: In the numerical simulation, we consider the following data: $m_0 = 0.2\,\text{kg}$, $M = 3\,\text{kg}$, $l = 1.5\,\text{m}$, $L = 0.5\,\text{m}$. The initial position of the cart is at the origin and the track length is up to $0.5\,\text{m}$ on either side of the origin. The initial position of the states is $x(0) = (0, 0.035, 0, 0)^\top$ where the angle made by the pendulum with the vertical is two degrees to the right. The system is simulated for $T = 10$ sec. In this experiment, the number of basis functions is considered to be $N = 51$ and the bounds on the states are $\varepsilon_1 = 0.5$, $\varepsilon_2 = 0.07$, $\varepsilon_3 = 0.5$ and $\varepsilon_4 = 0.1$.

We solve the ensuing optimization problem using Algorithm 1, employing the simulated annealing routine for solving the global optimization, and the CPU time with 51 basis functions was 1466.23 secs for 2600 iterations of the global optimization. This is a practical number given the length of time horizon $T = 10$ secs. The state trajectories are depicted in Fig. 2 and Fig. 3; we can see that all the states reach the equilibrium position and stay within the specified bounds.

In Fig. 4, we observe hints of the turnpike property in the near-optimal $x_2(\cdot)$ trajectory which is the angle made by the inverted pendulum with the vertical axis from an initial state as it reaches the steady state and departs close to the bound at the end of the time interval $[0, 10]$. This observation points strongly to the justification of employing MPC in continuous-time.
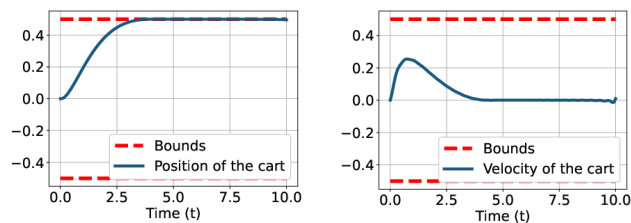

Fig. 2: Time evolution of the state $x_1(\cdot)$ and $x_3(\cdot)$ with 51 basis functions employed in (5). It can be noticed that the path constraint on the states are respected.
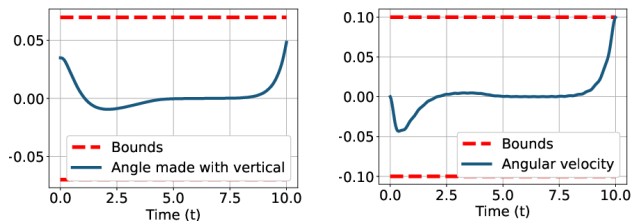

Fig. 3: Time evolution of the state $x_2(\cdot)$ and $x_4(\cdot)$ with 51 basis functions employed in (5) and it can be seen that the path constraints on the states are respected.
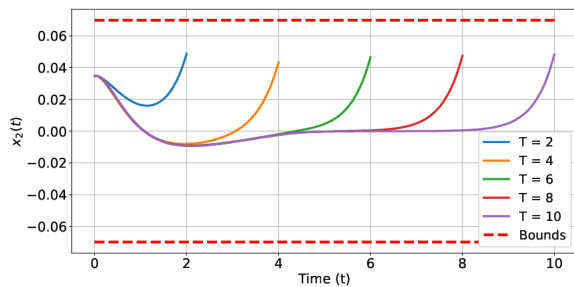

Fig. 4: Time evolution of the angle made by the pendulum with the vertical for different time horizons with 51 basis functions in (5). For each T, the angle reaches equilibrium and tends towards the bound at the end of the time horizon.

VI. CONCLUSION

This article established a numerical algorithm to solve a class of continuous-time optimal control problems in a tractable manner, with a focus on tackling model predictive control problems directly in the continuous-time regime satisfying the constraints *for every time*. A numerical example was included which demonstrates the numerical fidelity of the algorithm. Future work involves leveraging the tools developed herein to construct an offline explicit feedback map for a wide class of finite horizon robust optimal control problems.

REFERENCES

[1] J. M. Borwein. Direct theorems in semi-infinite convex programming. *Mathematical Programming*, pages 301–318, 12 1981.
[2] D. Chatterjee, A. Patra, and H. K. Joglekar. Swing-up and stabilization of a cart–pendulum system under restricted cart track length. *Systems & Control Letters*, 47:355–364, 2002.
[3] S. Das, A. Aravind, A. Cherukuri, and D. Chatterjee. Near-optimal solutions of convex semi-infinite programs via targeted sampling. *Annals of Operations Research*, 318(1):129–146, 2022. doi: `https://doi.org/10.1007/s10479-022-04810-4`.
[4] S. Das, S. Ganguly, M. Anjali, and D. Chatterjee. Towards continuous-time MPC: a novel trajectory optimization algorithm, 2023. doi: `https://doi.org/10.48550/arXiv.2306.07107`.
[5] M. Diehl and S. Gros. *Numerical Optimal Control*. Optimization in Engineering Center (OPTEC), 2011.
[6] G. N. Galbraith and R. B. Vinter. Lipschitz continuity of optimal controls for state constrained problems. *SIAM Journal on control and optimization*, 42(5):1727–1744, 2003.
[7] S. Ganguly and D. Chatterjee. Explicit feedback synthesis driven by quasi-interpolation for nonlinear robust model predictive control, 2023. doi: `https://doi.org/10.48550/arXiv.2306.03027`.
[8] S. Ganguly, N. Randad, D. Chatterjee, and R. Banavar. Constrained trajectory synthesis via quasi-interpolation. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4533–4538, 2022.
[9] D. Garg, M. Patterson, W. W. Hager, A. V. Rao, D. A. Benson, and G. T. Huntington. A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica*, 46(11):1843–1851, 2010.
[10] L. Grüne. Dissipativity and Optimal Control: Examining the Turnpike Phenomenon. *IEEE Control Systems Magazine*, 42(2):74–87, 2022.
[11] L. Magni and R. Scattolini. Stabilizing model predictive control of nonlinear continuous time systems. *Annual Reviews in Control*, 28(1):1–11, 2004.
[12] D. Q. Mayne and P. Falugi. Stabilizing conditions for model predictive control. *International Journal of Robust and Nonlinear Control*, 29(4):894–903, 2019. doi: `https://doi.org/10.1002/rnc.4409`.
[13] A. Muthyala, D. Chatterjee, and M. Nagahara. Control-theoretic splines under uncountably many constraints: fast and exact solutions. *to appear in IFAC World Congress 2023, Yokohama, Japan*.
[14] G. Pannocchia, J. B. Rawlings, D. Q. Mayne, and G. M. Mancuso. Whither discrete time model predictive control? *IEEE Transactions on Automatic Control*, 60(1):246–252, 2015. doi: `https://doi.org/10.1109/TAC.2014.2324131`.
[15] P. Paruchuri and D. Chatterjee. A numerical algorithm for attaining the chebyshev bound in optimal learning, 2023. doi: `https://doi.org/10.48550/arXiv.2307.01304`.