Automatica 188 (2026) 112896

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Control Lyapunov function design via configuration-constrained polyhedral computing[☆]

Boris Houska [a], Matthias A. Müller [b], Mario Eduardo Villanueva [c,*]

[a] ShanghaiTech University, China
[b] Leibniz University Hannover, Germany
[c] IMT School for Advanced Studies Lucca, Italy

ARTICLE INFO

Article history:
Received 10 March 2025
Received in revised form 29 September 2025
Accepted 19 December 2025

Keywords:
Linear systems
Control Lyapunov functions
Polyhedral computing
Convex optimization
Model predictive control
Uncertain control systems
Min–max model predictive control

ABSTRACT

This paper proposes novel approaches for designing control Lyapunov functions (CLFs) for constrained linear systems. We leverage recent configuration-constrained polyhedral computing techniques to devise piecewise affine convex CLFs. Additionally, we generalize these methods to uncertain systems with both additive and multiplicative disturbances. The proposed design methods are capable of approximating the infinite horizon value function of both nominal and min–max optimal control problems by solving a single, one-stage, convex optimization problem. As such, these methods find practical applications in explicit controller design as well as in determining terminal regions and value functions for nominal and min–max model predictive control (MPC). Numerical examples illustrate the effectiveness of this approach.

1. Introduction

Control Lyapunov Functions (CLFs) have long been recognized as foundational tools in the analysis of stabilizability characteristics of linear and nonlinear systems, dating back to the pioneering works by Zubov (1965), Artstein (1983), and Primbs et al. (1999). In particular, since their inception, piecewise affine CLFs and their polyhedral domains (Gutman & Cwikel, 1987) have received significant attention due to their flexibility and effectiveness in various control applications. By now, CLFs can be considered a standard tool in control theory and a long list of methods, design strategies, and applications exist, as surveyed in Giesl and Hafstein (2015).

In practical applications, however, the mere design of CLFs for stabilizing a system often falls short of the broader objectives. Instead, one is often tasked with the more intricate challenge of synthesizing CLFs that closely approximate the infinite horizon value functions of optimal control problems. As discussed in Bertsekas (2012), such approximately optimal CLFs are needed for achieving desired performance levels. Specifically, in the realm of model predictive control (MPC), CLFs are often used as terminal cost functions. Here, the objective is twofold: first, to leverage the inherent stability guarantees of CLFs; and second, to ensure that these functions accurately approximate the infinite horizon cost, thereby optimizing the closed-loop performance of the controller. This balance of computational tractability, stability, and optimality has been a focal point of the seminal work of Chen and Allgöwer on quasi-infinite horizon MPC (Chen & Allgöwer, 1998), as surveyed and generalized by Mayne et al. (2000) and subsequent researchers, forming the basis for numerous MPC designs and implementations (Rawlings & Mayne, 2009).

Furthermore, the control literature has witnessed a proliferation of methods for designing CLFs tailored to constrained linear systems. This trend is primarily driven by the insight that infinite-horizon costs in linear or linear-quadratic optimal control problems subject to polyhedral state and control constraints are either continuous piecewise linear or continuous piecewise quadratic functions (Bemporad et al., 2003). Explicit methodologies for constructing such cost functions alongside their associated polyhedral partitions are extensively documented in the explicit MPC literature (Bemporad et al., 2002; Grieder et al., 2004). Additionally, the inherent structural properties of these systems have spurred the development of approximate approaches, leveraging parametric optimization to devise piecewise affine infinite-horizon cost approximations (Bemporad & Filippi, 2006; Bemporad et al., 2011; Xu et al., 2025).

☆ The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Bin Zhou under the direction of Editor Florian Dorfler (Optimization in Systems and Control Bin Zhou).

* Corresponding author.
E-mail addresses: borish@shanghaitech.edu.cn (B. Houska), mueller@irt.uni-hannover.de (M.A. Müller), me.villanueva@imtlucca.it (M.E. Villanueva).

https://doi.org/10.1016/j.automatica.2026.112896
0005-1098/

Drawing parallels with the approximation of infinite horizon costs in traditional optimal control problems using CLFs, the design of min–max CLF functions serves a similar purpose in the realm of min–max optimal control problems. These robust CLFs find applications in approximating the infinite horizon cost associated with such problems. Existing techniques for constructing min–max CLFs often rely on the implementation of a min–max dynamic programming recursion, as pioneered by Witsenhausen (1968). Specifically, in the field of min–max MPC, researchers have developed a diverse array of iterative approximate dynamic programming approaches to design CLFs that cater to the specific requirements of this robust control framework. Works such as those by Boyd et al. (1994), Scokaert and Mayne (1998), Bemporad et al. (2003), and Kerrigan and Maciejowski (2004) have laid the foundation for this research. This line of work was further advanced in Diehl and Björnberg (2004) as well as in Lincoln and Rantzer (2006), who developed approximate dynamic programming methods for designing piecewise affine CLFs for min–max MPC.

Beyond these iterative and dynamic programming-based methods, polyhedral computing techniques rooted in convex and non-convex optimization have been proposed for constructing piecewise affine CLFs. These techniques leverage the properties of polyhedral sets to facilitate the design of stabilizing controllers. For instance, Gutman and Cwikel (1986, 1987) have conducted extensive research on stabilizing vertex control laws, which play a pivotal role in computing polyhedral invariant sets—a concept thoroughly explored in the works of Blanchini (1994, 1995), Blanchini and Miani (1999). The gauge- or Minkowski–Lyapunov functions, of these invariant sets naturally serve as CLFs, as demonstrated in recent studies by Raković and colleagues (Raković, 2020; Raković & Zhang, 2024). However, a notable limitation of these methods arises from the structural properties of gauge functions: their inherent rigidity makes them ill-suited for approximating value functions of general constrained linear-quadratic optimal control problems. This gap in polyhedral computing for CLF design motivates the approach presented in this article, which seeks to overcome these challenges by using polyhedral computing to parameterize the epigraph of a CLF function.

Additionally, in recent years, learning-based frameworks have gained popularity, leading to the proposal of neural networks with Rectified Linear Units (ReLU) for designing piecewise affine and other types of CLF functions. These methods typically rely on directly approximating the explicit solution of finite-horizon optimal control problems or using iterative dynamic programming techniques to synthesize the CLFs, as demonstrated in recent studies (Chen et al., 2022; Fabiani & Goulart, 2022; Karg & Lucia, 2020). Furthermore, neural network-based methods for constructing piecewise quadratic approximations of infinite-horizon value functions have been proposed in He et al. (2024).

The challenge in designing piecewise affine CLFs via direct single-stage convex optimization lies in the necessity to rely on so-called double description methods, which involve enumerating the facets and vertices of the CLF's polyhedral epigraph. In this context, an impressive and early attempt was made by Jones and Morari in Jones and Morari (2010), who proposed an implicit double description method for polyhedral CLF design. However, this early attempt relied on computationally expensive enumeration methods and did not result in a fully flexible convex optimization-based design. Nevertheless, as demonstrated in the current paper, it is feasible to modify and apply a recently proposed configuration-constrained polytopic computing technique (Houska et al., 2025; Villanueva et al., 2024) to develop such convex optimization-based CLF design techniques.

1.1. Contribution

The main contribution of this paper consists of novel polyhedral computing tools for constrained linear systems and min–max infinite horizon optimal control. These tools enable the design and optimization of parameterized polyhedral epigraphs of configuration-constrained piecewise affine convex CLFs through the solution of single-stage convex optimization problems. In detail, we provide a computationally tractable and fully convex characterization of such parametric CLFs. The main theoretical results supporting this characterization are presented in Theorems 1 and 2. The main application of these CLFs is that they can be used to find approximate solutions to the HJB equations of both nominal as well as min–max infinite-horizon optimal control problems for linear systems with convex stage cost and constraints.

1.2. Overview

The structure of the paper is as follows.

- Section 2 provides a concise review of the definition and key properties of CLFs.
- Section 3 proposes a novel class of piecewise affine CLFs with configuration-constrained polyhedral epigraphs.
- Section 4 explores the application of this new CLF parameterization for approximating solutions to HJB equations.
- Section 5 presents numerical case studies.
- And, finally, Section 6 concludes the paper.

1.3. Preliminaries and notation

We use the symbol $\mathbb{R}^n_+$ to denote the non-negative orthant in $\mathbb{R}^n$. The symbol $\mathbf{1} \in \mathbb{R}^n_+$ is used to denote the vector whose entries are all equal to 1. Moreover, we adopt basic notation and definitions from the field of polytopic computing (Fukuda, 2020). Generally, a polyhedron $\mathbb{P} \subseteq \mathbb{R}^n$ is a (potentially empty) set of the form

$$\mathbb{P} = \left\{ x \in \mathbb{R}^n \,\middle|\, Fx \leq z \right\}.$$

In this context, $F \in \mathbb{R}^{m \times n}$ is called the facet matrix of $\mathbb{P}$ and $z \in \mathbb{R}^m$ the facet parameter. Moreover, the set $\left\{ x \in \mathbb{R}^n \,\middle|\, Fx \leq 0 \right\}$ is called the recession cone of $\mathbb{P}$; see Houska et al. (2025, Section 2.1). Bounded polyhedra are called polytopes, which are obtained for facet matrices F for which $Fx \leq 0$ implies $x = 0$; see Ziegler (1995) for more details.

A function $M : \mathbb{R}^{n_x} \to \mathbb{R} \cup \{\infty\}$ is called positive definite if $M(0) = 0$ and $M(x) > 0$ for all $x \neq 0$. We denote the epigraph of M by

$$\mathrm{epi}(M) := \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \,\middle|\, M(x) \leq y \right\}.$$

We call M a proper compact convex function, if its domain $\mathrm{dom}(M)$ is non-empty and compact and $\mathrm{epi}(M)$ is a closed convex set. If $v_1, \ldots, v_m \in \mathbb{R}^n$, then

$$\mathrm{conv}\,(v_1, \ldots, v_m) := \left\{ \sum_{i=1}^m \theta_i v_i \,\middle|\, \begin{array}{l} \theta_1, \ldots, \theta_m \geq 0 \\ \sum_{i=1}^m \theta_i = 1, \end{array} \right\}$$

denotes the convex hull. If m is finite, $\mathrm{conv}\,(v_1, \ldots, v_m)$ is a polytope, whose vertices are a subset of these m points. Finally, we introduce the symbol

$$K_n := \left\{ (0, \ldots, 0, \theta)^\intercal \,\middle|\, \theta \geq 0 \right\} \subseteq \mathbb{R}^n$$

to denote the proper convex cone that is generated by the last unit vector in $\mathbb{R}^n$. This notation is motivated by the fact that every proper compact convex function M that is piecewise affine on a finite polytopic partition of its domain admits a vertex representation of its epigraph. That is, for $n = n_x + 1$, we have

$$\mathrm{epi}(M) \;=\; \mathrm{conv}(v_1, \ldots, v_m) \oplus K_n \;, \tag{1}$$

where $\oplus$ denotes the Minkowski sum, see also Minkowski (1989), Weyl (1934). We refer to (1) as the vertex representation of a piecewise affine convex function. Finally, the indicator function of a closed convex set $\mathbb{X} \subseteq \mathbb{R}^n$ is denoted by $I_{\mathbb{X}}$, such that $I_{\mathbb{X}}(x) = 0$ if $x \in \mathbb{X}$ and $I_{\mathbb{X}}(x) = \infty$ otherwise.

2. Control Lyapunov functions

This section reviews the basic definition of CLFs for both deterministic as well as uncertain constrained linear systems and highlights their importance in various applications.

2.1. Constrained linear systems

This paper is concerned with constrained linear systems with n_x states and n_u controls,

$$x_{k+1} = Ax_k + Bu_k \quad \text{subject to} \quad x_k \in \mathbb{X},\; u_k \in \mathbb{U}\;. \tag{2}$$

Here, $\mathbb{X} \subseteq \mathbb{R}^{n_x}$ and $\mathbb{U} \subseteq \mathbb{R}^{n_u}$ denote, respectively, closed and convex state and control constraint sets. Moreover, we may assume, without loss of generality, that B has full column rank, as redundant controls can be eliminated.[1]

Assumption 1. The sets $\mathbb{X}$ and $\mathbb{U}$ are closed and convex and such that $(0, 0) \in \mathbb{X} \times \mathbb{U}$. Moreover, we assume that $\mathrm{rank}(B) = n_u \le n_x$.

In many applications, the above system is equipped with a convex stage cost function L, as we are interested in the so-called control performance, $\sum_{k\in\mathbb{N}} L(x_k, u_k)$, typically over an infinite time horizon. We work with stage cost functions that satisfy the following assumption.

Assumption 2. The function $L : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ is convex, non-negative, and satisfies $L(0, 0) = 0$.

Note that Assumption 2 is an assumption on the stage cost L that is satisfied in many practical applications. In particular, it is satisfied when the control objective is the stabilization of a set-point or a more general set, with a quadratic cost function being a classical choice (Rawlings & Mayne, 2009). Assumption 2 is, however, not satisfied for more general economic control objectives, where stabilization is not the main control objective. Nevertheless, in applications that satisfy a dissipativity condition and have a known storage function, it is possible to replace the economic stage cost function by an equivalent convex and positive semi-definite tracking cost (Faulwasser et al., 2018). Additionally, we note that Assumption 2 does not exclude the case $L = 0$.

[1] When we assume without loss of generality that the matrix B has full column rank by eliminating redundant controls, it is understood that the control constraint set $\mathbb{U}$ must be appropriately adjusted to reflect this reduction. Specifically, the new constraint set can be derived from the original set by considering only those components of the control vector that remain after elimination. This adjustment can be accomplished through a straightforward projection of the original constraint set onto the subspace of the reduced control vector, or by keeping the eliminated controls as auxiliary variables in our representation of the control constraint set.

Definition 1. A compact convex set $X \subseteq \mathbb{X}$ with $0 \in X$ is called λ-control invariant,[2] with $\lambda \in [0, 1]$, if

$$\forall x \in X,\; \exists u \in \mathbb{U}, \quad Ax + Bu \in \lambda X.$$

A 1-control invariant set is called control invariant (CI).

2.2. Control Lyapunov Functions

Control Lyapunov Functions (CLFs) constitute a fundamental tool for analyzing the stabilizability of control systems. In general settings, these functions are required to be merely lower semi-continuous and positive definite. However, in the context of linear-convex system theory, it is often sufficient to focus on convex CLFs. Consequently, in this article, we adhere to the following formal definition of (convex) L-CLFs.

Definition 2. A function $M : \mathbb{R}^{n_x} \to \mathbb{R} \cup \{\infty\}$ is called an L-CLF if it is positive definite, proper compact convex, and satisfies the control Lyapunov inequality (CLI)[3]

$$M(x) \ge \min_{u\in\mathbb{U}} \; L(x, u) + I_{\mathbb{X}}(x) + M(Ax + Bu) \tag{3}$$

for all $x \in \mathbb{R}^{n_x}$.

Remark 1. In contrast to other existing definitions in the literature, state constraints of the form $Ax + Bu \in \mathbb{X}$ are not explicitly enforced in the CLI. They are, however, enforced implicitly: if M satisfies the CLI, then we must have $M \ge I_{\mathbb{X}}$.

Definition 2 contains the definition of convex control invariant sets as a special case. Namely, a compact convex set is CI if and only if it is the sublevel set of a 0-CLF, see Blanchini and Miani (2015). Moreover, if the inequality in (3) is replaced with an equality, it coincides with the HJB equation of an infinite horizon optimal control problem with stage cost L, as discussed in Bertsekas (2012). More precise statements regarding the relations between Definition 2, infinite horizon optimal control, and HJBs can be found in Section 4.

2.3. Robust control Lyapunov functions for uncertain linear systems

Definition 2 extends naturally to constrained uncertain linear systems of the form

$$x_{k+1} = A_k x_k + B_k u_k + w_k \quad \text{subject to} \quad \begin{cases} x_k \in \mathbb{X} \\ u_k \in \mathbb{U}, \end{cases} \tag{4}$$

where the sequence $(A_k, B_k, w_k) \in \mathbb{D}$ is unknown. In this context, the time-invariant uncertainty set,

$$\mathbb{D} \subseteq \mathbb{R}^{n_x \times n_x} \times \mathbb{R}^{n_x \times n_u} \times \mathbb{R}^{n_x},$$

is assumed to be non-empty, compact, and convex.

Definition 3. A compact convex set $X \subseteq \mathbb{X}$ is called robust control invariant (RCI) if

$$\forall x \in X,\; \exists u \in \mathbb{U}, \forall (A, B, w) \in \mathbb{D}, \quad Ax + Bu + w \in X.$$

[2] Sometimes λ-control invariant sets with $0 \le \lambda < 1$ are also called λ-contractive sets; see, e.g., Blanchini and Miani (2015, Section 4.2.4).

[3] If Assumptions 1 and 2 hold, the minimizer on the right side of (3) exists. This is because the map

$$u \to L(x, u) + I_{\mathbb{X}}(x) + M(Ax + Bu)$$

is, for every $x \in \mathbb{R}^{n_x}$, a compact, convex, and non-negative function. Here, the compactness statement follows from the requirement that M is proper compact convex and our assumption that B has full rank.

In the following, we will work with a non-standard generalization of the concept of λ-contractiveness to uncertain control systems. The role of this definition in the ongoing developments will be explained in Section 4.3.

Definition 4. Let $X_s \subseteq \mathbb{X}$ be a compact and convex RCI set and $\lambda \in [0,1)$ be a given constant. A compact convex set $X \subseteq \mathbb{X}$ is called robust λ-contractive[4] relative to X_s if

$$\forall x \in X,\ \exists u \in \mathbb{U}, \forall (A,B,w) \in \mathbb{D},$$
$$Ax + Bu + w \in \lambda X \oplus (1-\lambda) X_s,$$

where $\oplus$ denotes the Minkowski sum.

Uncertain linear systems are also often equipped with stage cost functions. While Assumption 2 is appropriate in the nominal case, slightly different assumptions must be considered for uncertain systems. This is due to the fact that, in general, a system with additive uncertainties cannot be stabilized at a single point. Instead, the goal is to steer the system's state into a feasible RCI set $X_s \subseteq \mathbb{X}$.

Definition 5. Let $X_s \subseteq \mathbb{X}$ be a given RCI set. We say that a convex function $L : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}_+$ is positive semi-definite with respect to X_s for System (4), if

$$\forall x \in X_s,\ \exists u \in \mathbb{U} : \begin{cases} L(x,u) = 0 \quad \text{and} \\ \forall (A,B,w) \in \mathbb{D}, \\ Ax + Bu + w \in X_s. \end{cases}$$

Assumption 3. The cost function $L : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}_+$ is convex and positive semi-definite with respect to at least one RCI set $X_s \subseteq \mathbb{X}$.

Note that, as in the nominal case, Assumption 3 does not exclude the case that we set $L = 0$. This assumption does, however, ensure that the following definition of min–max CLFs makes sense.

Definition 6. A function $M : \mathbb{R}^{n_x} \to \mathbb{R}_+ \cup \{\infty\}$ is called a robust L-CLF if it is non-negative, proper compact convex, satisfies $M(x_s) = 0$ for at least one $x_s \in \mathbb{X}$, and satisfies the robust CLI[5]

$$M(x) \geq \min_{u \in \mathbb{U}} \max_{(A,B,w) \in \mathbb{D}} L(x,u) + I_{\mathbb{X}}(x) + M(Ax + Bu + w)$$

for all $x \in \mathbb{R}^{n_x}$.

In analogy to Definition 2, Definition 6 contains RCI sets as a special case. That is, a compact convex set is RCI if and only if it is the sublevel set of a robust 0-CLF. Finally, the discrete-time min–max HJB for worst case optimal control on infinite time horizons can be obtained by replacing the inequality in the robust CLI with an equality; see Bertsekas (2012), Diehl and Björnberg (2004), Witsenhausen (1968).

3. CLF construction via polytopic computing

This section introduces a necessary, sufficient, and computationally tractable condition that characterizes the set of L-CLFs whose epigraph is a configuration-constrained polyhedron.

[4] The usual notion of robust λ-contractiveness corresponds to Definition 4 with $X_s = \{0\}$, compare (Blanchini & Miani, 2015, Definition 4.19).

[5] In contrast to the nominal case, a robust L-CLF function M is not required to be positive definite. The existence of the minimum on the right side of the robust CLI is, however, still guaranteed as long as Assumptions 1 and 3 are satisfied. Similarly, the maximum exists since $\mathbb{D}$ is assumed to be compact.

3.1. Representation of epigraphs

Let $J : \mathbb{R}^{n_x} \to \mathbb{R} \cup \{\infty\}$ be a proper compact convex function. Since its epigraph is a closed convex set, it can be approximated by a polyhedron sharing the same recession cone (Dörfler, 2022; Ney & Robinson, 1995). This is equivalent to approximating J with a piecewise affine convex function M, whose domain is a polytope. One possible strategy is to first approximate dom(J) by a polytope of the form

$$\left\{ x \in \mathbb{R}^{n_x} \;\middle|\; G_1 x \leq z_1 \right\},$$

by selecting a facet matrix $G_1 \in \mathbb{R}^{\mathsf{f}_1 \times n_x}$ and a parameter $z_1 \in \mathbb{R}^{\mathsf{f}_1}$. The second step is to construct an augmented facet matrix

$$F = (G\ h) = \begin{pmatrix} G_1 & 0 \\ G_2 & h_2 \end{pmatrix} \in \mathbb{R}^{\mathsf{f} \times n} \tag{5}$$

with $n = n_x + 1$ and $\mathsf{f} = \mathsf{f}_1 + \mathsf{f}_2$. Here, G is obtained by augmenting G_1 with a matrix $G_2 \in \mathbb{R}^{\mathsf{f}_2 \times n_x}$ and a vector h whose first f_1 components are equal to zero, while the remaining f_2 components are strictly negative.

Assumption 4. The matrix $F \in \mathbb{R}^{\mathsf{f} \times n}$ satisfies (5) with a matrix $G_1 \in \mathbb{R}^{\mathsf{f}_1 \times n_x}$ such that $G_1 x \leq 0$ implies $x = 0$.[6] Moreover, $\mathsf{f}_2 > 0$ and the coefficients of $h_2 \in \mathbb{R}^{\mathsf{f}_2}$ are strictly negative, that is $h_2 < 0$.

The matrix F can be used to construct polyhedra

$$P(z) := \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \in \mathbb{R}^n \;\middle|\; Gx + hy \leq z \right\},$$

with parameter $z = (z_1^\intercal\ z_2^\intercal)^\intercal \in \mathbb{R}^{\mathsf{f}}$. Assumption 4 ensures that $P(z)$ can be used to represent the epigraph of proper compact convex functions, see also Fig. 1.

Proposition 1. *Let F satisfy Assumption 4 and let z be such that $P(z)$ is non-empty. Then, there exists a unique proper compact convex function $M : \mathbb{R}^{n_x} \to \mathbb{R} \cup \{\infty\}$ such that* $\text{epi}(M) = P(z)$.

Proof. The set $P(z)$ is, by construction, a polyhedron and thus non-empty, convex, and closed. By Assumption 4, we have $h_2 < 0$. Thus, $P(z)$ must be the epigraph of a proper closed and convex function. This function is then also unique, as the given epigraph uniquely determines this function. Additionally, since $G_1 x \leq 0$ implies $x = 0$, the domain of this function is also bounded, implying the statement of the proposition. □

Remark 2. If the conditions of Proposition 1 are satisfied, the function M can be constructed explicitly, by a direct comparison of epi(M) and $P(z)$. It is given by

$$M(x) = \begin{cases} \max\limits_{i \in \{1,\ldots,\mathsf{f}_2\}} \dfrac{(z_2 - G_2 x)_i}{(h_2)_i} & \text{if } G_1 x \leq z_1 \\ \infty & \text{otherwise.} \end{cases}$$

3.2. Configuration templates

The class of polyhedra satisfying the conditions of Proposition 1 is combinatorially large and also does not admit a jointly affine parameterization of facets and vertices (Houska et al., 2025). As such, working with general polyhedra is not entirely practical

[6] If the matrix G_1 is such that $G_1 x \leq 0$ implies $x = 0$, the set $\{x \mid G_1 x \leq z_1\}$ is bounded. This is necessary in practice, as finite piecewise affine approximations of convex functions on unbounded domains cannot be expected to be accurate.

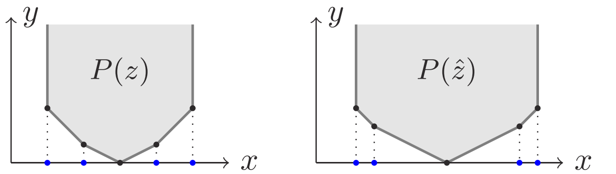

Fig. 1. Visualization of two polyhedral epigraph sets, $P(z)$ and $P(\hat{z})$, using the same facet matrix F but different parameters z and $\hat{z}$. The x-axis partitioning is implicitly determined by the parameter choice: the left epigraph yields equidistant partitioning, while the right does not. Unlike existing piecewise affine cost function representations — such as the canonical approach in Bemporad and Filippi (2006), which requires exhaustive partitioning for basis construction — we propose to optimize the polyhedral state-space partitioning by optimizing the epigraph parameter z. Note that this approach differs from less systematic randomized state-space partitioning techniques that are used in neural-network-based approximations (Xu et al., 2025).

for the development of efficient CLF design methods. Instead, we propose restricting the search to a specific set $\mathcal{P}(E)$ of polyhedra, which share the same face configuration. In detail, we propose the following definition of $\mathcal{P}(E)$, which is inspired by recent advancements in the field of polytopic computing (Villanueva et al., 2024).

Definition 7. The triplet (F, E, V) with facet matrix $F \in \mathbb{R}^{\mathsf{f}\times n}$, edge matrix $E \in \mathbb{R}^{\mathsf{e}\times\mathsf{f}}$, and vertex matrices $V_1, V_2, \ldots, V_{\mathsf{v}} \in \mathbb{R}^{n\times\mathsf{f}}$ is called a configuration triplet if

$$\left.\begin{array}{rcl} P(z) & = & \{x \in \mathbb{R}^n \mid Fx \leq z\} \text{ and} \\ P(z) & = & \mathrm{conv}\,(V_1 z, \ldots, V_{\mathsf{v}} z) \oplus K_n \end{array}\right\} \iff Ez \leq 0,$$

with $z \in \mathbb{R}^{\mathsf{f}}$. The associated set of configuration-constrained polyhedra is denoted by

$$\mathcal{P}(E) := \left\{ P(z) \,\middle|\, \exists z \in \mathbb{R}^{\mathsf{f}} : \; Ez \leq 0 \right\}.$$

The configuration constraints imposed by the matrix E in the above definition of the configuration triplet (F, E, V) are crucial as they restrict the parameter z to lie within the cone $\{z \in \mathbb{R}^{\mathsf{f}} \mid Ez \leq 0\}$. On this cone, the corresponding polyhedron $P(z)$ maintains a consistent face configuration, ensuring that it adheres to a predefined combinatorial structure, see Houska et al. (2025), Villanueva et al. (2024). This restriction serves two primary purposes. First, it guarantees the existence of a jointly affine facet and vertex representation (also known as jointly affine double description, see Houska et al. (2025, Sections 2.11–2.14)) of the polyhedron. Such a representation is essential for developing convex reformulations of CLF conditions for the epigraph of the polyhedron, see also Theorem 1 below. Second, by constraining the polyhedron to a specific configuration, we avoid the potential exponential growth in the number of vertices of general polyhedra in facet form (McMullen, 1971). In other words, configuration constraints allow us to freeze the number of vertices and, consequently, maintain full control over the vertex and facet complexity of the polyhedra. This control is particularly beneficial when managing the computational complexity of the underlying function, whose epigraph is described by the configuration-constrained polyhedron.

Remark 3. Configuration triplets can be constructed in various ways. One option is, for example, to start with a non-empty template polytope $P(\overline{z}) = \{ x \mid Fx \leq \overline{z} \}$, assumed to be simple and in minimal facet representation. In this case one can compute the vertex matrices $V_1, V_2, \ldots, V_{\mathsf{v}}$ and the associated edge matrix E as elaborated in Villanueva et al. (2024, Section 3.5). If it happens that $P(\overline{z})$ is not simple, one can add a small perturbation to the facet parameter such that $P(\overline{z})$ is simple after rescaling F, or directly use the considerations in Houska et al. (2024, Proposition 1). A survey of other efficient methods and tools for constructing such triplets can also be found in Houska et al. (2025, Sections 2.10 and 2.13). There, it is also explained that E can be constructed by enumerating the edges of the simple polyhedron $P(\overline{z})$, motivating the name "edge matrix" for E. Generally, it can be shown that the number e of rows of E is smaller than or equal to the number of edges of $P(\overline{z})$; see Houska et al. (2025, Section 2.13).

If (F, E, V) is a configuration triplet, if Assumption 4 holds, and if z satisfies $Ez \leq 0$, then the conditions of Proposition 1 are satisfied. This means that there exists a proper compact convex function M for which $\mathrm{epi}(M) = P(z)$. Moreover, under the same conditions, the set of minimizers of

$$\min_{x,y} \; y \quad \text{s.t.} \quad \begin{pmatrix} x \\ y \end{pmatrix} \in P(z) \tag{6}$$

must be a face of $P(z)$ (Ziegler, 1995). Of course, in the most basic case in which (6) has a unique minimizer, this set will simply be a vertex of $P(z)$. Thus, in this case, we can give this vertex the number 1, recalling that the vertex enumeration of $P(z) \in \mathcal{P}(E)$ does not depend on the choice of z.

Assumption 5. The triplet (F, E, V) is a configuration triplet, such that the vertex $V_1 z$ is the unique minimizer of (6) for any $z \in \mathbb{R}^{\mathsf{f}}$ with $Ez \leq 0$.

Note that the above construction and assumptions are such that the following statement holds.

Proposition 2. *Let Assumption 4 and 5 hold. Then the following statements are equivalent.*

(1) The polyhedron $P(z) \in \mathcal{P}(E)$ is the epigraph of a positive definite proper compact convex function M.
(2) We have $Ez \leq 0$ and $V_1 z = 0$.

Proof. As mentioned above, the statement of this proposition is a direct consequence of Proposition 1. Here, the constraint $Ez \leq 0$ is equivalent to the condition $P(z) \in \mathcal{P}(E)$. And, since $Ez \leq 0$ holds, the positive definiteness of the function M is equivalent to the constraint $V_1 z = 0$, since $V_1 z$ is the unique minimum of M due to Assumption 5. □

3.3. CLFs with polyhedral epigraphs

The following theorem establishes computationally tractable convex feasibility conditions that characterize the set of all L-CLF functions with bounded domain, whose epigraph is a polyhedron with a given face configuration.

Theorem 1. *Let Assumptions 1, 2 and 4, and 5 be satisfied and let the matrices R_i and s_i be such that $V_i = (R_i^\intercal, s_i)^\intercal$. Then, the following statements are equivalent.*

(1) The polyhedron $P(z) \in \mathcal{P}(E)$ is the epigraph of an L-CLF function.
(2) We have $Ez \leq 0$, $V_1 z = 0$, $R_i z \in \mathbb{X}$, and there exist control inputs $\upsilon_i \in \mathbb{U}$ and a vector $y \in \mathbb{R}^{\mathsf{v}}$ such that

$$\begin{array}{rrcl} & L(R_i z, \upsilon_i) + y_i & \leq & s_i^\intercal z \\ \text{and} & GAR_i z + GB\upsilon_i + hy_i & \leq & z \end{array}$$

hold for all $i \in \{1, \ldots, \mathsf{v}\}$.

Proof. First, Proposition 2 establishes the fact that $P(z) \in \mathcal{P}(E)$ is the epigraph of a positive definite and proper compact convex function M if and only if $Ez \le 0$ and $V_1 z = 0$. Thus, it is left to show that the remaining conditions in the second statement are equivalent to enforcing the descent condition encoded in the CLI. The proof of this statement is divided in three parts.

`Part I:` Since the right-hand expression in the CLI

$$M(x) \ge \min_{u \in \mathbb{U}} \; L(x,u) + I_{\mathbb{X}}(x) + M(Ax+Bu)$$

is convex, it is sufficient to check the descent condition at the points $R_i z$, which correspond to the projections of the vertices of $\mathrm{epi}(M)$ onto the domain of M. This implies that the CLI holds if and only if there exist control inputs $v_i \in \mathbb{U}$ with

$$M(R_i z) \;\ge\; L(R_i z, v_i) + I_{\mathbb{X}}(R_i z) + M(AR_i z + Bv_i)$$

for all vertex indices $i \in \{1,\ldots,\mathsf{v}\}$. Note that the derivation of the latter inequality can be regarded as a generalization of the key derivation step for the vertex control conditions for control invariant polytopes, which were originally introduced by Gutman and Cwikel (1986, 1987). Moreover, we use the equivalence

$$\mathrm{dom}(M) \subseteq \mathbb{X} \qquad \Longleftrightarrow \qquad \left\{ \begin{array}{l} \forall i \in \{1,\ldots,\mathsf{v}\}, \\ R_i z \in \mathbb{X}, \end{array} \right.$$

to eliminate the indicator function from the inequality above by enforcing the constraints $R_i z \in \mathbb{X}$ explicitly.

`Part II:` Since we have $V_i = (R_i^\intercal, s_i)^\intercal$, the relation $M(R_i z) = s_i^\intercal z$ holds for all z with $Ez \le 0$. Thus, together with the consideration in the first part of the proof, we find that the CLI holds if and only if there exist vectors $v_i \in \mathbb{U}$ with

$$s_i^\intercal z \;\ge\; L(R_i z, v_i) + M(AR_i z + Bv_i)$$

and $R_i z \in \mathbb{X}$ for all vertex indices $i \in \{1,\ldots,\mathsf{v}\}$.

`Part III:` It remains to use the fact that the inequality $M(AR_i z + Bv_i) \le y_i$ holds if and only if

$$(G,h)\begin{pmatrix} AR_i z + Bv_i \\ y_i \end{pmatrix} \;\le\; z.$$

Notice this is an immediate consequence of the fact that $P(z) = \mathrm{epi}(M)$. Thus, by writing out this inequality and substituting y_i in the inequality from the second part of the proof, the statement of the theorem follows. □

3.4. Robust CLFs for min-max control

The developments from the previous sections can be extended for min–max control tasks for systems with given compact and convex uncertainty set $\mathbb{D}$. To facilitate this extension, it is helpful to introduce the worst-case robust counterpart functions

$$D_{i,j}(z,v) \;:=\; \max_{(A,B,w)\in\mathbb{D}} \; G_j A R_i z + G_j B v + G_j w$$

for all $i \in \{1,\ldots,\mathsf{v}\}$ and $j \in \{1,\ldots,\mathsf{f}\}$. The functions $D_{i,j}$ are well-defined and convex, since we assume that $\mathbb{D}$ is compact and convex, recalling that the maximum over affine functions is convex. In the following, however, we simply assume that these functions are given, as they can be worked out explicitly for typical choices of $\mathbb{D}$ using known results from robust convex optimization (Bertsimas et al., 2011).

Example 1. Let us consider the practically relevant case that the uncertainty set $\mathbb{D} = \mathbb{W}_1 \times \mathbb{W}_2$ is a direct product of the polytopes

$$\mathbb{W}_1 = \mathrm{conv}((A_1,B_1),\ldots,(A_{\bar{l}},B_{\bar{l}}))$$
$$\text{and} \quad \mathbb{W}_2 = \{w \in \mathbb{R}^{n_x} \mid Gw \le \overline{w}\},$$

where $\bar{l} \in \mathbb{N}$ denotes the number of vertices of the matrix polytope $\mathbb{W}_1$, while $\mathbb{W}_2$ is given in facet representation. Then, the robust counterpart functions are given by

$$D_{i,j}(z,v) = \max_{l \in \{1,\ldots,\bar{l}\}} \; G_j A_l R_i z + G_j B_l v + \overline{w}_j$$

for all $i \in \{1,\ldots,\mathsf{v}\}$ and $j \in \{1,\ldots,\mathsf{f}\}$; see Bertsimas et al. (2011), Houska et al. (2025).

An additional consideration is that for uncertain systems it is often impossible to find a robust L-CLF that is positive definite with respect to 0. Instead, one needs to work with configuration templates that explicitly permit the set of minimizers of (6) to be a non-singleton set. Specifically, a — preferably small — RCI set. Note that this discussion is closely related to the corresponding construction in Definition 5 and Assumption 3. Consequently, Assumption 5 is, in general, not adequate. An adequate assumption must, instead, reflect the fact that one facet of $P(z)$ needs to model the zero-level RCI set, as achieved by the following construction.

Assumption 6. The triplet (F,E,V) is a configuration triplet. Moreover, the last row of the matrix $F \in \mathbb{R}^{\mathsf{f}\times n}$ is equal to the negative nth standard unit vector in $\mathbb{R}^{1\times n}$. That is, $F_{\mathsf{f}} = -e_n^\intercal$.

After replacing Assumption 5 with Assumption 6, in analogy to Proposition 2, the following statement holds.

Proposition 3. *Let Assumption 4 and 6 hold. Then, the following statements are equivalent.*

(1) The polyhedron $P(z) \in \mathcal{P}(E)$ is the epigraph of a non-negative proper compact convex function M with $M(x_{\mathrm{s}}) = 0$ for at least one $x_{\mathrm{s}} \in \mathrm{dom}(M)$.
(2) We have $Ez \le 0$ and $z_{\mathsf{f}} = 0$.

Based on this modified requirement on the choice of the triplet (F,E,V), we are now able to characterize robust L-CLFs, whose epigraphs are configuration-constrained polyhedra.

Theorem 2. *Let Assumptions 1, 3 and 4, and 6 be satisfied, let $\mathbb{D}$ be compact and convex, and let the matrices R_i and s_i be such that $V_i = (R_i^\intercal, s_i)^\intercal$. Then, the following statements are equivalent.*

(1) The polyhedron $P(z) \in \mathcal{P}(E)$ is the epigraph of a robust L-CLF.
(2) We have $Ez \le 0$, $z_{\mathsf{f}} = 0$, $R_i z \in \mathbb{X}$, and there exist control inputs $v_i \in \mathbb{U}$ and a vector $y \in \mathbb{R}^{\mathsf{v}}$ such that

$$L(R_i z, v_i) + y_i \;\le\; s_i^\intercal z \tag{7}$$
$$\text{and} \quad D_{i,j}(z,v_i) + h_j y_i \;\le\; z_j \tag{8}$$

for all $i \in \{1,\ldots,\mathsf{v}\}$ and all $j \in \{1,\ldots,\mathsf{f}\}$.

Proof. The proof of this theorem is essentially analogous to the proof of Theorem 1 after replacing the vertex constraint $V_1 z = 0$ with the facet condition $z_{\mathsf{f}} = 0$. Apart from this, one needs to modify the argument from the third part of the proof of Theorem 1, which is, however, straightforward as the inequality

$$D_{i,j}(z,v_i) + h_j y_i \;\le\; z_j$$

holds for all $i \in \{1,\ldots,\mathsf{v}\}$ and all $j \in \{1,\ldots,\mathsf{f}\}$, if and only if

$$M(AR_i z + Bv_i + w) \;\le\; y_i$$

holds for all $(A,B,w) \in \mathbb{D}$. This is by definition of $D_{i,j}$, since they are constructed as the robust counterparts of the rows of the above inequality. □

4. CLF design via convex optimization

The conditions specified in Theorems 1 and 2 are convex, enabling the application of convex optimization methods for designing L-CLFs. In theory, one could formulate arbitrary, ideally convex, design objectives.

In practical applications, however, one is often interested in identifying CLFs that closely approximate the solutions of the infinite horizon (min–max) Hamilton–Jacobi–Bellman (HJB) equation. Therefore, the current section primarily focuses on design objectives that lead to CLFs approximating such HJB solutions.

4.1. Hamilton–Jacobi-Bellman equation

The primary goal of this section is to establish conditions under which the minimizer of a particular optimization problem solves the HJB equation. This is crucial for designing CLFs that approximate solutions to the HJB equation, which in turn aids in the development of optimal control strategies. We begin by introducing the necessary mathematical framework and then proceed to state and prove the main lemma that connects the optimization problem to the HJB equation.

Let $X \subseteq \mathbb{X}$ be a compact convex λ-control invariant set with $0 \in X$. Moreover, let $\mathcal{M}(X)$ be the set of functions $M : \mathbb{R}^{n_x} \to \mathbb{R}$ such that M is an L-CLF in the sense of Definition 2 and such that $\mathrm{dom}(M) = X$.

Consider an arbitrary probability measure μ on X with full support, $\mathrm{supp}(\mu) = X$. The goal of the following lemma is to establish conditions under which the minimizer $M^\star$ of

$$\min_{M \in \mathcal{M}(X)} \int_X M \,\mathrm{d}\mu \tag{9}$$

is well-defined and solves the HJB equation on X. That is, we have

$$M^\star(x) \;=\; \min_{u \in \mathbb{U}} \; L(x,u) + I_X(x) + M^\star(Ax+Bu) \tag{10}$$

for all $x \in X$. Note that in terms of interpreting the above discrete-time HJB equation, the λ-control invariant set $X \subseteq \mathbb{X}$, not the set $\mathbb{X}$, should be regarded as our state constraint set, since we are interested in finding a CLF function, whose domain is equal to X. By definition, all functions $M \in \mathcal{M}(X)$ satisfy $M + I_X = M$. Consequently, the comments from Remark 1 also apply here—after replacing $\mathbb{X}$ with X. This explains why enforcing the state constraint $Ax + Bu \in X$ in (10) is then not necessary.

Lemma 1. *Let Assumptions 1 and 2 hold, let $X \subseteq \mathbb{X}$ be a compact convex λ-control invariant set with $0 \in X$ and $0 \le \lambda < 1$, let μ be any probability measure on X with full support and let L be positive definite. Then, the minimizer $M^\star \in \mathcal{M}(X)$ of (9) exists and is unique. Moreover, $M^\star$ satisfies the HJB equation (10) for all $x \in X$.*

Hamilton–Jacobi–Bellman equations are among the most extensively studied equations in optimal control theory, as evidenced in the literature (Bardi & Capuzzo-Dolcetta, 1997; Bellman, 1957; Bertsekas, 2012). The above mentioned relation between the optimization problem (9) and the HJB has, however, received much less attention. Hence, for the sake of maintaining the current article self-contained and given that the statement of the aforementioned lemma has — to the best of our knowledge — not been presented in this specific form, which is required in our context, we include a brief proof below.

Proof. Assumption 2 ensures that L is convex and hence Lipschitz continuous on X. Since we assume that X is λ-control invariant with $0 \le \lambda < 1$, this is already sufficient to conclude that $\mathcal{M}(X)$ is non-empty, because the system is exponentially stabilizable on X. Next, for any given CLF function $M \in \mathcal{M}(X)$, the associated finite horizon cost functions $J_N^M : \mathbb{R}^{n_x} \to \mathbb{R}_+ \cup \{\infty\}$, given by

$$J_N^M(x_0) = \min_{x,u} \sum_{k=0}^{N-1} L(x_k,u_k) + M(x_N)$$
$$\text{s.t.} \quad \begin{cases} \forall k \in \{0,\ldots,N-1\} \\ x_{k+1} = Ax_k + Bu_k \\ x_k \in X, \; u_k \in \mathbb{U}, \end{cases}$$

are well-defined for all $x_0 \in \mathbb{R}^{n_x}$ and all horizon lengths $N \in \mathbb{N}$. This follows from an inductive argument using the positive definiteness of M and the fact that a minimizer of $L(x,u)+M(Ax+Bu)$ exists in $\mathbb{U}$, see Footnote 2. Let us summarize the three most important properties of these cost functions

(1) $J_N^M \in \mathcal{M}(X)$ for all $N \in \mathbb{N}$ and all $M \in \mathcal{M}(X)$;
(2) $J_{N+1}^M \le J_N^M$ for all $N \in \mathbb{N}$, and all $M \in \mathcal{M}(X)$; and
(3) the infinite horizon cost

$$J_\infty(x) = \lim_{N \to \infty} J_N^M(x)$$

exists and is finite for all $x \in X$. Moreover, this limit does not depend on the choice of $M \in \mathcal{M}(X)$.

Note that the first two properties of the functions J_N^M are immediate consequences of Bellman's principle of optimality. Moreover, the (pointwise) existence of the infinite horizon limit follows from the second statement by applying the monotone convergence theorem. Since M is proper compact convex with $M(0) = 0$, and hence continuous on its domain, for all $M \in \mathcal{M}(X)$, the infinite horizon limit J_∞ does not depend on M. Additionally, J_∞ satisfies the HJB equation (10) (Bertsekas, 2012; Rawlings & Mayne, 2009). Together with our assumption that L is positive definite, this implies that we also have $J_\infty \in \mathcal{M}(X)$.

Now, let us assume that we can find an $M \in \mathcal{M}(X)$ for which

$$\int_X M \,\mathrm{d}\mu \;<\; \int_X J_\infty \,\mathrm{d}\mu. \tag{11}$$

In this case, it follows from the definition of $J_0^M = M$ and the monotonicity of the cost function sequence that we have

$$\int_X M \,\mathrm{d}\mu \;=\; \int_X J_0 \,\mathrm{d}\mu \;\ge\; \int_X J_\infty \,\mathrm{d}\mu \;>\; \int_X M \,\mathrm{d}\mu.$$

But this is a contradiction. Thus, J_∞ must be a minimizer of (9). The fact that J_∞ is not only a minimizer, but actually the unique minimizer of (9) follows by using a very similar argument: if there was another minimizer $M \neq J_\infty$, we can find at least one probability measure μ for which (11) holds, which then also yields a contradiction. Thus, the statement of the lemma holds. □

Remark 4. Note that the existence and uniqueness of the solution of an HJB on a given domain is typically analyzed under weaker assumptions compared to those in Lemma 1. This is due to the fact that Lemma 1 seeks to establish a condition under which the HJB equation (10) not only possesses a solution but also guarantees that this solution is an L-CLF that satisfies the requirements from Definition 2. For instance, consider the case where $L = 0$, rendering $M = I_X$ a trivial solution to the HJB equation (10). However, if $X \neq \{0\}$, this function is not positive definite and, consequently, does not qualify as an L-CLF. Such scenarios are excluded in Lemma 1 by imposing the sufficient condition that L is positive definite.

In summary, Lemma 1 establishes that under the given assumptions, the minimizer of the optimization problem (9) exists, is unique, and satisfies the HJB equation on the set X. This result is foundational for the subsequent development of CLF design methods via convex optimization, as it guarantees that the designed CLFs will approximate solutions to the HJB equation.

4.2. Optimal CLFs with polyhedral epigraph

This section focuses on constructing approximate solutions to the optimization problem (9) by minimizing a suitable objective over the set of L-CLFs whose epigraph is a configuration-constrained polyhedron. We leverage the conditions from Theorem 1 to ensure computational tractability. The goal is to develop a systematic method for designing CLFs that approximate the minimizer of (9) while maintaining control over the complexity of the resulting polyhedra.

In general, the idea is to solve optimization problems of the form

$$\min_{z,y,v} \sigma(z) \quad \text{s.t.} \quad \begin{cases} \forall i \in \{1,\ldots,\mathsf{v}\}, \\ L(R_i z, v_i) + y_i \leq s_i^\intercal z \\ G_1 A R_i z + G_1 B v_i \leq \lambda z_1 \\ G_2 A R_i z + G_2 B v_i + h_2 y_i \leq z_2 \\ E z \leq 0,\ V_1 z = 0,\ v_i \in \mathbb{U},\ R_i z \in \mathbb{X}, \end{cases} \tag{12}$$

where $\sigma : \mathbb{R}^{\mathsf{f}} \to \mathbb{R}$ denotes an L-CLF design objective and $\lambda \in [0,1]$ a contraction parameter. If σ is chosen to be a convex function, Problem (12) is a convex optimization problem with $\mathsf{f}+\mathsf{v}(1+n_u)$ variables and $\mathsf{v}(1+\mathsf{f}+n_{\mathbb{X}}+n_{\mathbb{U}})+\mathsf{e}+n_x$ constraints. Here, $n_{\mathbb{X}}$ and $n_{\mathbb{U}}$ denote, respectively, the number of convex constraints needed to represent the state and control constraints, $x \in \mathbb{X}$ and $u \in \mathbb{U}$.

Note that, a relatively simple and effective heuristic for choosing σ is to set

$$\sigma(z) = c^\intercal z,$$

where $c \in \mathbb{R}^{\mathsf{f}}$ is a strictly negative weight vector, with $c < 0$. Maximizing the facet parameters z_i, each with weight $-c_i > 0$, amounts to one way of maximizing some measure of the size of the epigraph $P(z) \in \mathcal{P}(E)$ of the L-CLF that is represented by these constraints.

There are, however, a couple of details that require a more careful discussion. First of all, if one wishes to interpret the solution $P(z)$ of (12) as the epigraph of an L-CLF M that approximates the minimizer of (9), one needs to keep in mind that the statement of Lemma 1 in general only holds on λ-control invariant sets X with $0 \leq \lambda < 1$. In this context, the constraints

$$G_1 A R_i z + G_1 B v_i \leq \lambda z_1 \tag{13}$$

in (12) enforce the domain of the L-CLF M,

$$\mathrm{dom}(M) = X := \left\{ x \in \mathbb{X} \,\middle|\, G_1 x \leq z_1 \right\},$$

to be a λ-control invariant polytope. Then, a possible variant of (12) can be obtained by first freezing z_1 for which (13) is feasible and, ideally, for which $\mathrm{dom}(M)$ is a large set and then minimize over v and z_2 only. In this case, the domain X is frozen and it can then indeed be justified to minimize $\sigma(z) = c_2^\intercal z$ for an arbitrary strictly negative weight, for example, $c_2 = -\mathbf{1}$. This is because, in this case, the goal is to approximate the minimizer of (9), which is known to be independent of the choice of the probability measure μ.

Finally, a computationally more demanding but also significantly more systematic method for constructing a suitable CLF design objective σ proceeds by pre-computing the optimal values

$$\zeta_i^N = \max_{x,u} G_i x_0 + \sum_{k=0}^{N-1} h_i L(x_k,u_k) + h_i \overline{M}(x_N) \quad \text{s.t.} \quad \begin{cases} \forall k \in \{0,\ldots,N-1\}, \\ x_{k+1} = A x_k + B u_k, \\ x_k \in \mathbb{X},\ u_k \in \mathbb{U}, x_N \in \mathbb{X} \end{cases} \tag{14}$$

for all $i \in \{1,\ldots,\mathsf{f}\}$ and a large but finite time horizon N. Here, $0 \leq \overline{M} \leq M^\star$ denotes a suitable convex under-estimator of $M^\star$. For example, we can set $\overline{M} = 0$, or, if the stage cost L is quadratic, we can choose $\overline{M}$ as the quadratic cost function of the corresponding unconstrained infinite-horizon Linear Quadratic Regulator (LQR), which also underestimates $M^\star$. Since we assume that $h_i < 0$ for all $i \in \{\mathsf{f}_1+1,\ldots,\mathsf{f}\}$–compare Assumption 4, the above maximization problems are convex optimization problems that can be solved efficiently, even if N is relatively large. Moreover, under the rather mild additional assumption that $\mathbb{X}$ and $\mathbb{U}$ are bounded, the above maxima are bounded, too. The corresponding polyhedron $P(\zeta^N)$ is then, by construction, for all $N \in \mathbb{N}$, an outer approximation of the epigraph of the exact infinite horizon cost function $M^\star$ that solves the HJB under the assumption of Lemma 1. Since any feasible point z of (12) yields an inner approximation, $P(z)$,

$$P(\zeta^N) \supseteq \mathrm{epi}(M^\star) \supseteq P(z),$$

one can use the function

$$\sigma(z) = \left\| \mathsf{W}(z - \zeta^N) \right\|_\infty ,$$

with an invertible weight matrix $\mathsf{W} \in \mathbb{R}^{\mathsf{f}\times\mathsf{f}}$, as design objective in order to directly minimize the (weighted) maximum distance between the upper and the lower bound on the exact infinite horizon cost $M^\star$ in the given directions (G,h).

4.3. Approximate solutions to min-max HJBs

In this section, we extend the previous considerations to uncertain constrained linear control systems. The goal is to establish conditions under which the minimizer of an optimization problem involving robust CLFs solves the min–max HJB equation. This extension is crucial for developing optimal control strategies in the presence of uncertainties. Specifically, it turns out that an extension of the considerations from the previous sections for uncertain constrained linear control systems is possible by replacing the conditions from Theorem 1 by the corresponding conditions of Theorem 2.

In order to generalize the corresponding constructions from the previous section, however, we first need to choose a compact and convex RCI set $X_{\mathrm{s}} \subseteq \mathbb{X}$ and then choose a stage cost function L that satisfies the positive semi-definiteness requirements from Assumption 3 for this choice of X_{s}. Moreover, we assume that X is a compact and convex set that is robust λ-contractive relative to the same RCI set X_{s}, with $\lambda \in [0,1)$; see Definition 4. Finally, in analogy to the nominal case, we use the symbol $\widehat{\mathcal{M}}(X)$ to denote the set of robust L-CLFs with domain X.

The following lemma establishes conditions under which the optimization problem

$$\min_{M \in \widehat{\mathcal{M}}(X)} \int_X M \,\mathrm{d}\mu \tag{15}$$

admits a unique minimizer $M^\star$ that solves the min–max HJB equation on X,

$M(x) =$

$$\min_{u\in\mathbb{U}} \max_{(A,B,w)\in\mathbb{D}} L(x,u)+I_X(x)+M(Ax+Bu+w) \tag{16}$$

for all $x \in X$. As for the nominal case, the set X should here be regarded as a state constraint set. It needs to be taken into account in the HJB equation, as we are interested in designing CLFs with domain X.

Lemma 2. *Let $X_s \subseteq \mathbb{X}$ be a given compact and convex RCI set, let Assumptions 1 and 3 hold for this choice of X_s, and let $X \subseteq \mathbb{X}$ be a compact and convex set that is robust λ-contractive with respect to X_s. Then the minimizer $M^\star \in \widehat{\mathcal{M}}(X)$ of (15) exists, is unique, and satisfies the min–max HJB equation (16) for all $x \in X$. This statement holds for any probability measure μ on X with full support.*

Proof. The proof of this lemma is largely analogous to the proof of Lemma 1. Nevertheless, one difference is that, for reasons that have been discussed in Section 2.3, robust CLF functions are not required to be positive definite. As such, the positive semi-definiteness requirement on L in Assumption 3 is in this case already sufficient for our purposes. In detail, since L is convex, it is also Lipschitz continuous. Moreover, as X is robust λ-contractive, there exists for any $x_0 \in X$ a sequence of control laws $\kappa_k : \mathbb{X} \to \mathbb{U}$ that keeps the state x_k of the uncertain closed-loop control system

$$x_{k+1} = Ax_k + B\kappa_k(x_k) + w_k$$

in the sets $X_k = \lambda^k X \oplus (1-\lambda^k)X_s$. For example, if $\kappa : \mathbb{X} \to \mathbb{U}$ denotes a control law that can be used to steer the system safely from X to $\lambda X \oplus (1-\lambda)X_s$ and if $\kappa_s : \mathbb{X} \to \mathbb{U}$ can be used to keep the system's state within X_s, then $\kappa_k = \lambda^k\kappa + (1-\lambda^k)\kappa_s$ is such a sequence of control laws. This statement is a direct consequence of the condition from Definition 4. This means that we have $x_k \in X_k$ for all $k \in \mathbb{N}$, independent of the realization of the uncertainty sequence. In other words, the distance of the system's state to the target set X_s is robustly exponentially stabilizable on X, which implies — together with Lipschitz continuity and positive semi-definiteness of L — that $\widehat{\mathcal{M}}(X)$ is non-empty. In this context, another principal difference to the nominal case is that we define the cost functions J_N^M by solving min–max (see e.g. Kerrigan and Maciejowski (2004), Scokaert and Mayne (1998)) instead of nominal optimal control problems. In detail, we formulate the min–max optimal control problem

$$J_N^M(x_0) = \min_{x,u} \max_{\ell\in\mathcal{I}_N} \sum_{k=0}^{N-1} L(x_k^\ell,u_k^\ell) + M(x_N^\ell)$$
$$\text{s.t.} \begin{cases} \forall k \in \{0,\ldots,N-1\}, \forall \ell \in \mathcal{I}_N, \\ x_{k+1}^\ell = A_k^\ell x_k^\ell + B_k^\ell u_{j(k,\ell)} + w_k^\ell, \\ x_k^\ell \in X,\ u_{j(k,\ell)} \in \mathbb{U},\ x_0^\ell = x_0 \ . \end{cases} \tag{17}$$

Here, ℓ indexes admissible uncertainty sequences, such that we have

$$\bigcup_{\ell\in\mathcal{I}_N} ((A_0^\ell,B_0^\ell,w_0^\ell),\ldots,(A_{N-1}^\ell,B_{N-1}^\ell,w_{N-1}^\ell)) \;=\; \mathbb{D}^N,$$

where $\mathcal{I}_N$ denotes the (possibly uncountable) set of all such indices. Moreover, the control vectors are enumerated by the index $j(k,\ell) \in \mathbb{N}$ with $j(k,\ell) = j(k,i)$ whenever

$$((A_0^\ell,B_0^\ell,w_0^\ell),\ldots,(A_{k-1}^\ell,B_{k-1}^\ell,w_{k-1}^\ell))$$
$$= ((A_0^i,B_0^i,w_0^i),\ldots,(A_{k-1}^i,B_{k-1}^i,w_{k-1}^i))$$

and $j(k,\ell) \neq j(k,i)$ otherwise, for all $i,\ell \in \mathcal{I}_N$ with $i \neq \ell$. The index notation in (17) was introduced in Scokaert and Mayne (1998) in the context of min–max model predictive control. In this context, the optimization variables, $x_k^\bullet$ and $u_\bullet$ are functions on $\mathcal{I}_N$, rendering (17) a non-trivial functional optimization problem. Despite this complication, the existence of a min–max solution can be guaranteed under the listed assumptions. This can be proven using a standard min–max dynamic programming argument, see Bertsekas (2012). Namely, by a simple induction over N, it follows that the functions J_N^M are compact and convex. Consequently, as B has full-rank, the min–max dynamic programming recursion for J_N^M can be used to establish the existence of min–max solutions. Thus, the finite-horizon cost functions J_N^M are, as in the nominal case, well defined for all $x_0 \in \mathbb{R}^{n_x}$ and all horizon lengths $N \in \mathbb{N}$.

The argument of the proof of Lemma 1 can now be recovered step by step as neither Bellman's principle of optimality nor the monotonicity argument used are affected by this change of definition, implying that the limit J_∞ exists, does not depend on M, and satisfies the min–max HJB equation (Diehl & Björnberg, 2004). The last contradiction step can be applied analogously, establishing the statement of the lemma. □

4.4. Optimal robust CLFs with polyhedral epigraph

This section generalizes the CLF design procedures from the previous sections to the min–max problem setting. For this aim, we may assume that a configuration-constrained RCI polytope,

$$X_s \;=\; \{x \in \mathbb{R}^{n_x} \mid G_1 x \le z^s\} \subseteq \mathbb{X}, \tag{18}$$

with $Ez^s \le 0$, is given. Methods for computing such configuration-constrained polytopic target sets are surveyed in Houska et al. (2025). Next, an associated configuration-constrained polyhedral robust CLF approximation of the solution of the min–max HJB (16) can be found by solving the convex optimization problem

$$\min_{z,y,v} \sigma(z)$$
$$\text{s.t.} \begin{cases} \forall i \in \{1,\ldots,\mathsf{v}\}, \\ L(R_i z, v_i) + y_i \le s_i^\intercal z, \\ Ez \le 0,\ \ v_i \in \mathbb{U},\ \ R_i z \in \mathbb{X},\ \ z_\mathsf{f} = 0, \\ \forall j \in \{1,\ldots,\mathsf{f}_1\}, \\ D_{i,j}(z,v_i) \le \lambda z_j + (1-\lambda)z_j^s, \\ \forall j \in \{\mathsf{f}_1+1,\ldots,\mathsf{f}_1+\mathsf{f}_2\}, \\ D_{i,j}(z,v_i) + h_j y_i \le z_j. \end{cases} \tag{19}$$

The constraints $D_{i,j}(z,v_i) \le \lambda z_j + (1-\lambda)z_j^s$ enforce the λ-contractivity condition from Definition 4, as elaborated in Houska et al. (2024). Similar to the nominal case, the above problem has $\mathsf{f} + \mathsf{v}(1+n_u)$ optimization variables. In general, however, computing CLFs for min–max problems is more expensive than in the nominal case, as the number of convex inequality constraints in (19) is equal to $\mathsf{v}(1+\mathsf{f}n_\mathbb{D}+n_\mathbb{U}+n_\mathbb{X})+\mathsf{e}+\mathsf{f}$. Here, $n_\mathbb{D}$ denotes the number of constraints needed to evaluate the functions $D_{i,j}$. For example, if our uncertainty set $\mathbb{D}$ is chosen as in Example 1, we have $n_\mathbb{D} = \bar{l}$, recalling that $\bar{l}$ denotes the number of vertices of the matrix polytope.

The discussion regarding the design objective σ parallels the one in the previous section. The primary difference lies when one considers minimizing the weighted maximum distance, $\sigma(z) = \|W(z-\zeta^N)\|_\infty$ between the epigraph $P(z)$ and an outer approximation $P(\zeta^N)$ of the epigraph of an optimal L-CLF that solves the min–max HJB equation. In such a case, the components of ζ^N are computed as

$$\zeta_i^N = \max_{x,u} \min_{\ell\in\mathcal{I}_N} G_i x_0^\ell + \sum_{k=0}^{N-1} h_i L(x_k^\ell,u_k^\ell) + h_i\overline{M}(x_N^\ell)$$
$$\text{s.t.} \begin{cases} \forall k \in \{0,\ldots,N-1\}, \forall \ell \in \mathcal{I}_N, \\ x_{k+1}^\ell = A_k^\ell x_k^\ell + B_k^\ell u_{j(k,\ell)} + w_k^\ell \\ x_k^\ell \in \mathbb{X},\ u_{j(k,\ell)} \in \mathbb{U},\ x_N^\ell \in \mathbb{X}\ , \end{cases} \tag{20}$$

where we use the same notation as in the proof of Lemma 2 and $0 \leq \overline{M} \leq M^\star = J_\infty^0$ is a convex under-estimator of the optimal infinite horizon cost. Unfortunately, this optimization problem is rarely ever tractable, as $\mathcal{I}_N$ is typically uncountable. An exception is the case where $\mathbb{D}$ is a polytope with given vertices, where it is sufficient to enumerate the extreme scenarios, as elaborated in Kerrigan and Maciejowski (2004). Nevertheless, even in this case solving (20) remains computationally intractable for larger N due to an exponentially large number of uncertainty scenarios. Thus, in summary, while Problem (19) is computationally tractable for a broad class of practical min–max optimal control problems and can serve as a viable means to obtain L-CLF approximations of the optimal min–max infinite horizon cost, achieving an exact computational verification of the precision of these approximations remains challenging.

5. Numerical illustration

Throughout this section, we present numerical results for one nominal and for one min–max test problem, which are used to illustrate the L-CLF design techniques in this paper.

5.1. First test problem

Our first test problem corresponds to a modified version of a standard benchmark problem for robust control design that has originally been proposed in Goulart et al. (2006). In the nominal setting, using the same notation as in (2), the system matrices are given by

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} \frac{1}{2} \\ 1 \end{pmatrix} . \tag{21}$$

The state and control constraint sets are given by

$$\mathbb{X} = [-1, 2]^2 \quad \text{and} \quad \mathbb{U} = \left[-\frac{1}{2}, \frac{1}{2} \right].$$

Moreover, the stage cost function $L : \mathbb{R}^2 \times \mathbb{R} \to \mathbb{R}$ is given by $L(x,u) = \|x\|_{\mathsf{Q}}^2 + \|u\|_{\mathsf{R}}^2$ with weights

$$\mathsf{Q} = \begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{10} \end{pmatrix} \quad \text{and} \quad \mathsf{R} = \frac{1}{10} . \tag{22}$$

Note that this problem setting is used below to illustrate the proposed nominal L-CLF design methods.

5.2. Second test problem

Our second test problem considers an uncertain system of the form (4) with $\mathbb{D} = \{(A,B)\} \times \mathbb{W}$. Here, the matrices A and B are not uncertain, but given by (21). Moreover, the uncertainty set $\mathbb{W}$ of the additive process noise is set to

$$\mathbb{W} = \left\{ \begin{pmatrix} 1 \\ 1 \end{pmatrix} \omega \;\middle|\; \omega \in \left[-\frac{1}{40}, \frac{1}{40} \right] \right\}.$$

Finally, in order to construct a suitable stage cost function, we first compute a polyhedral approximation X_{s} of the minimal robust positive invariant set for the closed loop controlled system $x_{k+1} = (A+BK)x_k + w_k$, with

$$K = -[0.895, 1.367] \quad \text{such that} \quad (A+BK)X_{\mathrm{s}} + \mathbb{W} \subseteq X_{\mathrm{s}}.$$

A configuration-constrained RCI polytope X_{s} of the form (18) can be constructed by solving a single linear program, as explained in Villanueva et al. (2024, Remark 5). A more thorough discussion on the construction of configuration-constrained RCI polytopes can also be found in Houska et al. (2025, Section 4.4). The final set-regulation stage cost function $L : \mathbb{R}^2 \times \mathbb{R} \to \mathbb{R}$ is then defined by

$$L(x,u) = \min_{\xi \in X_{\mathrm{s}}} \|x - \xi\|_{\mathsf{Q}}^2 + \|u - Kx\|_{\mathsf{R}}^2 ,$$

where Q and R are selected as in (22). Note that this setting will be used below to illustrate the proposed robust L-CLF design methods.

5.3. Construction of configuration templates

In order to set up the proposed CLF design optimization problems, we first need to choose a configuration triplet (F,E,V). As discussed in Remark 3, if a facet matrix F is given, such that Assumption 4 is satisfied, an associated triplet (F,E,V) can be constructed by enumerating the edges and vertices of a selected non-empty polyhedron $P(\overline{z}) = \{x \mid Fx \leq \overline{z}\}$ and then apply the existing methods from Houska et al. (2025), Villanueva et al. (2024) to compute E and V. We discuss several strategies for constructing F and $\overline{z}$, all of which are heuristic:

- (S1) The most basic strategy proceeds by selecting the first f_1 rows of F from the boundary of the $(n-1)$-hemisphere $\left\{ c \in \mathbb{R}^n \;\middle|\; \|c\|_2 = 1,\ c_n \leq 0 \right\}$ and the remaining f_2 rows from its interior, in such a way that Assumption 4 holds. Now, one option is to set $\overline{z} = \mathbf{1}$.
- (S2) Another option is to select F as above but then set $\overline{z} = \zeta^N$. Here, we compute $\zeta_1^N, \ldots, \zeta_{\mathsf{f}}^N$ by solving (14) or (20). This has the advantage that the outer approximation $P(\zeta^N) \supseteq \mathrm{epi}(M^\star)$ defines the configuration template.
- (S3) Note that strategies for selecting the first f_1 rows of F and strategies for pre-computing a (robust) λ-control invariant polytope X are discussed in Houska et al. (2025, Section 3.3 and 3.5)). Based on this preparation, one possible strategy for selecting the remaining f_2 rows of F and $\overline{z}$ proceeds by evaluating the finite horizon cost functions $J_N^{\overline{M}}$ at selected sample points $\xi_1, \ldots, \xi_{\mathsf{f}_2} \in X$. One can then choose F and $\overline{z}$ such that the polyhedron $P(\overline{z})$ corresponds to the epigraph of a continuous piecewise affine interpolation of the points

  $$(\xi_1, J_N^{\overline{M}}(\xi_1)), \ldots, (\xi_{\mathsf{f}_2}, J_N^{\overline{M}}(\xi_{\mathsf{f}_2})).$$

  Different methods for computing such a piecewise affine interpolation are discussed in Boyd and Vandenberghe (2004, Section 6.5.5) as well as in Bemporad and Filippi (2006, Proposition 3.2).

Note that the first strategy, which sets $\overline{z} = \mathbf{1}$, results in a generic configuration template that may not always be adequate for approximating optimal value functions. In contrast, the second and third strategies guarantee that the templates can capture at least some approximation of $J_N^{\overline{M}} \approx M^\star$. However, it should be clear that all strategies are heuristics whose effectiveness is problem-dependent.

5.4. Nominal L-CLF

Our first numerical illustration concerns the nominal problem setting from Section 5.1. We apply Strategy S3) from the previous section to construct a facet matrix with $\mathsf{f}_1 = 8$ and $\mathsf{f}_2 = 37$ by interpolating sampling points on a pre-computed control invariant set, with $N = 5$ and with $\overline{M}$ being the optimal LQR cost function for the unconstrained problem. This leads to a configuration triplet (F,E,V) with $\mathsf{f} = 45$, $\mathsf{v} = 88$, and $\mathsf{e} = 132$.

Fig. 2 shows $P(z^\star)$, the epigraph of the optimal L-CLF M for the chosen template. Its facet parameter $z^\star$ was computed by solving the convex quadratically constrained optimization problem

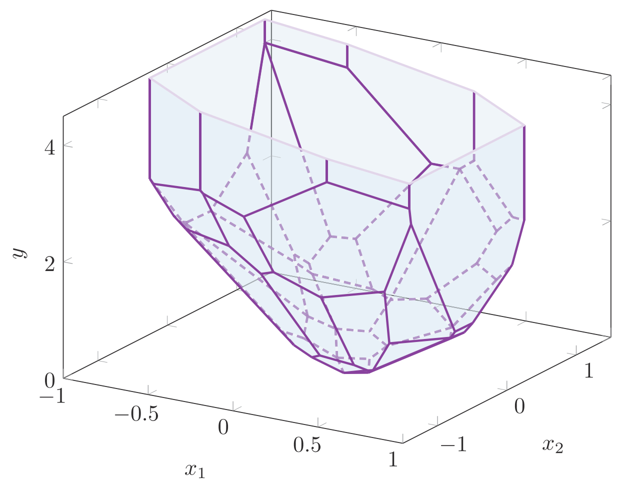

Fig. 2. Polyhedral epigraph $P(z^\star)$ of the L-CLF function M.

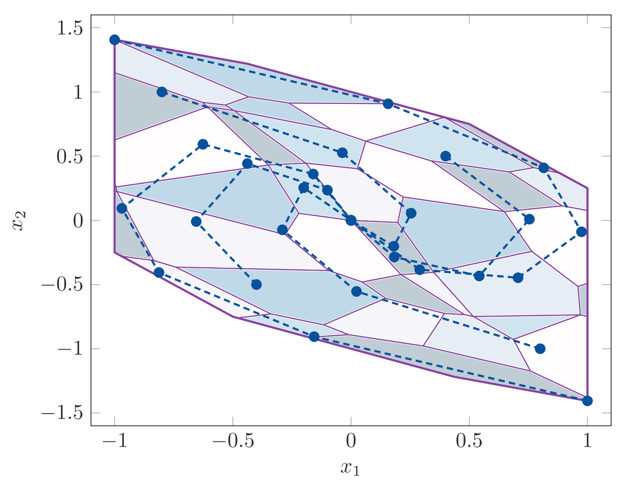

Fig. 3. Trajectories of the closed-loop system under μ_M and polyhedral partition of $\mathrm{dom}(M)$.

from (12) with 221 decision variables and 4710 constraints, using the CLF design objective

$$\sigma(z) = \left\| \mathsf{W}\left(z - \zeta^5\right)\right\|_\infty \quad \text{with} \quad \mathsf{W} = \begin{pmatrix} 100 \cdot \mathbb{1}_8 & 0 \\ 0 & \mathbb{1}_{37} \end{pmatrix}$$

with ζ^5 computed as in (14) with $\overline{M}$ defined as above. Here, $\mathbb{1}_8$ and $\mathbb{1}_{37}$ denote the unit matrices of sizes 8×8 and 37×37, respectively. The optimization problem is sparse. For example, the vertex matrices have only between 7 and 9 (out of 135) non-zero elements while the edge matrix has only 515 (out of 5940) non-zero elements. In fact, when aggregating all linear constraints, approximately 99.5% of its entries are zeros.

Fig. 3 shows the polyhedral partition of $\mathrm{dom}(M)$ together with selected trajectories of the closed-loop system under the feedback law $\mu_M : \mathrm{dom}(M) \to \mathbb{U}$, given by

$$\mu_M(x) \in \underset{u \in \mathbb{U}}{\mathrm{argmin}} \; L(x,u) + M(Ax+Bu) \; . \tag{23}$$

The feedback law μ_M can be evaluated directly, without solving the above problem, since we have $\mu(R_i z^\star) = u_i^\star$ for all $i \in \{1, \ldots, \mathsf{v}\}$. Thus $\mu(x)$ can be evaluated by solving a point location problem for x followed by a linear interpolation of the optimal vertex control inputs $u_i^\star$ for the corresponding region.

Notice that the construction above is of similar nature as explicit model predictive controllers (Herceg et al., 2013). In particular, (23) provides an approximation to the optimal, infinite horizon, feedback law whose complexity, unlike that of an exact explicit MPC controller, can be bounded a priori and does not directly depend on the prediction horizon.

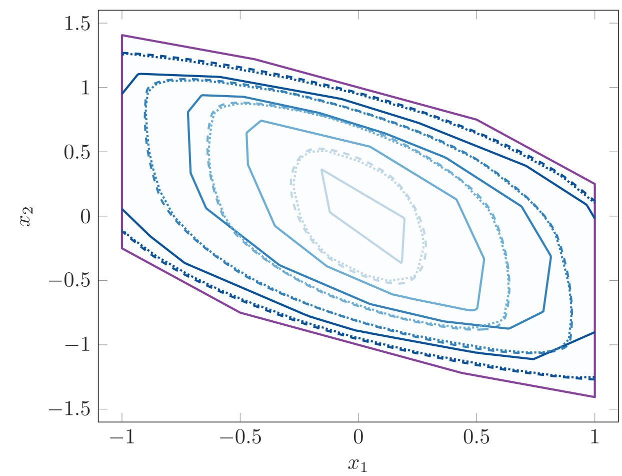

Fig. 4. (0.1, 0.5, 1, and 2)-Contours of M (continuous lines), M' (dotted lines), and $M^\star$ (dashed lines).

Finally, Fig. 4 shows the (0.1, 0.5, 1, 2)-contours of the function M (solid lines) and the function $M^\star$ solving the HJB equation (dashed lines). By construction, each sublevel set of M is contained in the same sublevel set of $M^\star$ since $M(x) \geq M^\star(x)$ for all $x \in \mathrm{dom}(M)$. The maximum approximation error relative to the maximum value $M^\star_{\max}$ of $M^\star$ on its domain satisfies

$$\frac{1}{M^\star_{\max}} \left(\max_{x \in \mathrm{dom}(M)} |M(x) - M^\star(x)| \right) \approx 0.41 \; .$$

This error can be reduced, for example, by refining the polyhedral template. As an illustration, a function M' was constructed by using a different configuration triple, which was generated by applying Strategy S1) and randomly selected directions from the hemisphere. The dimensions of the corresponding random configuration triplet are $\mathsf{f} = 469$, $\mathsf{e} = 1401$, and $\mathsf{v} = 934$. Notice that the contour lines (dotted lines) of M' approximately match those of $M^\star$. The maximum relative error of this approximation is approximately 0.08.

5.5. Robust L-CLF

This section illustrates the design of robust CLF functions based on the problem setting from Section 5.2. For the sake of illustration, we apply Strategy S2) in this instance. In detail, this strategy is used to construct an initial facet matrix F, here with $\mathsf{f}_1 = 8$ and $\mathsf{f}_2 = 83$. Moreover, we have set $N = 7$ and chosen $\overline{M}$ to be the value function of the unconstrained LQR corresponding to the nominal system in order to solve (20). Finally, this construction leads to a configuration triplet (F, E, V) with $\mathsf{f} = 91$, $\mathsf{v} = 180$, and $\mathsf{e} = 270$.

Fig. 5 depicts the polyhedral epigraph $P(z^\star)$ of the optimal robust L-CLF M. The corresponding optimal facet parameter $z^\star$ was computed by solving the convex quadratically constrained program (19), with 451 variables and 18002 constraints. As in the nominal setting, this problem is very sparse, too. For instance, the linear constraints are encoded in a matrix with approximately 99.7% zeros.

Fig. 6 shows the polyhedral partition of $\mathrm{dom}(M)$ and selected trajectories of the uncertain closed loop system under the feedback law

$$\mu_M(x) \in \underset{u \in \mathbb{U}}{\mathrm{argmin}} \; \max_{w \in \mathbb{W}} \; L(x,u) + M(Ax+Bu+w) \; .$$

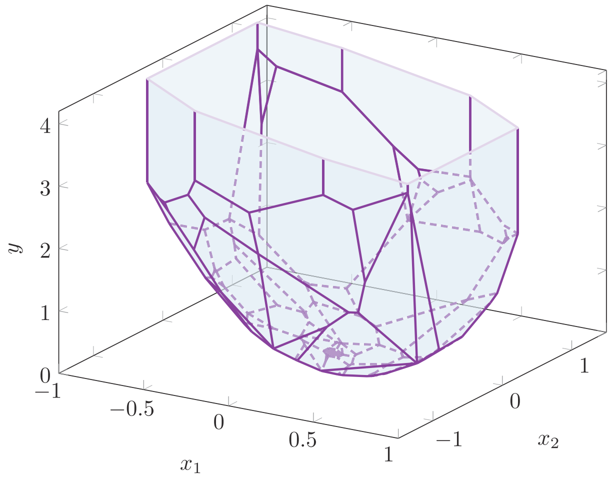

Fig. 5. Polyhedral epigraph $P(z^\star)$ of M, the robust L-CLF function. The dark (purple) facet satisfies $z_f^\star = 0$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

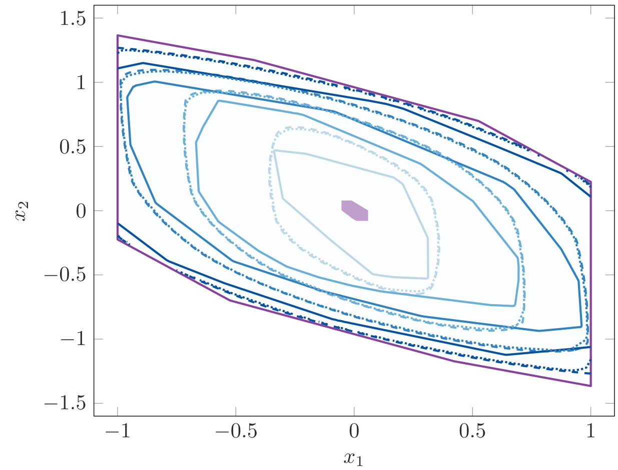

Fig. 7. (0.1, 0.5, 1, and 2)-Contours of M (continuous lines), M' (dotted lines), and $M^\star$ (dashed lines). The 0-sublevel set corresponds, for all functions, to X_s.

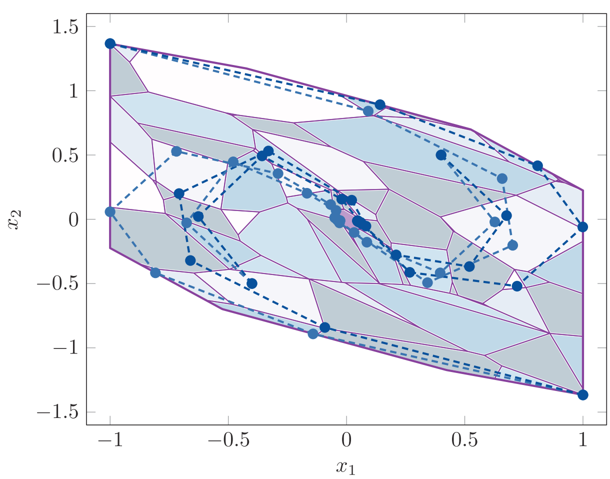

Fig. 6. Trajectories of the closed-loop system under μ_M and polyhedral partition of dom(M). In (dark) purple, the set X_s. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

As explained in the previous section, the feedback law can be evaluated based on the optimal vertex control inputs $v_i^\star$ from (19). For each initial point, two trajectories are plotted corresponding to the extreme (but potentially not worst-case) uncertainty scenarios $w_k = \pm\frac{1}{40}$ for all $k \in \mathbb{N}$. Notice that, as predicted, all trajectories remain in dom(M) and converge to X_s. As in the nominal case, the above construction provides a means to construct an explicit controller with fixed complexity that approximates the optimal min–max infinite horizon optimal control law.

Finally, Fig. 7 shows selected contours of M, a very accurate numerical approximation of $M^\star$, and a second, finer, polyhedral approximation M'. This last approximation was constructed using a configuration triple with $\mathsf{f} = 914$, $\mathsf{e} = 2487$, and $\mathsf{v} = 1658$, based on a random matrix F; see Strategy S1). Here, the relative approximation error of M is approximately 0.47, while that of M' is approximately 0.11.

5.6. Comparison with explicit MPC

In the preceding sections, we have demonstrated the application of our proposed CLF design method to construct approximations of the optimal value function for various constrained linear-quadratic optimal control problems. As these problems can, in principle, be solved exactly using parametric optimization techniques developed within the Explicit MPC framework (Bemporad et al., 2002; Grieder et al., 2004), a natural question arises: how does our proposed method compare to these established approaches?

To address this, we conduct a comparative analysis using two settings derived from the examples in Sections 5.1 and 5.4. In the first setting, we consider state constraints $\mathbb{X} = [-25, 25] \times [-5, 5]$ and a horizon $N = 28$, ensuring that the finite-horizon optimal control problem is equivalent to its infinite-horizon counterpart when an appropriate LQR-based terminal cost is used. Explicit enumeration reveals that 1351 regions are required to represent the exact partitioning of the infinite-horizon cost. In contrast, our proposed method, using a random template with parameters $\mathsf{f} = 212$, $\mathsf{e} = 624$, and $\mathsf{v} = 418$, computes a piecewise affine CLF approximation with a maximum relative error of 0.6. Furthermore, the relative maximum distance between the domain of the computed CLF approximation and the maximal control invariant set is only 0.008, indicating a high degree of accuracy. These approximations may be deemed acceptable in the sense that the approximate feedback comes with a stability guarantee. Given that $212 \ll 1351$, evaluating the approximate feedback is significantly less computationally intensive than evaluating the explicit representation of the optimal feedback.

Next, we set the state constraints to $\mathbb{X} = [-30, 30] \times [-10, 10]$ and the prediction horizon to $N = 35$. This results in 2063 regions for an exact explicit MPC approach. Applying our proposed method with a random template, setting $\mathsf{f} = 375$, $\mathsf{e} = 1119$, and $\mathsf{v} = 746$, yields a piecewise affine CLF that approximates the exact infinite-horizon cost with a maximum relative error of approximately 0.75. As $375 \ll 2063$, the approximate feedback is also in this case less accurate but less costly to store and evaluate.

It is important to note that the above comparisons consider scenarios where evaluating the exact partitioning using explicit MPC is feasible. However, in general, explicit MPC methods face significant computational challenges due to the combinatorial explosion of regions in parametric quadratic programs, often making it infeasible to compute the explicit partitioning. In contrast, our proposed single-stage polyhedral approximation method offers users full control over the computational complexity of the CLF representation. This is particularly relevant in MPC applications, where the primary goal is often to compute CLFs for guaranteeing recursive feasibility and stability. Generally, in the context of designing terminal costs for MPC, especially in the context of controlling large-scale systems, a significant approximation error relative to the optimal infinite horizon cost may be acceptable, as performance can still be fine-tuned by increasing the prediction horizon of the MPC controller.

6. Conclusions

This paper introduced a novel method for the design of control Lyapunov functions (CLFs) for constrained linear systems using configuration-constrained polyhedral computing techniques. The main contribution of this work is the development of convex conditions that characterize CLFs and robust CLFs with configuration-constrained polyhedral epigraphs, as detailed in Theorems 1 and 2. These conditions enable the construction of piecewise affine convex CLFs that approximate the infinite horizon value function for both nominal and min–max optimal control problems through the solution of a single convex optimization problem.

A notable feature of the proposed design techniques is that optimizing the epigraph of the polyhedral CLF inherently determines the associated control inputs at the vertices of the polyhedral domain partition. The resulting control inputs provide a computationally efficient approximation of the infinite horizon optimal control law, which can be evaluated using point location and convex interpolation techniques, as discussed and illustrated in Section 5.

The versatility of this approach holds promise for applications in determining terminal regions and cost functions in MPC. Moreover, it provides a valuable bridge between implicit and explicit control synthesis approaches. Specifically, the proposed methodology offers a means to construct an explicit controller that approximates the optimal infinite horizon feedback control law, with a complexity that is fixed a priori. This characteristic ensures that the resulting controller is both computationally predictable and efficient, making it particularly well-suited for real-time applications, where fixed computational requirements are critical. The numerical examples presented in this paper highlight the practicality and effectiveness of the proposed CLF design strategies, underscoring their potential for wide-ranging applications in control theory and practice.

References

Artstein, Z. (1983). Stabilization with relaxed controls. *Nonlinear Analysis. Theory, Methods & Applications*, *7*(11), 1163–1173.
Bardi, M., & Capuzzo-Dolcetta, I. (1997). *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Birkhäuser.
Bellman, R. E. (1957). *Dynamic programming*. Princeton University Press.
Bemporad, A., Borrelli, F., & Morari, M. (2003). Min-max control of constrained uncertain discrete-time linear systems. *IEEE Transactions on Automatic Control*, *48*(9), 1600–1606.
Bemporad, A., & Filippi, C. (2006). An algorithm for approximate multiparametric convex programming. *Computational Optimization and Applications*, *35*, 87–108.
Bemporad, A., Morari, M., Dua, V., & Pistikopoulos, E. (2002). The explicit linear quadratic regulator for constrained systems. *Automatica*, *38*(1), 3–20.
Bemporad, A., Oliveri, A., Poggi, T., & Storace, M. (2011). Ultra-fast stabilizing model predictive control via canonical piecewise affine approximations. *IEEE Transactions on Automatic Control*, *56*(12), 2883–2897.
Bertsekas, D. P. (2012). *Dynamic programming and optimal control* (3rd). Belmont, Massachusetts: Athena Scientific Dynamic Programming and Optimal Control.
Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review*, *53*(3), 464–501.
Blanchini, F. (1994). Ultimate boundedness control for uncertain discrete-time systems via set-induced Lyapunov functions. *IEEE Transactions on Automatic Control*, *39*(2), 428–433.
Blanchini, F. (1995). Nonquadratic Lyapunov functions for robust control. *Automatica*, *31*(3), 451–461.
Blanchini, F., & Miani, S. (1999). Numerical computation of polyhedral Lyapunov functions for robust synthesis. *IFAC Proceedings*, *32*(2), 2065–2070.
Blanchini, F., & Miani, S. (2015). Set-theoretic methods in control. *Systems & control: foundations & applications*, (p. xvi+481). Boston, MA: Birkhäuser Boston, Inc..
Boyd, S., Ghaoui, L. E., Feron, E., & Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. SIAM.
Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.
Chen, H., & Allgöwer, F. (1998). A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, *34*(10), 1205–1217.
Chen, S. W., Wang, T., Atanasov, N., Kumar, V., & Morari, M. (2022). Large scale model predictive control with neural networks and primal active sets. *Automatica*, *135*, Article 109947.
Diehl, M., & Björnberg, J. (2004). Robust dynamic programming for min-max model predictive control of constrained uncertain systems. *IEEE Transactions on Automatic Control*, *49*(12), 2253–2257.
Dörfler, D. (2022). On the approximation of unbounded convex sets by polyhedra. *Journal of Optimization Theory and Applications*, *194*(1), 265–287.
Fabiani, F., & Goulart, P. J. (2022). Reliably-stabilizing piecewise-affine neural network controllers. *IEEE Transactions on Automatic Control*, *68*(9), 5201–5215.
Faulwasser, T., Grüne, L., Müller, M. A., et al. (2018). Economic nonlinear model predictive control. *Foundations and Trends® in Systems and Control*, *5*(1), 1–98.
Fukuda, K. (2020). *Polyhedral computation*. ETH Zürich Research Collection.
Giesl, P., & Hafstein, S. (2015). Review on computational methods for Lyapunov functions. *Discrete and Continuous Dynamical Systems. Series B*, *20*(8), 2291–2331.
Goulart, P. J., Kerrigan, E. C., & Maciejowski, J. M. (2006). Optimization over state feedback policies for robust control with constraints. *Automatica*, *42*(4), 523–533. http://dx.doi.org/10.1016/j.automatica.2005.08.023.
Grieder, P., Borrelli, F., Torrisi, F., & Morari, M. (2004). Computation of the constrained infinite time linear quadratic regulator. *Automatica*, *40*(4), 701–708.
Gutman, P. O., & Cwikel, M. (1986). Admissible sets and feedback control for discrete-time linear dynamical systems with bounded controls and states. *IEEE Transactions on Automatic Control*, *31*(4), 373–376.
Gutman, P.-O., & Cwikel, M. (1987). An algorithm to find maximal state constraint sets for discrete-time linear dynamical systems with bounded controls and states. *IEEE Transactions on Automatic Control*, *32*, 251–254.
He, K., Shi, S., van den Boom, T., & De Schutter, B. (2024). Approximate dynamic programming for constrained linear systems: A piecewise quadratic approximation approach. *Automatica*, *160*, Article 111456.
Herceg, M., Kvasnica, M., Jones, C., & Morari, M. (2013). Multi-parametric toolbox 3.0. In *Proceedings of European control conference* (pp. 502–510).
Houska, B., Müller, M. A., & Villanueva, M. E. (2024). On stabilizing terminal costs and regions for configuration-constrained tube MPC. *IEEE Control Systems Letters*, *8*, 1961–1966.
Houska, B., Müller, M. A., & Villanueva, M. E. (2025). Polyhedral control design: Theory and methods. *Annual Reviews in Control*, *60*, Article 100992.
Jones, C. N., & Morari, M. (2010). Polytopic approximation of explicit model predictive controllers. *IEEE Transactions on Automatic Control*, *55*(11), 2542–2553.
Karg, B., & Lucia, S. (2020). Efficient representation and approximation of model predictive control laws via deep learning. *IEEE Transactions on Cybernetics*, *50*(9), 3866–3878.
Kerrigan, E. C., & Maciejowski, J. M. (2004). Feedback min-max model predictive control using a single linear program: Robust stability and the explicit solution. *International Journal on Robust and Nonlinear Control*, *14*, 395–413.
Lincoln, B., & Rantzer, A. (2006). Relaxing dynamic programming. *IEEE Transactions on Automatic Control*, *51*(8), 1249–1260.
Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Scokaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, *36*(6), 789–814.
McMullen, P. (1971). The maximum numbers of faces of a convex polytope. *Mathematika*, *17*, 179–184.
Minkowski, H. (1989). Allgemeine Lehrsätze über die konvexen Polyeder. In *Ausgewählte Arbeiten zur Zahlentheorie und zur Geometrie: in D. Hilberts Gedächtnisrede auf H. Minkowski, Göttingen 1909* (pp. 121–139). Vienna: Springer Vienna.
Ney, P. E., & Robinson, S. M. (1995). Polyhedral approximation of convex sets with an application to large deviation probability theory. *Journal of Convex Analysis*, *2*(1–2), 229–240.
Primbs, J. A., Nevistić, V., & Doyle, J. C. (1999). Nonlinear optimal control: A control Lyapunov function and receding horizon perspective. *Asian Journal of Control*, *1*(1), 14–24.
Raković, S. V. (2020). Robust Minkowski-Lyapunov functions. *Automatica*, *120*(109168).

Raković, S. V., & Zhang, S. (2024). Polyhedral robust Minkowski–Lyapunov functions. *IEEE Transactions on Automatic Control*, *69*(6), 3576–3588.
Rawlings, J. B., & Mayne, D. Q. (2009). *Model predictive control: theory and design*. Madison, WI: Nob Hill Publishing.
Scokaert, P. O. M., & Mayne, D. Q. (1998). Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, *43*(8), 1136–1142.
Villanueva, M. E., Müller, M. A., & Houska, B. (2024). Configuration-constrained tube MPC. *Automatica*, *163:111543*.
Weyl, H. (1934). Elementare Theorie der konvexen Polyeder. *Commentarii Mathematici Helvetici*, *7*, 290–306.
Witsenhausen, H. S. (1968). A minimax control problem for sampled linear systems. *IEEE Transactions on Automatic Control*, *13*(1), 5–21.
Xu, J., Lou, Y., De Schutter, B., & Xiong, Z. (2025). Error-free approximation of explicit linear MPC through lattice piecewise affine expression. *IEEE Transactions on Automatic Control*, *70*(3), 1745–1760.
Ziegler, G. M. (1995). *Lectures on polytopes*. Springer.
Zubov, V. I. (1965). Methods of A.M. Lyapunov and their application. *Mathematics of Computation*, *19*, 349.

Boris Houska is an associate professor at ShanghaiTech University. He received a diploma in mathematics from the University of Heidelberg in 2007 and a Ph.D. in electrical engineering from KU Leuven in 2011. From 2012 to 2013, he was a postdoctoral researcher at the Centre for Process Systems Engineering at Imperial College London. Boris Houska's research interests include numerical optimization, optimal control, robust and global optimization, and fast model predictive control algorithms.

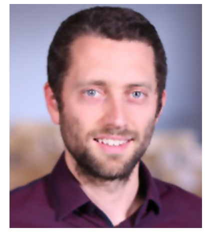

Matthias A. Müller received a Diploma degree in engineering cybernetics from the University of Stuttgart, Germany, an M.Sc. in electrical and computer engineering from the University of Illinois at Urbana-Champaign (both in 2009), and a Ph.D. in mechanical engineering from the University of Stuttgart in 2014. Since 2019, he is Director of the Institute of Automatic Control and Full Professor at the Leibniz University Hannover, Germany. His research interests include nonlinear control and estimation, model predictive control, and data- and learning-based control, with application in different fields including biomedical engineering and robotics. He has received various awards for his work, including the European Systems & Control Ph.D. Thesis Award, an ERC Starting Grant from the European Research Council, the IEEE CSS George S. Axelby Outstanding Paper Award, the Brockett-Willems Outstanding Paper Award, and the Journal of Process Control Paper Award. He serves/d as an associate editor for Automatica and as an editor of the International Journal of Robust and Nonlinear Control.

Mario Eduardo Villanueva is an associate professor at the IMT School for Advanced Studies Lucca. He received his M.Sc. and Ph.D. in Chemical Engineering from Imperial College London in 2011 and 2016, respectively. From 2016 to 2017, he was a postdoctoral researcher at Texas A&M University, and from 2017 to 2021 he held a postdoctoral position at the School of Information Science and Technology, ShanghaiTech University. He is the recipient of the Dudley Newitt Prize (Imperial College London, 2016), the Excellent Postdoc Awards (ShanghaiTech University, 2018 and 2019), and the Journal of Process Control Paper Prize (IFAC, 2020). His research interests include set-based computing, robust and stochastic control, global optimization, and system-theoretic tools for data analysis.