

Geometric Data-Driven Dimensionality Reduction in MPC with Guarantees

Roland Schurig¹, Andreas Himmel¹, Rolf Findeisen¹

Abstract—We address the challenge of dimension reduction in the discrete-time optimal control problem which is solved repeatedly online within the framework of model predictive control. Our study demonstrates that a reduced-order approach, aimed at identifying a suboptimal solution within a low-dimensional subspace, retains the stability and recursive feasibility characteristics of the original problem. We present a necessary and sufficient condition for ensuring initial feasibility, which is seamlessly integrated into the subspace design process. Additionally, we employ techniques from optimization on Riemannian manifolds to develop a subspace that efficiently represents a collection of pre-specified high-dimensional data points, all while adhering to the initial admissibility constraint.

I. INTRODUCTION

In model predictive control (MPC) [1], [2], [3], the control input is computed as the solution of an optimal control problem. The optimization must be performed online and in real time at every sampling instance of the control system. This limits its applicability for systems with small sampling times and/or limited hardware capacities.

Existing work to counteract this problem aims at alleviating the computational burden of solving the optimal control problem. A prominent idea is to approximate the problem in a lower-dimensional subspace. The design of the subspace can range from conceptually simple approaches like *move blocking* [4], [5] to more elaborated algorithms that involve *singular value decomposition* [6], [7], [8].

Many of these ideas share the paradigm that they are *data-driven*, i.e., the high-dimensional problem has to be sampled and the subspace is designed based on this data to capture its the most important patterns. Furthermore, some results for the stability of such schemes exist, e.g., [4], [6].

The goal of this contribution is to recall, collect and unify existing results for subspace-based dimensionality reduction in MPC. We carefully review the dimensionality reduction problem in the linear MPC context and present a general scheme, based on [6], [7], [9], which enjoys stability and feasibility guarantees. We provide self-contained proofs within our setting and notation, and add an upper bound for the closed-loop cost. Then, we relate the scheme to the geometric framework for dimensionality reduction introduced in [10].

We put special focus on the aspect of *initial* feasibility. We build upon that and add necessary and sufficient conditions

for initial feasibility in a given initial set and translate them into our geometric framework. While [10] mainly focused on the geometry of the subspace design, we also relate it lengthily to the MPC formulation.

We firmly believe that the geometric framework has numerous advantages, since it exploits the geometry of the subspace approximation and allows for a conceptually conclusive problem formulation. The growing number of methods and algorithms for solving (constrained) optimization problems over Riemannian manifolds pave the way for our formulation [11], [12], [13].

Notation: To denote nonnegative real numbers, we use $\mathbb{R}_{\geq 0} := \{x \in \mathbb{R} \mid x \geq 0\}$. The positive integers are \mathbb{N} , and $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. We use the standard notion of class \mathcal{K} and \mathcal{K}_∞ comparison functions [2, Def. 2.13].

II. MPC FORMULATION

Consider the problem of steering the discrete-time linear time-invariant system

$$x(t+1) = Ax(t) + Bu(t), \quad x(0) = x_s, \quad (1)$$

to the origin, with the pair (A, B) stabilizable, while obeying the constraints $x(t) \in \mathbb{X} \subseteq \mathbb{R}^n$, $u(t) \in \mathbb{U} \subseteq \mathbb{R}^m$, $\forall t \in \mathbb{N}_0$, where \mathbb{X} and \mathbb{U} are compact polyhedral sets containing the origin in their interior.

A suitable method to tackle this problem is MPC [1], [2], [3]. It relies on repeatedly solving an open-loop optimal control problem over a finite prediction horizon $N \in \mathbb{N}$.

In MPC, a *stage cost* $\ell: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$, $(x, u) \mapsto \ell(x, u)$ that penalizes the distance of the pair (x, u) from the origin is considered. Then, at each sampling time, MPC predicts N time steps into the future and chooses a control sequence that minimizes the summed stage cost over the horizon. The first part of the control sequence is applied to the plant, and the procedure is repeated at the next time step. We assume that the stage cost ℓ is quadratic with $\ell(x, u) := x^\top Qx + u^\top Ru$. The weighting matrices satisfy $Q \succ 0$ and $R \succ 0$.

Definition 1. We identify the set of input sequences over the horizon N with \mathbb{R}^d , where $d := Nm$. Therefore, we call $z \in \mathbb{R}^d$ an *input sequence* of length N . Given an input sequence $z = (z_0, \dots, z_{N-1}) \in \mathbb{R}^d$, we use $k \mapsto \text{pr}_k(z) := z_k \in \mathbb{R}^m$ for $k \in \{0, \dots, N-1\}$ to select its k^{th} component.

It is well-known that numerical problems and ill-conditioning in representing open-loop trajectories over a (long) horizon N arise if the plant (1) is open-loop unstable. Often, this problem is avoided by *pre-stabilizing* the system

¹ Technical University of Darmstadt, Control and Cyber-Physical Systems Laboratory, 64289 Darmstadt, Germany, {roland.schurig, andreas.himmel, rolf.findeisen}@iat.tu-darmstadt.de

The authors acknowledge support within the frame of the LOEWE initiative emergenCITY.

with a linear controller $x \mapsto \kappa(x) := Kx$ for some stabilizing feedback gain $K \in \mathbb{R}^{m \times n}$ [4], [6], [8].

Definition 2. Given an initial value $x_0 \in \mathbb{R}^n$ and an input sequence $z \in \mathbb{R}^d$, we define the pre-stabilized *trajectory* $x_z(\cdot, x_0)$ iteratively via $x_z(0, x_0) := x_0$ and

$$x_z(k+1, x_0) := Ax_z(k, x_0) + B(\kappa(x_z(k, x_0)) + \mathbf{pr}_k(z))$$

for $k \in \{0, \dots, N-1\}$. We refer to

$$\mathbf{u}_z(k, x_0) := \kappa(x_z(k, x_0)) + \mathbf{pr}_k(z), \quad k \in \{0, \dots, N-1\}$$

as the corresponding *control sequence* that produces the trajectory $x_z(\cdot, x_0)$ according to the system dynamics (1).

We further introduce the following standard terminal ingredients that are frequently used in MPC to establish formal stability guarantees [2], [3].

Assumption 1. There exists a closed polyhedral terminal set $\mathbb{X}_f \subseteq \mathbb{X}$ containing the origin in its interior, a quadratic terminal penalty $V_f: \mathbb{X}_f \rightarrow \mathbb{R}_{\geq 0}$ and a terminal controller $\kappa_f: \mathbb{X}_f \rightarrow \mathbb{U}$ that satisfy $Ax + B\kappa_f(x) \in \mathbb{X}_f$ and $V_f(x) - V_f(Ax + B\kappa_f(x)) \geq \ell(x, \kappa_f(x))$ for all $x \in \mathbb{X}_f$.

Definition 3. For a state $x \in \mathbb{X}$, we call the input sequence $z \in \mathbb{R}^d$ *admissible* if

$$\forall k \in \{0, \dots, N-1\}: (x_z(k, x), \mathbf{u}_z(k, x)) \in \mathbb{X} \times \mathbb{U}$$

and $x_z(N, x) \in \mathbb{X}_f$ holds [2]. The set of admissible input sequences for x is denoted by $\mathbb{U}^N(x)$.

Additionally, we introduce the *feasible set* $\mathcal{X}_N = \{x \in \mathbb{X} \mid \mathbb{U}^N(x) \neq \emptyset\}$ as the set of states for which we can find at least one admissible input sequence.

Finally, for $x \in \mathbb{R}^n$ and $z \in \mathbb{R}^d$, we can introduce the cost function over the horizon N as $J_N(x, z) := V_f(x_z(N, x)) + \sum_{k=0}^{N-1} \ell(x_z(k, x), \mathbf{u}_z(k, x))$. Then, we define the *constrained finite time optimal control problem*

$$\mathcal{P}_N(x): \min_{z \in \mathbb{R}^d} J_N(x, z) \quad \text{s.t. } z \in \mathbb{U}^N(x) \quad (2)$$

on which the standard MPC scheme is based. We introduce the following objects to describe $\mathcal{P}_N(x)$: The *optimal value function* $V_N: \mathcal{X}_N \rightarrow \mathbb{R}_{\geq 0}$, which maps any $x \in \mathcal{X}_N$ to the corresponding optimal value of (2), i.e., $V_N(x) \leq J_N(x, z)$ holds for all $z \in \mathbb{U}^N(x)$. The *optimizer mapping* $\mu_N: \mathcal{X}_N \rightarrow \mathbb{R}^d$, which maps any $x \in \mathcal{X}_N$ to the corresponding optimal input sequence of (2), i.e., $V_N(x) = J_N(x, \mu_N(x))$.

The MPC controller asymptotically stabilizes (1) on the feasible set \mathcal{X}_N under the assumptions made. For more details, we refer to, e.g., [2, Th. 5.13]. Because the region of attraction is \mathcal{X}_N , we introduce an additional assumption on the initial state of the plant.

Assumption 2. The initial value x_s of (1) is contained in the compact set $\mathbb{X}_0 = \text{conv}(X_0) \subseteq \mathcal{X}_N$, where $X_0 := \{\bar{x}_1, \dots, \bar{x}_s\}$, $\bar{x}_j \in \mathcal{X}_N$ and conv denotes the convex hull of a set.

Remark 1. Problem (2) is a *multiparametric quadratic program* with parameter x and optimization variable z [3, Sec. 6.3]. The admissible sets $\mathbb{U}^N(x)$ are bounded polyhedra.

III. DIMENSIONALITY REDUCTION IN MPC

We aim to simplify optimization problem (2), which must be solved at every sampling instance of an MPC scheme. A classical approach to diminish the computational burden in $\mathcal{P}_N(x)$ is to reduce the number of (scalar) decision variables from d to r , with $r \ll d$; see, e.g., [4], [5], [6], [7], [10].

One of the most established techniques and a cornerstone in statistics for dimensionality reduction and identifying patterns in data is *principal component analysis* (PCA); see, e.g., [12], [14]. PCA offers a systematic way to uncover the r *dominant directions* of variation of the data [12]. The underlying assumption is that the data is located on or near a r -dimensional affine subspace in \mathbb{R}^d . By transforming to this new coordinate system, a low-dimensional approximation of the high-dimensional data is available [14].

Motivated by PCA, we aim for a data-driven design of suitable subspaces in the context of MPC. We consider a data set of the form

$$\mathcal{D} = \{(x_i, z_i)\}_{i=1}^L \subset \mathcal{X}_N \times \mathbb{R}^d, \quad (3)$$

where $z_i := \mu_N(x_i)$, i.e., our data consists of pairs of states *and* corresponding optimal input sequences. We seek to find an affine subspace – represented by a *basis* and an *offset* – that best captures the dominant directions of variation in \mathcal{D} .

The *compact Stiefel manifold*

$$\mathcal{ST}(r, d) := \{U \in \mathbb{R}^{d \times r} \mid U^\top U = I_r\}$$

offers a convenient way to represent a basis, since the columns of $U \in \mathcal{ST}(r, d)$ form an orthonormal basis of an r -dimensional subspace of \mathbb{R}^d . We refer to [15] for an introduction to the Stiefel manifold.

In order to capture the state information in \mathcal{D} , we allow for a *state-dependent* offset $\sigma: \mathcal{X}_N \rightarrow \mathbb{R}^d$, i.e., we seek to identify a basis $U \in \mathcal{ST}(r, d)$ whose column space best approximates $z_i - \sigma(x_i)$. We restrict the class of possible offsets to

$$\mathcal{F}_N := \{\sigma: \mathcal{X}_N \rightarrow \mathbb{R}^d \mid \sigma: x \mapsto \Gamma x + \xi\}.$$

This has desirable properties when it comes to guaranteeing feasibility in the MPC context and combines ideas from previous works [7], [9]. For a discussion on how to obtain $(U, \sigma) \in \mathcal{ST}(r, d) \times \mathcal{F}_N$, we refer to Sec. IV.

Thus, for a given pair (U, σ) , an intuitive, low-dimensional reformulation of (2) is to replace the optimization variable z with $U\alpha + \sigma(x)$. Note that this reduces the number of scalar decision variables from d to r . However, this intuitive reformulation does not yet fulfill the requirements we have if the optimization problem is to be solved in closed-loop in a receding horizon fashion. In particular, we must put a special focus on the issue of *feasibility*. More precisely, we shall discuss if and when for all states $x \in \mathcal{X}_N$ there exists an $\alpha_x \in \mathbb{R}^r$ such that $U\alpha_x + \sigma(x) \in \mathbb{U}^N(x)$ holds.

In general, the answer to this question depends on how the pair (U, σ) is chosen. However, a common approach in reduced MPC to circumvent this problem is to hand an admissible guess $\tilde{z} \in \mathbb{U}^N(x)$ – however obtained – to the solver as a “fall-back” option [7].

Therefore, given $x \in \mathcal{X}_N$ and $\tilde{z} \in \mathbb{R}^d$, we introduce the *reduced-order MPC problem*

$$\begin{aligned} \widetilde{\mathcal{P}}_N(x, \tilde{z}): \quad & \min_{(\alpha, \tau)} \quad J_N(x, U\alpha + \tau\sigma(x) + (1 - \tau)\tilde{z}) \\ & \text{s.t.} \quad U\alpha + \tau\sigma(x) + (1 - \tau)\tilde{z} \in \mathbb{U}^N(x) \end{aligned} \quad (4)$$

to replace $\mathcal{P}_N(x)$, with $\tau \in \mathbb{R}$. Note that the optimization problem accepts an additional input \tilde{z} and we have added an extra degree of freedom in τ .

Remark 2. If $\tilde{z} \in \mathbb{U}^N(x)$, then $\widetilde{\mathcal{P}}_N(x, \tilde{z})$ is guaranteed to have the (suboptimal) solution $\tilde{\alpha} = 0$ and $\tilde{\tau} = 0$. Thus, we have introduced τ to ensure that in this case the solver can always return an admissible input sequence by scaling the influence of the basis and the offset to zero.

Similar to the full-order case, we use $\widetilde{V}_N(x, \tilde{z})$ to denote the resulting optimal value, which depends on the choice of (U, σ) . However, to not overload notation and maintain readability, we suppress this dependence here.

Suppose that the optimizer of $\widetilde{\mathcal{P}}_N(x, \tilde{z})$ is (α^*, τ^*) . Then, the corresponding optimal input sequence is given by $\tilde{\mu}_N(x, \tilde{z}) := U\alpha^* + \tau^*\sigma(x) + (1 - \tau^*)\tilde{z}$, i.e., $\widetilde{V}_N(x, \tilde{z}) = J_N(x, \tilde{\mu}_N(x, \tilde{z}))$.

We now shift our focus to the question of how admissible guesses $\tilde{z} \in \mathbb{U}^N(x)$ can be obtained. In MPC with terminal conditions, a common approach to construct an admissible solution at time $t \in \mathbb{N}$ in *closed-loop* is to consider the admissible optimal solution from the *last* time step $t - 1$, c.f. [4], [5], [6], [7]. The procedure is based on the terminal ingredients in Assumption 1 and relies on standard arguments.

Proposition 1. Given a state $x \in \mathcal{X}_N$ and an input sequence $z \in \mathbb{U}^N(x)$, we create an input sequence $\mathbf{s}_f(x, z) \in \mathbb{U}^N(x_z(1, x))$. If $x_z(1, x) = 0$, then we set $\mathbf{s}_f(x, z) := 0$. Otherwise, we define

$$\mathbf{pr}_k(\mathbf{s}_f(x, z)) := \begin{cases} \mathbf{pr}_{k+1}(z), & k \in \{0, \dots, N - 2\} \\ \kappa_f(x_N) - \kappa(x_N), & k = N - 1 \end{cases}$$

with $x_N = x_z(N, x)$. We call $\mathbf{s}_f(x, z)$ the *admissibly shifted input sequence* for x and z .

Proof. We show that the inclusion $\mathbf{s}_f(x, z) \in \mathbb{U}^N(x_z(1, x))$ is true. For $x_z(1, x) = 0$, it trivially holds. For $x_z(1, x) \neq 0$, the inclusion is implied by construction of \mathbf{s}_f that exploits the properties in Assumption 1; see, e.g., [2, Lemma 5.10]. \square

Therefore, if $x \in \mathcal{X}_N$ and $z^* \in \mathbb{U}^N(x)$ denote the state and the optimal input sequence at time $t - 1 \in \mathbb{N}$, we can hand $\tilde{z}^+ = \mathbf{s}_f(x, z^*)$ to (4) as the admissible guess for the state $x^+ = x_{z^*}(1, x)$ at time t .

With this reasoning, however, we *cannot* rely on an optimal solution from the last time step at time $t = 0$. From Assumption 2 we know that the initial state x_s at time $t = 0$ is in \mathbb{X}_0 . Consequently, we must establish the existence of admissible input sequences in (4) for all $x_s \in \mathbb{X}_0$ by just relying on U and σ .

We shift the discussion on how this can be incorporated into the design of U and σ to Sec. IV. For now, we content ourselves with the following definition.

Algorithm 1 Reduced-order MPC scheme.

Require: Initially admissible pair $(U, \sigma) \in \mathcal{ST}(r, d) \times \mathcal{F}_N$.

- 1: Initialize $\tilde{z} \leftarrow 0$.
- 2: **repeat** At each sampling time $t \in \mathbb{N}_0 \dots$
- 3: Measure the system state $x(t) \in \mathbb{X}$ and set $x := x(t)$.
- 4: Solve $\widetilde{\mathcal{P}}_N(x, \tilde{z})$ in (4) with (U, σ) and denote the resulting input sequence by $z^* := \tilde{\mu}_N(x, \tilde{z})$.
- 5: Update $\tilde{z} \leftarrow \mathbf{s}_f(x, z^*)$.
- 6: Define the reduced-order MPC feedback $\tilde{\kappa}_N(x, \tilde{z}) := \mathbf{u}_{z^*}(0, x)$ and apply it to the system.
- 7: **until** system state has converged to the origin;

Definition 4. We call the pair (U, σ) given by a basis $U \in \mathcal{ST}(r, d)$ and an offset $\sigma \in \mathcal{F}_N$ *initially admissible (IA)* if

$$\forall x_s \in \mathbb{X}_0: \exists \alpha \in \mathbb{R}^r: U\alpha + \sigma(x_s) \in \mathbb{U}^N(x_s).$$

Throughout the remainder of the section, we assume that the pair (U, σ) is initially admissible. This notion allows the introduction of Alg. 1 as our reduced-order MPC scheme. It is based on the reduced problem (4) and the admissible shift \mathbf{s}_f . In the following, we analyze the closed-loop system resulting from controlling (1) with Alg. 1.

The first important point to recognize is that the computation of the control value at time t does *not* only depend on the current state x of the plant, but also on the admissible guess \tilde{z} that is updated throughout the algorithm based on the solution of (4). This is the case because for $\tilde{z}_1, \tilde{z}_2 \in \mathbb{U}^N(x)$ with $\tilde{z}_1 \neq \tilde{z}_2$, we have $\tilde{\mu}_N(x, \tilde{z}_1) \neq \tilde{\mu}_N(x, \tilde{z}_2)$ in general. Therefore, the state information x of the plant (1) is *not* sufficient to describe what its next state $x^+ = Ax + B\tilde{\kappa}_N(x, \tilde{z})$ in closed-loop will be; the information \tilde{z} on the current admissible guess is required as well.

For this reason, we work with an extended state, which consists of the plant's state *and* the admissible guess. For $x_s \in \mathbb{X}_0$, we introduce $\chi(0, x_s) := x_s$, $z(0, x_s) := 0$ and

$$\begin{aligned} \chi(t + 1, x_s) &:= A\chi(t, x_s) + B\tilde{\kappa}_N(\chi(t, x_s), z(t, x_s)), \\ z(t + 1, x_s) &:= \mathbf{s}_f(\chi(t, x_s), \tilde{\mu}_N(\chi(t, x_s), z(t, x_s))). \end{aligned}$$

Here, $\chi(t, x_s)$ denotes the state of the plant in closed-loop at time $t \in \mathbb{N}_0$ when the initial state at time $t = 0$ was x_s , while $z(t, x_s)$ denotes the admissible guess at the same time.

We define the state space of the closed-loop dynamics as $M := M_1 \cup M_2 \subseteq \mathcal{X}_N \times \mathbb{R}^d$ with $M_1 := \{(x_s, 0) \mid x_s \in \mathbb{X}_0\}$ and $M_2 := \{(x, \tilde{z}) \mid x \in \mathcal{X}_N \setminus \{0\}, \tilde{z} \in \mathbb{U}^N(x)\} \cup \{(0, 0)\}$.

Remark 3. This somewhat overelaborate definition serves two purposes. On the one hand, it explicitly gives prominence to the nature of the initial extended state in Alg. 1. On the other hand, it highlights the fact that by construction of \mathbf{s}_f , once the plant state x reaches the equilibrium, it stays there. (We show that in our next result.) We choose this definition to stress that (i) the solution of (4) is special at time $t = 0$ in the sense that feasibility solely originates from (U, σ) and (ii) the entire scheme only works under the initial admissibility assumption on (U, σ) . For subsequent times, feasibility of (4) will stem from Proposition 1.

We introduce the map $g: M \rightarrow \mathbb{R}^n \times \mathbb{R}^d$, $p = (x, \tilde{z}) \mapsto g(p) = (x^+, \tilde{z}^+)$, with $x^+ := Ax + B\tilde{\kappa}_N(x, \tilde{z})$ and $\tilde{z}^+ := \mathbf{s}_f(x, \tilde{\mu}_N(x, \tilde{z}))$, and establish some basic properties.

Lemma 1. *The value $g(p)$ exists for each $p \in M$ and the inclusion $g(p) \in M_2$ holds for each $p \in M$.*

Proof. Assume first that $(x, \tilde{z}) \in M_2$. Then, $\tilde{\mathcal{P}}_N(x, \tilde{z})$ is feasible by construction, so $\tilde{\mu}_N(x, \tilde{z})$ exists. If $(x, \tilde{z}) \in M_1$, then $\tilde{\mu}_N(x, \tilde{z})$ is guaranteed to exist under the IA assumption on (U, σ) . Thus, we have the existence of $\tilde{\mu}_N(x, \tilde{z}) \in \mathbb{U}^N(x)$ and hence also of $\tilde{\kappa}_N(x, \tilde{z})$ for all $(x, \tilde{z}) \in M$. Also, for $(x, \tilde{z}) \in M$, we have $x \in \mathcal{X}_N$, so $\mathbf{s}_f(x, \tilde{\mu}_N(x, \tilde{z}))$ is well-defined. This establishes the existence of $g(p)$ for all $p \in M$.

For the second part, we have to show that $g(p)$ is never of the form $(0, \tilde{z})$ with $\tilde{z} \neq 0$, since $(0, \tilde{z})$ implies $\tilde{z} = 0$ by construction of M_2 . We first consider $p = (0, 0)$. Then, we have $\tilde{\mu}_N(0, 0) = 0$, which implies $g(p) = (0, 0) \in M_2$. Next, we consider $p = (x, \tilde{z})$ with $Ax + B\tilde{\kappa}_N(x, \tilde{z}) = 0$. Then, $\mathbf{s}_f(x, \tilde{\mu}_N(x, \tilde{z})) = 0$ by definition of \mathbf{s}_f , which implies $g(p) = (0, 0) \in M_2$. Finally, for all remaining $p = (x, \tilde{z}) \in M$, we have $g(p) = (x^+, \tilde{z}^+)$ with $x^+ \neq 0$. By Proposition 1, the inclusion $\tilde{z}^+ \in \mathbb{U}^N(x^+)$ holds, which confirms the claim $g(p) \in M_2$ for all $p \in M$. \square

This result justifies to introduce the difference equation

$$p(t+1) = g(p(t)), \quad p(0) = p_s \in M. \quad (5)$$

We denote the solution of (5) with initial state $p_s \in M$ by $p(t, p_s)$. Note that for $p_s = (x_s, 0) \in M_1$, the equality $p(t, p_s) = (\chi(t, x_s), z(t, x_s))$ holds, i.e., $p(\cdot, p_s)$ represents the trajectory resulting from controlling (1) with Alg. 1.

Hence, we have established that (5) is well-defined and that it describes the closed-loop dynamics of our reduced MPC scheme. Next, we turn our attention to the closed-loop properties. A contribution of this work is to explicitly state an upper bound of the closed-loop cost of our reduced MPC scheme. For $p = (x, \tilde{z}) \in M$, we introduce $\tilde{\ell}(p) := \ell(x, \tilde{\kappa}_N(x, \tilde{z}))$ as a short-hand notation for the stage cost in closed-loop for the reduced MPC controller. Using the closed-loop dynamics (5) with initial state $p_s = (x_s, 0) \in M_1$, we introduce $\tilde{J}_{\text{cl}}(x_s) := \sum_{t=0}^{\infty} \tilde{\ell}(p(t, p_s))$ as the *closed-loop cost* of the reduced MPC controller for $x_s \in \mathbb{X}_0$. We first state an intermediate result, c.f. [6], [9].

Lemma 2. *The inequality $\tilde{V}_N(g(p)) \leq \tilde{V}_N(p) - \tilde{\ell}(p)$ holds for all $p = (x, \tilde{z}) \in M$.*

Proof. We denote $z^* = \tilde{\mu}_N(p)$, $u = \tilde{\kappa}_N(p)$, $x_N = \mathbf{x}_{z^*}(N, x)$ and $g(p) = (x^+, \tilde{z}^+) \in M$. By construction, $x_N \in \mathbb{X}_f$, $\mathbf{u}_{\tilde{z}^+}(N-1, x^+) = \kappa_f(x_N)$ and $\mathbf{x}_{\tilde{z}^+}(k, x^+) = \mathbf{x}_{z^*}(k+1, x)$ for $k \in \{0, \dots, N-1\}$. This implies $J_N(g(p)) = V_f(Ax_N + B\kappa_f(x_N)) + \tilde{V}_N(p) - \ell(x, u) - V_f(x_N) + \ell(x_N, \kappa_f(x_N))$. Thus, with Assumption 1 and $\tilde{V}_N(g(p)) \leq J_N(g(p))$, we obtain $\tilde{V}_N(g(p)) \leq \tilde{V}_N(p) - \ell(x, u)$. \square

Theorem 1. *The closed-loop cost $\tilde{J}_{\text{cl}}(x_s)$ is upper bounded by $\tilde{V}_N(x_s, 0)$ for all $x_s \in \mathbb{X}_0$.*

Proof. We follow the proof of [2, Th. 4.11]. By Lemma 2, the dynamic programming inequality $\tilde{V}_N(p) - \tilde{V}_N(g(p)) \geq \tilde{\ell}(p)$ holds for all $p \in M$. Then, under Lemma 1, considering a closed-loop trajectory $p(\cdot, p_s)$ starting from $p_s = (x_s, 0) \in M_1$ and summing over t gives for all $K \in \mathbb{N}$

$$\begin{aligned} \sum_{t=0}^K \tilde{\ell}(p(t, p_s)) &\leq \sum_{t=0}^{K-1} \tilde{V}_N(p(t, p_s)) - \tilde{V}_N(p(t+1, p_s)) \\ &= \tilde{V}_N(p(0, p_s)) - \tilde{V}_N(p(K, p_s)) \leq \tilde{V}_N(p_s) = \tilde{V}_N(x_s, 0). \end{aligned}$$

Since ℓ is nonnegative, the term on the left is monotone increasing and bounded. Therefore, for $K \rightarrow \infty$, it converges to $\tilde{J}_{\text{cl}}(x_s)$. The assertion follows because the right-hand side of the inequality is independent of K . \square

Finally, we focus on stability. Similar results can be found in [6], [9]. We show that the plant's state converges to the origin in closed-loop. To exclude that the state deviates too far from to origin before it eventually converges to it, we show that Alg. 1 leads to a stable closed-loop system in the sense of [2, Def. 10.24]. While this result does not establish asymptotic stability, it suffices within the context of this work.

Theorem 2. *The solution $p(\cdot, p_s) = (\chi(\cdot, x_s), z(\cdot, x_s))$ of (4) satisfies $\lim_{t \rightarrow \infty} \chi(t, x_s) = 0$ for all $p_s = (x_s, 0) \in M_1$. In addition, consider any compact set $S \subseteq \mathbb{X}_0$ with $0 \notin S$. There exists $\alpha_S \in \mathcal{K}$ such that $\|\chi(t, x_s)\| \leq \alpha_S(\|x_s\|)$ holds for all $x_s \in S$ and $t \in \mathbb{N}_0$.*

Proof. We follow the proof of [2, Proposition 10.25]. First, recall from the proof of Th. 1 that $\lim_{K \rightarrow \infty} \sum_{t=0}^K \tilde{\ell}(p(t, p_s)) \leq \tilde{V}_N(p_s)$ holds. Nonnegativity of $\tilde{\ell}$ then implies $\lim_{t \rightarrow \infty} \tilde{\ell}(p(t, p_s)) = 0$. In general, we have $\tilde{\ell}(p) = x^\top Qx + \tilde{\kappa}_N(p)^\top R\tilde{\kappa}_N(p) \geq x^\top Qx \geq \lambda_{\min}(Q)\|x\|^2$ for any $p = (x, \tilde{z}) \in M$, where we have used $Q \succ 0$. This implies $\chi(t, x_s) \rightarrow 0$.

For the second part, we first show that $\tilde{V}_N(\cdot, 0)$ is upper-bounded by a \mathcal{K}_∞ function on S , for which we apply ideas taken from the proof of [2, Proposition 5.7]. The value $\tilde{V}_N(x_s, 0)$ exists for each $x_s \in S$ by the initial admissibility assumption. The positive value $l_{\min} := \min_{x_s \in S} \|x_s\|$ exists by compactness of S . We define $N_l := B_l(0) \cap S$, where $B_l(0)$ denotes the closed ball around 0 with radius l . Note that N_l is compact, which implies that $N_l \times \mathbb{U}^N$ is compact. Thus, for $l \geq l_{\min}$, we can compute

$$\hat{\alpha}_2(l) := \max\{V_f(x_N) + \sum_{k=0}^{N-1} \ell(x_k, u_k) \mid$$

$$x_0 \in N_l, (u_0, \dots, u_{N-1}) \in \mathbb{U}^N, x_{k+1} = Ax_k + Bu_k\}$$

which implies $\tilde{V}_N(x_s, 0) \leq \hat{\alpha}_2(\|x_s\|)$ for all $x_s \in S$. The resulting function $\hat{\alpha}_2$ is continuous and monotone increasing in l . Next, we pick an $\alpha'_2 \in \mathcal{K}_\infty$ with $\alpha'_2(l_{\min}) = \hat{\alpha}_2(l_{\min})$ and any $\bar{\alpha}_2 \in \mathcal{K}_\infty$. With this, we define $\tilde{\alpha}_2 \in \mathcal{K}_\infty$ by

$$\tilde{\alpha}_2(l) := \begin{cases} \alpha'_2(l), & l \in [0, l_{\min}), \\ \hat{\alpha}_2(l) + \bar{\alpha}_2(l - l_{\min}), & l \geq l_{\min}, \end{cases}$$

which satisfies $\tilde{V}_N(x_s, 0) \leq \tilde{\alpha}_2(\|x_s\|)$ for all $x_s \in S$.

Secondly, we note that $S \times \{0\} \subseteq M_1$ holds. Hence, using Lemma 2, for any initial value $p_s = (x_s, 0) \in S \times \{0\}$ we obtain $\tilde{V}_N(p(t, p_s)) \geq \tilde{\ell}(p(t, p_s)) \geq \lambda_{\min}(Q)\|\chi(t, x_s)\|^2 =: \tilde{\alpha}_1(\|\chi(t, x_s)\|)$ and $\tilde{V}_N(p(t, p_s)) \leq \tilde{V}_N(p_s)$ for all $t \in \mathbb{N}_0$. This implies

$$\begin{aligned} \|\chi(t, x_s)\| &\leq \tilde{\alpha}_1^{-1}(\tilde{V}_N(p(t, p_s))) \leq \tilde{\alpha}_1^{-1}(\tilde{V}_N(p_s)) \\ &\leq (\tilde{\alpha}_1^{-1} \circ \tilde{\alpha}_2)(\|x_s\|) =: \alpha_S(\|x_s\|) \end{aligned}$$

for all $x_s \in S$ and $t \in \mathbb{N}_0$. \square

Remark 4. From a computational perspective, in Alg. 1, once $\chi(t, x_s) \in \mathbb{X}_f$, it is possible to replace $\tilde{\kappa}$ with the terminal controller κ_f , in case κ_f is easier to evaluate. This is, e.g., the case when the terminal ingredients are based on the unconstrained LQR controller. Then, $\kappa_f(x) = K_\infty x$, c.f. [3].

IV. DATA-DRIVEN SUBSPACE DESIGN

As introduced in Sec. III, it is possible to design a reduced-order and stable MPC scheme based on a low-dimensional representation of the high-dimensional optimizers $\mu_N(x)$. A major assumption, however, is that the pair $(U, \sigma) \in \mathcal{ST}(r, d) \times \mathcal{F}_N$ in Alg. 1 is initially admissible. This has to be incorporated into the design of U and σ . To this end, we state an important observation for linear systems, c.f. [9].

Lemma 3. *Let $(U, \sigma) \in \mathcal{ST}(r, d) \times \mathcal{F}_N$. Then:*

$$(U, \sigma) \text{ is IA} \Leftrightarrow \forall \bar{x}_j \in X_0: \exists \tilde{\alpha}_j \in \mathbb{R}^r: U\tilde{\alpha}_j + \sigma(\bar{x}_j) \in \mathbb{U}^N(\bar{x}_j).$$

Proof. The \Rightarrow -direction is trivial, since by Assumption 2, $\bar{x}_j \in \mathbb{X}_0$ for all $j \in \{1, \dots, s\}$. For the \Leftarrow -direction, consider any $x_s \in \mathbb{X}_0$. Since $\mathbb{X}_0 = \text{conv}(X_0)$, there is $\theta_j \geq 0$ for $j \in \{1, \dots, s\}$ with $\sum_{j=1}^s \theta_j = 1$ such that $x_s = \sum_{j=1}^s \theta_j \bar{x}_j$. It follows from convexity arguments that $U\tilde{\alpha} + \sigma(x_s) \in \mathbb{U}^N(x_s)$, with $\tilde{\alpha} = \sum_{j=1}^s \theta_j \tilde{\alpha}_j$. \square

Remark 5. Checking if the pair (U, σ) is initially admissible thus boils down to checking whether it offers admissible input sequences for a *finite* number of states – namely for the vertices of \mathbb{X}_0 .

While initial admissibility is an assumption that *must* be met in Alg. 1, we also seek to design (U, σ) such that they are able to explain the most important patterns of the high-dimensional optimizers in \mathcal{D} , so that the re-formulation from (2) to (4) is justified.

In the remainder of the paper, though, we focus on identifying the *optimal basis* $U \in \mathcal{ST}(r, d)$, i.e., we fix the offset to a specific $\sigma_0 \in \mathcal{F}_N$. A possible way to select the data (Γ_0, ξ_0) that describes the offset σ_0 is $\xi_0 = \sum_{i=1}^L \omega_i z_i$, with $\omega_i \geq 0$, $\sum_{i=1}^L \omega_i = 1$, and $\Gamma_0 = Z(X_s)^\dagger$, where $Z = [z_1, \dots, z_L]$, $X_s = [x_1 - \xi_0, \dots, x_L - \xi_0]$ and $(\cdot)^\dagger$ denotes the Moore-Penrose pseudoinverse [9]¹.

As the offset is now fixed, we introduce – based on the original data set (3) – a *shifted* data set

$$\mathcal{D}_s := \{z_i - \sigma_0(x_i) \mid (x_i, z_i) \in \mathcal{D}\},$$

¹In the light of Lemma 3, the weights ω_i could be used to favor optimal input sequences associated to states close to the vertices of \mathbb{X}_0 .




Fig. 1: Illustration of the subspace approach for $d = 2, r = 1$.

and the *shifted* admissible sets $\mathcal{P}_j := \{z - \sigma_0(\bar{x}_j) \mid z \in \mathbb{U}^N(\bar{x}_j)\}$ for $j \in \{1, \dots, s\}$.

We start our discussion by recalling an important fact concerning suboptimality and admissibility of (4) from [10, Lemma 1]. It states that whether $\tilde{\mathcal{P}}_N(x_0, 0)$ can be solved with the pair $(U, \sigma_0) \in \mathcal{ST}(r, d) \times \mathcal{F}_N$ solely depends on the *subspace* spanned by the columns of U , and not directly on U itself. To see that, consider $\tilde{U} = UR \in \mathcal{ST}(r, d)$ with $R \in O(r)$, where $O(r)$ denotes the orthogonal group of dimension r . Note that $\text{span}(\tilde{U}) = \text{span}(U)$. Solving $\tilde{\mathcal{P}}_N(x, 0)$ with either (U, σ_0) or (\tilde{U}, σ_0) leads to the same solution $\tilde{V}_N(x, \tilde{z})$, since the part $U\alpha^*$ of the optimizer can also be represented by \tilde{U} via $\tilde{U}\tilde{\alpha}^* = \tilde{U}(R^\top\alpha^*) = U\alpha^*$.

This gives rise to the *equivalence class* $[U] := \{UR \in \mathcal{ST}(r, d) \mid R \in O(r)\}$, and motivates the introduction of the so-called *Grassmann manifold* [12], [15], [16] as the quotient space $\mathcal{GR}(r, d) := \mathcal{ST}(r, d)/O(r) = \{[U] \mid U \in \mathcal{ST}(r, d)\}$.

Intuitively, the Grassmann manifold $\mathcal{GR}(r, d)$ is the set of all r -dimensional subspaces of \mathbb{R}^d . A subspace $[U] \in \mathcal{GR}(r, d)$ can be equivalently expressed by the unique orthogonal projector $P = \pi(U) := UU^\top$ onto it [16]; see [10] for a more elaborate discussion.

The upcoming discussion is supported by Fig. 1. We consider $\delta_i \in \mathcal{D}_s$, and an arbitrary subspace $[U] \cong P = \pi(U)$. Firstly, we can *project* δ_i onto $\text{span}(U)$ by virtue of the projector P . As we can see, $\delta_{p,i} := P\delta_i$ is the point in $\text{span}(U)$ with the minimal Euclidean distance to δ_i . The point $\delta_{p,i} \in \text{span}(U)$ also has a low-dimensional representation $\alpha_i := U^\top\delta_i \in \mathbb{R}^r$. We refer to α_i as the *latent variable* for δ_i with respect to U . Note that $U\alpha_i = \delta_{p,i}$.

An important observation, though, is that the latent variable depends on the *basis* that is chosen to span the subspace, and not on the subspace itself. Consider $\tilde{U} = UR \in [U]$ with $\tilde{U} \neq U$. Surely, there exists $\tilde{\alpha}_i \in \mathbb{R}^r$ such that $\tilde{U}\tilde{\alpha}_i = \delta_{p,i}$, but we have the relation $\tilde{\alpha}_i = R^\top\alpha_i \neq \alpha_i$.

Because the feasibility of $\tilde{\mathcal{P}}_N(x, 0)$ solely depends on the subspace, the latent variable is not well suited to represent a point in $\text{span}(U)$ in the design phase of the subspace since it loses its meaning and interpretation under a change of basis. Hence, we suggest the following design task.

Problem 1. Design a subspace $[U] \in \mathcal{GR}(r, d)$ such that

- (i) The subspace $[U]$ intersects all shifted admissible sets: $\forall j \in \{1, \dots, s\}: \exists \tilde{z}_j \in \text{span}(U): \tilde{z}_j \in \mathcal{P}_j$.
- (ii) The subspace $[U]$ minimizes the Euclidean distance of δ_i to $\text{span}(U)$ for all $\delta_i \in \mathcal{D}_s$.

Remark 6. More elaborate optimality requirements than (ii) are possible on $\mathcal{GR}(r, d)$. We restrict our analysis to this simple one for the sake of an easy interpretation and exposition.

A. Design of the optimal subspace

We present an optimization problem over $\mathcal{GR}(r, d)$ as a possible solution to Problem 1, and compare our approach to an existing one that is based on convex programming.

1) *Riemannian approach:* We aim to minimize the distance between the points in \mathcal{D}_s and the subspace $[U] \in \mathcal{GR}(r, d)$. Following our previous discussion, we directly consider the projection $P\delta_i$ of δ_i onto $\text{span}(U)$ via the projector $P = \pi(U)$. This retains the interpretability of the formulation under a change of basis $\tilde{U} \in [U]$. Thus, we introduce $f: \mathcal{GR}(r, d) \rightarrow \mathbb{R}_{\geq 0}$, $P \mapsto \sum_{i=1}^L \|\delta_i - P\delta_i\|^2$ as objective function, c.f. Fig. 1.

We can now turn to the initial admissibility constraint, which is illustrated in Fig. 2a. We show the shifted admissible set \mathcal{P}_j for one $j \in \{1, \dots, s\}$ and a subspace $[U] \cong P = \pi(U) \in \mathcal{GR}(r, d)$. We again rely on the concept of projection. For admissibility, we consider the maximum volume ellipsoid center $\bar{\delta}_j$ of \mathcal{P}_j [17, Sec. 8.5.2]; this center can also be seen in Fig. 2a. The idea is to demand that the point on $\text{span}(U)$ with the smallest Euclidean distance to the center $\bar{\delta}_j$ of \mathcal{P}_j should still be in \mathcal{P}_j . As discussed before, $P\bar{\delta}_j \in \text{span}(U)$ is exactly this point, so we require $P\bar{\delta}_j \in \mathcal{P}_j$ for all $j \in \{1, \dots, s\}$. In conclusion, to tackle Problem 1, we propose to solve


$$\min_{P \in \mathcal{GR}(r, d)} f(P) \quad \text{s.t.} \quad \forall j \in \{1, \dots, s\}: P\bar{\delta}_j \in \mathcal{P}_j. \quad (6)$$

Corollary 1. *Assume that $P = \pi(U) \in \mathcal{GR}(r, d)$ is a feasible point of (6). Then, the pair $(U, \sigma_0) \in \mathcal{ST}(r, d) \times \mathcal{F}_N$ is initially admissible.*

Proof. For each $j \in \{1, \dots, s\}$, we have $U\bar{\alpha}_j \in \mathcal{P}_j$ with $\bar{\alpha}_j = U^\top \bar{\delta}_j$ by assumption. Then, $U\bar{\alpha}_j + \sigma_0(\bar{x}_j) \in \mathbb{U}^N(\bar{x}_j)$ by definition of \mathcal{P}_j . \square

2) *Comparison to a Euclidean approach:* In [9], the authors proposed an approach similar to (6). We refer to the approach in [9] as an ‘‘Euclidean approach’’, since the rigorous viewpoint of optimizing over subspaces was not pursued. We present a slightly adapted version of the Euclidean approach for comparison with our Riemannian approach, to highlight its conceptual advantages.

In the Euclidean approach, the idea is to *alternately* optimize the basis and the latent representation of points in \mathcal{D}_s and \mathcal{P}_j . Starting from an initial subspace $[U_{(0)}] \in \mathcal{GR}(r, d)$, we construct $[U_{(k+1)}] \in \mathcal{GR}(r, d)$ from $[U_{(k)}]$ in an iterative fashion. Once a subspace $[U_{(k)}]$ is available, we seek to find a latent representation $\bar{\alpha}_j^{(k)}$ for a point on $\text{span}(U_{(k)})$ which is also in \mathcal{P}_j , or at least close to it. Following our previous discussion, we again choose the point on $\text{span}(U_{(k)})$ with the smallest Euclidean distance to the maximum volume ellipsoid center $\bar{\delta}_j$ of \mathcal{P}_j . This choice leads to the latent



(a) Riemannian approach. (b) Infeasibility of the Euclidean approach.

Fig. 2: Illustration of the admissibility problem.

variables $\bar{\alpha}_j^{(k)} := U_{(k)}^\top \bar{\delta}_j$. The convex program

$$\hat{U}_{(k+1)} = \arg \min_{U \in \mathbb{R}^{d \times r}} f^{\text{eucl}}(U) \quad \text{s.t.} \quad U\bar{\alpha}_j^{(k)} \in \mathcal{P}_j \quad (7)$$

is solved, with $j \in \{1, \dots, s\}$. Then, a subspace $[U_{(k+1)}] \in \mathcal{GR}(r, d)$ can be obtained by orthogonalizing the columns of $\hat{U}_{(k+1)}$, the latent variables $\bar{\alpha}_j^{(k+1)}$ can be computed as described above and the iteration counter k can be increased. Here, f^{eucl} stands for any convex cost function; we concentrate on the feasibility of (7).

The above procedure enjoys the benefit of entirely relying on the (repetitive) solution of a convex program, for which powerful tools exist. However, we argue that our subspace perspective offers further potential *conceptually*. The key difference to (6) is that (7) relies on the latent representation $\bar{\alpha}_j^{(k)}$ of $\bar{\delta}_j$ with respect to $U_{(k)}$. We have already discussed that the latent representation loses its meaning under a change of basis. As we see next in a simple example, it is also not suited for comparing high-dimensional points among different subspaces.

Example 1. We assume $d = s = 2$, and consider the polyhedral sets $\mathcal{P}_1 = \{z \in \mathbb{R}^2 \mid (2, 0) \leq z \leq (4, 2)\}$ and $\mathcal{P}_2 = \{z \in \mathbb{R}^2 \mid (2, 4) \leq z \leq (4, 6)\}$. They have the maximum volume ellipsoidal centers $\bar{\delta}_1 = (3, 1)$ and $\bar{\delta}_2 = (3, 5)$, respectively. This is illustrated in Fig. 2b. We consider an initial subspace spanned by $U_{(0)}^\top = [1 \ 0]$, which implies $\bar{\alpha}_1^{(0)} = \bar{\alpha}_2^{(0)} = 3$. The optimization variable in (7) is $U^\top = [u_1 \ u_2]$, and for our specific initial subspace the constraints for u_2 are given by $0 \leq 3u_2 \leq 2 \wedge 4 \leq 3u_2 \leq 6$, i.e., (7) is infeasible. Note that the subspace spanned by $U^\top = \frac{1}{\sqrt{2}} [1 \ 1]$ does intersect both sets, c.f. Fig. 2b, so that $[U]$ is a possible solution to Problem 1. This fact is just not captured by the Euclidean formulation (7); it is, however, captured by the Riemannian formulation (6).

While this example is admittedly constructed, it shows the major flaw of the Euclidean approach: the latent variables $\bar{\alpha}_j^{(k)}$ *only* have a meaning in the coordinates given by $U_{(k)}$, and not necessarily in the coordinates given by a different U .

Remark 7. While (6) enjoys the aforementioned conceptual advantages, it comes at the price of having to optimize over the Riemannian manifold $\mathcal{GR}(r, d)$. This can be dealt with by using tailored algorithms [11], [12]; see also [10].




Fig. 3: Initial set and data points for pendulum example.

Further, we consider *additional* constraints of the form $P\bar{\delta}_j \in \mathcal{P}_j$, which can be tackled by applying methods from [13]. A detailed discussion, though, is out of scope here.

V. SIMULATION RESULTS

We consider the system $(\dot{x}^1, \dot{x}^2) = (x^2, x^1 + u)$, which is obtained by linearizing a simple inverted pendulum model around its upright position. Here, x^1 and x^2 describe the angular position and velocity, and u the applied torque. We discretize with sampling time $T_s = 0.1$ s to obtain a discrete-time model. We impose the state and input constraints $\mathbb{X} = \{x \in \mathbb{R}^2 \mid |x^1| \leq 1, |x^2| \leq 0.35\}$ and $\mathbb{U} = \{u \in \mathbb{R} \mid |u| \leq 1\}$. The objective is to stabilize the linearized system in the origin. Hence, we apply an MPC scheme with stage cost $\ell(x, u) = x^\top x + 0.1u^2$. The unconstrained LQR controller K_∞ is used to derive the terminal cost and terminal set. The MPT3 toolbox and YALMIP are used to formulate (2).

Using Alg. 1, we seek to design a reduced MPC scheme that features a subspace of dimension $r = 2$. As our baseline MPC, we consider $\mathcal{P}_{13}(x)$ in (2) for the prediction horizon $N = 13$. We use K_∞ to pre-stabilize the system.

We aim to achieve the same domain of attraction that the full-order MPC controller would have for a prediction horizon of $N_{\text{desired}} = 12$, i.e., we set the initial set to $\mathbb{X}_0 = \mathcal{X}_{12}$. We obtain the data set \mathcal{D} in (3) by sampling $L = 450$ states $x_i \in \mathbb{X}_0 \setminus \mathbb{X}_f$, and computing the corresponding optimal input sequences $z_i = \mu_{13}(x_i) \in \mathbb{U}^{13}(x_i)$. Fig. 3 illustrates the situation.

We determine the fixed offset $\sigma_0 \in \mathcal{F}_{13}$ as described in Sec. IV with $\omega_i = 1/L$, and apply our Riemannian framework (6) to design a subspace $[U^*] \in \mathcal{GR}(2, 13)$ such that the pair (U^*, σ_0) is initially admissible. For this task, we implement [13, Alg. 1], and the optimization problem(s) on the Grassmann manifold are solved with Manopt [18]. We note that the Euclidean counterpart in (7) was infeasible in the first iteration $k = 0$.

In Fig. 3, we report the relative difference $\varepsilon: \mathbb{X}_0 \rightarrow \mathbb{R}$ for the closed-loop cost \tilde{J}_{cl} with respect to the full-order MPC controller with $N = 12$. The mean of ε over the shown grid is 0.31%, with a standard deviation of 0.34%.

To conclude, with a two-dimensional subspace, we designed a reduced MPC scheme which uses three degrees of freedom in closed-loop to guarantee stability and feasibility. The scheme has the same domain of attraction as the full-order controller with a prediction horizon of twelve, and only a minor decrease in closed-loop performance.

VI. CONCLUSIONS

We introduce a reduced-order MPC scheme for which stability and feasibility can be established. The required subspace is computed via constrained optimization over a Riemannian manifold. The design process ensures initial feasibility of the MPC problem in the reduced space, which paves the way for the closed-loop properties.

There are many possible next developments. In this work, we have fixed all design parameters but the subspace with a heuristic, so approaches can be discussed that take all parameters – such as the offset or the pre-stabilizing feedback gain – and their interconnections into account. Moreover, the current design process does not consider closed-loop information. Reinforcement learning or Bayesian optimization are possible tools for tackling the outlined challenges.

REFERENCES

- [1] R. Findeisen and F. Allgöwer, “An Introduction to Nonlinear Model Predictive Control,” in *21st Benelux Meeting on Systems and Control*, pp. 119–141, 2002.
- [2] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control: Theory and Algorithms*. Communications and Control Engineering, Springer, 2, ed., 2017.
- [3] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.
- [4] R. Cagienard, P. Grieder, E. Kerrigan, and M. Morari, “Move blocking strategies in receding horizon control,” *J. Proc. Cont.*, vol. 17, no. 6, pp. 563–570, 2007.
- [5] R. C. Shekhar and C. Manzie, “Optimal move blocking strategies for model predictive control,” *Automatica*, vol. 61, pp. 27–34, 2015.
- [6] G. Pan and T. Faulwasser, “NMPC in active subspaces: Dimensionality reduction with recursive feasibility guarantees,” *Automatica*, vol. 147, p. 110708, 2023.
- [7] A. Bemporad and G. Cimini, “Variable elimination in model predictive control based on K-SVD and QR factorization,” *IEEE Transactions on Automatic Control*, vol. 68, no. 2, pp. 782–797, 2021.
- [8] O. J. Rojas, G. C. Goodwin, M. M. Serón, and A. Feuer, “An SVD based strategy for receding horizon control of input constrained linear systems,” *Int. J. Rob. Cont.*, vol. 14, no. 13–14, pp. 1207–1226, 2004.
- [9] G. Goebel and F. Allgöwer, “Semi-explicit mpc based on subspace clustering,” *Automatica*, vol. 83, pp. 309–316, 2017.
- [10] R. Schurig, A. Himmel, and R. Findeisen, “Toward grassmannian dimensionality reduction in mpc,” *IEEE Control Systems Letters*, vol. 7, pp. 3187–3192, 2023.
- [11] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008.
- [12] N. Boumal, *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- [13] C. Liu and N. Boumal, “Simple algorithms for optimization on riemannian manifolds with constraints,” *Applied Mathematics & Optimization*, vol. 82, pp. 949–981, 2020.
- [14] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.
- [15] A. Edelman, T. A. Arias, and S. T. Smith, “The geometry of algorithms with orthogonality constraints,” *SIAM J. Matrix Anal. & Appl.*, vol. 20, no. 2, pp. 303–353, 1998.
- [16] T. Bendokat, R. Zimmermann, and P.-A. Absil, “A Grassmann manifold handbook: Basic geometry and computational aspects,” *arXiv preprint arXiv:2011.13699*, 2020.
- [17] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [18] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, “Manopt, a Matlab toolbox for optimization on manifolds,” *Journal of Machine Learning Research*, vol. 15, no. 42, pp. 1455–1459, 2014.