On the Stability of a Class of Robust Receding Horizon Control Laws for Constrained Systems *

Paul J. Goulart and Eric C. Kerrigan

Department of Engineering
University of Cambridge
Cambridge CB2 1PZ
United Kingdom
{pjg42,eck21}@eng.cam.ac.uk

August 4, 2005

Abstract

This paper is concerned with the stability of a class of robust and constrained optimal control laws for linear discrete-time systems subject to bounded state disturbances and arbitrary convex constraints on the states and inputs. The paper considers the class of feedback control policies parameterized as affine functions of the system state, calculation of which has recently been shown to be tractable via a suitable convex reparameterization. When minimizing the expected value of a quadratic cost, we show that the resulting value function in the optimal control problem is convex. When used in the design of a robust receding horizon controller, we provide sufficient conditions to establish that the closed-loop system is input-to-state stable (ISS). The paper further shows that the resulting control law has an interesting interpretation as the projection of the optimal unconstrained linear-quadratic control law onto the set of constraint-admissible control policies.

1 Introduction

This paper is concerned with the stability of a class of robust and constrained optimal control laws for linear discrete-time systems subject to bounded state disturbances, and subject to *arbitrary* convex constraints on the states and inputs. We consider the class of feedback control policies parameterized as affine functions of the system state, calculation of which has been shown to be tractable via a suitable convex reparameterization [1]. When minimizing the expected value of a quadratic cost, we show that the resulting value function in the optimal control problem is convex, and provide sufficient conditions, when used in the design of a robust receding horizon controller, to establish that the closed-loop system is input-to-state stable (ISS).

It is generally accepted that, if one wishes to account for disturbances when designing finite- or receding-horizon control laws for constrained systems, then the optimization must be done over state feedback policies, rather than over open-loop control sequences [2]. However, the difficulty is that proposals for optimization over arbitrary feedback

*Technical Report CUED/F-INFENG/TR.532, Dept. of Engineering, University of Cambridge, UK.

laws, such as [3], are generally intractable, particularly when one wishes to guarantee constraint satisfaction for all possible realizations of the disturbance.

As a result, many compromise solutions have been proposed in the predictive control literature. One common approach is to pre-compute one or more pre-stabilizing linear feedback control laws off-line, and then calculate on-line perturbations to this control law [4–6]. Approaches of this sort, though computationally attractive, are problematic in the sense that it is not obvious how one should select the particular pre-stabilizing feedback law to employ.

An obvious improvement to this approach is to optimize over state feedback policies *on-line* at each time instant. However, this is seemingly problematic since it generally requires the solution of a non-convex optimization problem at each step, because the predicted sequence of states is a nonlinear function of the gains to be optimized.

In a recent publication [1], the authors demonstrate that the *non-convex* state feedback optimization problem can be reparameterized as an *equivalent* but *convex* problem by recasting the optimization problem in terms of affine disturbance or error feedback laws. They further demonstrate that, when implemented in a receding horizon fashion with a particular cost function, the closed loop system is input-to-state stable (ISS) when the constraints and disturbance sets are polytopic.

In this paper we present a generalization of this result, using the expected value of a quadratic cost. We demonstrate that, for systems with *arbitrary* convex state and input constraints and disturbance sets, the resulting value function is convex and lower semicontinuous when optimizing over state feedback policies, and provide conditions under which input-to-state stability can be established for such systems using convex Lyapunov functions. Since the optimization problems we consider are performed over arbitrary convex sets, the proofs differ substantially from those required in the case where the constraints and disturbance sets are polytopic, as in [1]. This is of particular interest, for example, in the case where the disturbance or constraint sets are 2−norm bounded, and the resultant optimization problem can be solved as a tractable second-order-cone program (SOCP), but for which no proof of stability exists at present.

The paper further shows that the resulting control law has an interesting interpretation as the projection of the optimal unconstrained linear-quadratic control law onto the set of constraint-admissible control policies.

Notation: A continuous function $\gamma : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a $\mathcal{K}$-function if it is strictly increasing and $\gamma(0) = 0$; it is a $\mathcal{K}_\infty$-function if, in addition, $\gamma(s) \to \infty$ as $s \to \infty$. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a $\mathcal{KL}$-function if for all $k \geq 0$, the function $\beta(\cdot, k)$ is a $\mathcal{K}$-function and for each $s \geq 0$, $\beta(s, \cdot)$ is decreasing with $\beta(s, k) \to 0$ as $k \to \infty$. $\mathbb{Z}_{[k,l]}$ represents the set of integers $\{k, k+1, \ldots, l\}$. $\mathbb{E}$ is the expectation operator. Given sets X and Y, $X + Y := \{x + y \mid x \in X,\ y \in Y\}$, $\text{int}X$ represents the interior of X, $\text{rint}X$ its relative interior, $\text{lin}X$ its linear hull (i.e. the smallest subspace the contains X), and ∂X its boundary. $\bar{\mathbb{R}}$ represents the extended real line $[-\infty, \infty]$. Given a vector x and matrices A and B, $A \otimes B$ is the Kronecker product of A and B, $\mathcal{N}(A)$ is the null space of A, $\text{tr}(A)$ is the trace of A, $\text{vec}(A)$ denotes the vector formed by stacking the columns of A into a vector, $\|x\|_A^2 := x^T A x$ and $\|x\| := \sqrt{x^T x}$.

2 Definitions and Standing Assumptions

Consider the following discrete-time linear time-invariant system:

$$x^+ = Ax + Bu + w, \tag{1}$$

where $x \in \mathbb{R}^n$ is the system state at the current time instant, x^+ is the state at the next time instant, $u \in \mathbb{R}^m$ is the control input and $w \in \mathbb{R}^n$ is the disturbance. It is assumed that (A, B) is stabilizable and that at each sample instant a measurement of the state is available. It is assumed that the current and future values of the disturbance are unknown and may change from one time instant to the next, but are contained in a compact set W containing the origin in its relative interior. We further assume that, in addition to lying in the set W, the disturbances are independent and identically distributed with mean $\mathbb{E}[w] = 0$ and positive semidefinite covariance $\mathbb{E}\left[ww^T\right] = C_w$. Finally, we assume that the covariance C_w is sensibly defined with respect to the set W, i.e. we assume that $\mathcal{N}(C_w) \cap \mathrm{lin} W = \{0\}$.

The system is subject to mixed convex constraints on the state and input, so that the system must satisfy $(x, u) \in \mathcal{Z}$ where $\mathcal{Z} \subset \mathbb{R}^n \times \mathbb{R}^m$ is a convex and compact set containing the origin in its interior. A design goal is to guarantee that the state and input of the closed-loop system remain in $\mathcal{Z}$ for all time and for all allowable disturbance sequences.

In addition to $\mathcal{Z}$, a target/terminal constraint set $X_f \subset \mathbb{R}^n$ is given, which is convex, compact and contains the origin in its interior. The set X_f can, for example, be used as a target set in time-optimal control or, if chosen to be robust positively invariant, to design a receding horizon controller with guaranteed invariance and stability properties [1].

Before proceeding, we define some additional notation. In the sequel, predictions of the system's evolution over a finite control/planning horizon will be used to define a number of suitable control policies. Let the length N of this planning horizon be a positive integer and define stacked versions of the predicted input, state and disturbance vectors $\mathbf{u} \in \mathbb{R}^{mN}$, $\mathbf{x} \in \mathbb{R}^{n(N+1)}$ and $\mathbf{w} \in \mathbb{R}^{nN}$, respectively, as

$$\mathbf{x} := \mathrm{vec}(x_0, \ldots, x_{N-1}, x_N), \tag{2a}$$

$$\mathbf{u} := \mathrm{vec}(u_0, \ldots, u_{N-1}), \tag{2b}$$

$$\mathbf{w} := \mathrm{vec}(w_0, \ldots, w_{N-1}), \tag{2c}$$

where $x_0 = x$ denotes the current measured value of the state and $x_{i+1} := Ax_i + Bu_i + w_i$, $\forall i \in \mathbb{Z}_{[0,N-1]}$ denotes the prediction of the state after i time instants. We let the set $\mathcal{W} := W^N := W \times \cdots \times W$, so that $\mathbf{w} \in \mathcal{W}$. We define the matrix $\mathbf{C}_w := I \otimes C_w$, so that $\mathbb{E}\left[\mathbf{w}\mathbf{w}^T\right] = \mathbf{C}_w$ and $\mathcal{N}(\mathbf{C}_w) \cap \mathrm{lin}\mathcal{W} = \{0\}$. We define a convex and compact set $\mathcal{Y}$, appropriately constructed from $\mathcal{Z}$ and X_f, such that the constraints to be satisfied are equivalent to $(\mathbf{x}, \mathbf{u}) \in \mathcal{Y}$, i.e.

$$\mathcal{Y} := \left\{ (\mathbf{x}, \mathbf{u}) \;\middle|\; \begin{array}{c} (x_i, u_i) \in \mathcal{Z},\ \forall i \in \mathbb{Z}_{[0,N-1]} \\ x_N \in X_f \end{array} \right\}. \tag{3}$$

Finally, we construct matrices $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{E}$ (defined in the Appendix) using the relation (1) such that $\mathbf{x} = \mathbf{A}x + \mathbf{B}\mathbf{u} + \mathbf{E}\mathbf{w}$.

3 Affine Feedback Control Policies

3.1 State Feedback Parameterization

One natural approach to controlling the system in (1), while ensuring the satisfaction of the constraints (3) for all allowable disturbance sequences, is to search over the set of time-varying affine state feedback control policies. We thus consider policies of the form:

$$u_i = \sum_{j=0}^{i} L_{i,j} x_j + g_i, \quad \forall i \in \mathbb{Z}_{[0,N-1]}, \tag{4}$$

where each $L_{i,j} \in \mathbb{R}^{m\times n}$ and $g_i \in \mathbb{R}^m$. For notational convenience, we also define the block lower triangular matrix $\mathbf{L} \in \mathbb{R}^{mN\times n(N+1)}$ and stacked vector $\mathbf{g} \in \mathbb{R}^{mN}$ as

$$\mathbf{L} := \begin{bmatrix} L_{0,0} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ L_{N-1,0} & \cdots & L_{N-1,N-1} & 0 \end{bmatrix}, \mathbf{g} := \begin{bmatrix} g_0 \\ \vdots \\ g_{N-1} \end{bmatrix} \tag{5}$$

so that the control input sequence can be written as $\mathbf{u} = \mathbf{L}\mathbf{x} + \mathbf{g}$. For a given initial state x (since the system is time-invariant, the current time can always be taken as zero), we say that the pair $(\mathbf{L}, \mathbf{g})$ is admissible if the control policy (4) guarantees that, for all allowable disturbance sequences of length N, the constraints (3) are satisfied over the horizon $i = 0, \ldots, N$. More precisely, the set of admissible $(\mathbf{L}, \mathbf{g})$ is defined as

$$\Pi_N^{sf}(x) := \left\{ (\mathbf{L}, \mathbf{g}) \;\middle|\; \begin{array}{r} (\mathbf{L}, \mathbf{g}) \text{ satisfies (5)} \\ \mathbf{x} = \mathbf{A}x + \mathbf{B}\mathbf{u} + \mathbf{E}\mathbf{w} \\ \mathbf{u} = \mathbf{L}\mathbf{x} + \mathbf{g} \\ (\mathbf{x}, \mathbf{u}) \in \mathcal{Y},\ \forall \mathbf{w} \in \mathcal{W} \end{array} \right\} \tag{6}$$

and the set of initial states x for which an admissible control policy of the form (4) exists is defined as

$$X_N^{sf} := \left\{ x \in \mathbb{R}^n \;\middle|\; \Pi_N^{sf}(x) \neq \emptyset \right\}. \tag{7}$$

As noted in [1], it is generally *not possible* to select a single pair $(\mathbf{L}, \mathbf{g})$ such that $(\mathbf{L}, \mathbf{g}) \in \Pi_N^{sf}(x)$ for all $x \in X_N^{sf}$. Additionally, such a control policy is seemingly very difficult to compute, since the set $\Pi_N^{sf}(x)$ is non-convex. However, for a given $x \in X_N^{sf}$, an admissible pair $(\mathbf{L}, \mathbf{g})$ may be found via convex optimization through an appropriate reparameterization. This parameterization is introduced in the following section.

3.2 Disturbance Feedback Parameterization

An alternative to (4) is to parameterize the control policy as an affine function of the sequence of past disturbances, so that

$$u_i = \sum_{j=0}^{i-1} M_{i,j} w_j + v_i, \quad \forall i \in \mathbb{Z}_{[0,N-1]}, \tag{8}$$

where each $M_{i,j} \in \mathbb{R}^{m\times n}$ and $v_i \in \mathbb{R}^m$. It should be noted that, since full state feedback is assumed, the past disturbance sequence is easily calculated as the difference between the predicted and actual states at each step, i.e.

$$w_i = x_{i+1} - Ax_i - Bu_i, \quad \forall i \in \mathbb{Z}_{[0,N-1]}. \tag{9}$$

Define the variable $\mathbf{v} \in \mathbb{R}^{mN}$ and the block lower triangular matrix $\mathbf{M} \in \mathbb{R}^{mN \times nN}$ such that

$$\mathbf{M} := \begin{bmatrix} 0 & \cdots & \cdots & 0 \\ M_{1,0} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ M_{N-1,0} & \cdots & M_{N-1,N-2} & 0 \end{bmatrix}, \mathbf{v} := \begin{bmatrix} v_0 \\ v_1 \\ \vdots \\ v_{N-1} \end{bmatrix} \tag{10}$$

so that the control input sequence can be written as $\mathbf{u} = \mathbf{M}\mathbf{w} + \mathbf{v}$. Define the set of admissible $(\mathbf{M}, \mathbf{v})$, for which the constraints (3) are satisfied, as:

$$\Pi_N^{df}(x) := \left\{ (\mathbf{M}, \mathbf{v}) \,\middle|\, \begin{array}{r} (\mathbf{M}, \mathbf{v}) \text{ satisfies (10)} \\ \mathbf{x} = \mathbf{A}x + \mathbf{B}\mathbf{u} + \mathbf{E}\mathbf{w} \\ \mathbf{u} = \mathbf{M}\mathbf{w} + \mathbf{v} \\ (\mathbf{x}, \mathbf{u}) \in \mathcal{Y}, \ \forall \mathbf{w} \in \mathcal{W} \end{array} \right\}, \tag{11}$$

and define the set of initial states x for which an admissible control policy of the form (8) exists as

$$X_N^{df} := \{x \in \mathbb{R}^n \mid \Pi_N^{df}(x) \neq \emptyset\}. \tag{12}$$

We are interested in this control policy parameterization primarily due to the following two properties, partial proof of which may be found in [1]:

Theorem 1 (Convexity). *For a given state $x \in X_N^{df}$, the set of admissible affine disturbance feedback parameters $\Pi_N^{df}(x)$ is closed and convex. Furthermore, the set of states X_N^{df}, for which at least one admissible affine disturbance feedback parameter exists, is also closed and convex.*

Theorem 2 (Equivalence). *The set of admissible states $X_N^{df} = X_N^{sf}$. Additionally, given any $x \in X_N^{sf}$, for any admissible $(\mathbf{L}, \mathbf{g})$ an admissible $(\mathbf{M}, \mathbf{v})$ can be constructed that yields the same input and state sequence for all allowable disturbances, and vice-versa.*

Remark 1. The method of proof for convexity in [1] is insufficient for the rather general convex state and input constraints presented here. However, proof of convexity is straightforward by noting that the set

$$\mathcal{C}_N := \bigcap_{\mathbf{w} \in \mathcal{W}} \left\{ (\mathbf{M}, \mathbf{v}, x) \,\middle|\, \begin{array}{r} (\mathbf{M}, \mathbf{v}) \text{ satisfies (10)} \\ \mathbf{x} = \mathbf{A}x + \mathbf{B}\mathbf{u} + \mathbf{E}\mathbf{w} \\ \mathbf{u} = \mathbf{M}\mathbf{w} + \mathbf{v} \\ (\mathbf{x}, \mathbf{u}) \in \mathcal{Y} \end{array} \right\} \tag{13}$$

is closed and convex, since it is the intersection of closed and convex sets. The set $\Pi_N^{df}(x)$ is also closed and convex, since (11) can be rewritten in a similar manner. The set X_N^{df} is convex since it can be expressed as a projection of $\mathcal{C}_N$ onto an appropriate subspace. Closedness of X_N^{df} also follows directly from this projection when the set $\mathcal{C}_N$ is bounded; we prove the result in the more general unbounded case in Lemma 2 in the Appendix.

4 An Expected Value Cost Function

We consider a function $\Phi(x, \mathbf{u})$ which is quadratic in the state and control sequence, and seek a control policy that will minimize its expected value over the planning horizon.

We define

$$\Phi(x,\mathbf{u}) := \sum_{i=0}^{N-1}(\|x_i\|_Q + \|u_i\|_R) + \|x_N\|_P, \tag{14}$$

where, for all i, $x_{i+1} = Ax_i + Bu_i + w_i$, and Q, R and P are positive definite. We define an optimal policy pair $(\mathbf{L}^*(x), \mathbf{g}^*(x)) \in \Pi_N^{sf}(x)$ to be one which minimizes the expected value (over all disturbances) of this function over the set of feasible control policies (6). We thus define

$$V_N(x,\mathbf{L},\mathbf{g}) := \mathbb{E}\left[\Phi(x,\bar{\mathbf{u}})\right] \tag{15}$$

where $\bar{\mathbf{u}} := \mathbf{L}\bar{\mathbf{x}} + \mathbf{g}$ and $\bar{\mathbf{x}} := (I - \mathbf{BL})^{-1}(\mathbf{A}x + \mathbf{Bg} + \mathbf{w})$, and define an optimal policy pair as

$$(\mathbf{L}^*(x), \mathbf{g}^*(x)) := \underset{(\mathbf{L},\mathbf{g})\in\Pi_N^{sf}(x)}{\operatorname{argmin}} V_N(x,\mathbf{L},\mathbf{g}). \tag{16}$$

We assume for the moment that the minimizer in (16) exists and is well-defined. The receding horizon control policy $\mu_N : X_N^{sf} \to \mathbb{R}^m$ is defined by the first part of the optimal affine state feedback control policy, i.e.

$$\mu_N(x) := L_{0,0}^*(x)x + g_0^*(x) \tag{17}$$

Note that the control law $\mu_N(\cdot)$ is time-invariant and is, in general, a nonlinear function of the current state. The closed-loop system becomes

$$x^+ = Ax + B\mu_N(x) + w. \tag{18}$$

We also define the value function $V_N^* : X_N^{sf} \to \mathbb{R}_{\geq 0}$ to be

$$V_N^*(x) := \min_{(\mathbf{L},\mathbf{g})\in\Pi_N^{sf}(x)} V_N(x,\mathbf{L},\mathbf{g}). \tag{19}$$

As noted in [1], the difficulty with this scheme lies in the non-convexity of the set of feasible policies $\Pi_N^{sf}(x)$ and of the function $V_N(x,\cdot,\cdot)$, and thus in the non-convexity of the optimization problem (19). We therefore exploit the alternative parameterization (8), and define the analogous cost function

$$J_N(x,\mathbf{M},\mathbf{v}) := \mathbb{E}\left[\Phi(x,\hat{\mathbf{u}})\right] \tag{20}$$

where $\hat{\mathbf{u}} := \mathbf{Mw} + \mathbf{v}$. In this case we define an optimal policy as:

$$(\mathbf{M}^*(x), \mathbf{v}^*(x)) := \underset{(\mathbf{M},\mathbf{v})\in\Pi_N^{df}(x)}{\operatorname{argmin}} J_N(x,\mathbf{M},\mathbf{v}). \tag{21}$$

We again assume for the moment that the minimizer in (21) exists and is well-defined. It then follows from Theorem 2 that

$$V_N^*(x) = \min_{(\mathbf{M},\mathbf{v})\in\Pi_N^{df}(x)} J_N(x,\mathbf{M},\mathbf{v}). \tag{22}$$

The control policy (17) is then given by

$$\mu_N(x) = v_0^*(x) = L_{0,0}^*(x)x + g_0^*(x). \tag{23}$$

We first demonstrate that the function $J_N(x,\cdot,\cdot)$ is convex, so the problem (22) can be solved using standard techniques in convex optimization.

Proposition 1 (Convex Cost). *The function* $(x, \mathbf{M}, \mathbf{v}) \mapsto J_N(x, \mathbf{M}, \mathbf{v})$ *is convex and quadratic in the state* x *and parameter* $\mathbf{M}$*, and strictly convex and quadratic in the parameter* $\mathbf{v}$*.*

Proof. The function (20) can be rewritten as:

$$J_N(x, \mathbf{M}, \mathbf{v}) = \mathbb{E}\left[(\|(\mathbf{A}x + \mathbf{B}\mathbf{v}) + (\mathbf{E} + \mathbf{B}\mathbf{M})\mathbf{w}\|^2_{\mathcal{Q}} + \|\mathbf{M}\mathbf{w} + \mathbf{v}\|^2_{\mathcal{R}})\right]$$

where $\mathcal{Q} := \left[\begin{smallmatrix} I \otimes Q & \\ & P \end{smallmatrix}\right]$ and $\mathcal{R} := I \otimes R$. Since $\mathbb{E}[\mathbf{w}] = 0$ and $\mathbf{w}$ is independent of both $\mathbf{v}$ and $\mathbf{M}$, this simplifies to

$$J_N(x, \mathbf{M}, \mathbf{v}) = \mathbb{E}\left[\|\mathbf{v}\|^2_{\mathcal{S}} + x^T H_v \mathbf{v} + \mathbf{w}^T H_M \mathbf{M}\mathbf{w} + \|\mathbf{M}\mathbf{w}\|^2_{\mathcal{S}} + \|\mathbf{A}x\|^2_{\mathcal{Q}} + \|\mathbf{E}\mathbf{w}\|^2_{\mathcal{Q}}\right]$$

where $H_v := 2\mathbf{A}^T\mathcal{Q}\mathbf{B}$, $H_M := 2\mathbf{E}^T\mathcal{Q}\mathbf{B}$, and $\mathcal{S} := \mathbf{B}^T\mathcal{Q}\mathbf{B} + \mathcal{R}$. This can be further simplified noting that $\mathbb{E}\left[\mathbf{w}^T X \mathbf{w}\right] = \mathrm{tr}(X\mathbf{C}_w) = \mathrm{tr}(\mathbf{C}_w X)$ for any X, so that

$$J_N(x, \mathbf{M}, \mathbf{v}) = \|\mathbf{v}\|^2_{\mathcal{S}} + x^T H_v \mathbf{v} + \mathrm{tr}(\mathbf{C}_w H_M \mathbf{M}) + \mathrm{tr}(\mathbf{M}^T \mathcal{S}\mathbf{M}\mathbf{C}_w) + \gamma, \tag{24}$$

where $\gamma := \mathrm{tr}(\mathbf{E}^T\mathcal{Q}\mathbf{E}\mathbf{C}_w) + \|\mathbf{A}x\|^2_{\mathcal{Q}}$. Finally, recalling the matrix identities $\mathrm{vec}(AXB) = (B^T \otimes A)\mathrm{vec}(X)$, $(A \otimes B)(C \otimes D) = (AC \otimes BD)$, and $\mathrm{tr}(A^T B) = \mathrm{vec}(A)^T \mathrm{vec}(B)$, the above may be written in vectorized form as:

$$J_N(x, \mathbf{M}, \mathbf{v}) \ = \ \|\mathbf{v}\|^2_{\mathcal{S}} \ + \ x^T H_v \mathbf{v} \ + \|\mathrm{vec}(\mathbf{M})\|^2_{(\mathbf{C}_w \otimes \mathcal{S})} + \mathrm{vec}(H_M^T \mathbf{C}_w)\mathrm{vec}(\mathbf{M}) + \gamma. \tag{25}$$

The matrix $\mathbf{C}_w$ is positive (semi)definite, so $\mathbf{C}_w \otimes \mathcal{S}$ is positive (semi)definite, since it is the Kronecker product of two positive (semi)definite matrices. This follows directly from the eigenvalue property of Kronecker products; see [7, Thm. 4.2.12]. The matrix $\mathcal{S}$ is positive definite, since $\mathcal{R}$ is positive definite by assumption, so the function is strictly convex in $\mathbf{v}$. □

Since the function $J_N(x, \cdot, \cdot)$ is to be minimized over the potentially unbounded set $\Pi^{df}_N(x)$, it is not immediately obvious that a minimizer in (21) should exist. However, by exploiting the special structure of the set $\Pi^{df}_N(x)$ and of the function $J_N(x, \cdot, \cdot)$ in (24), we may state the following result:

Proposition 2. *The function* $J_N(x, \cdot, \cdot)$ *attains its minimum on the set* $\Pi^{df}_N(x)$*.*

Proof. See the Appendix. □

Note that in conjunction with Theorem 2, this implies that $V_N(x, \cdot, \cdot)$ also attains its minimum on the set $\Pi^{sf}_N(x)$ in (16).

Remark 2. Since the function to be minimized in (22) is convex in the decision parameters $\mathbf{M}$ and $\mathbf{v}$, and the minimization is over the convex set $\Pi^{df}_N(x)$, the optimization problem (22) is easily solved using standard techniques from convex optimization. For example, it can be shown that if the constraint sets $\mathcal{Z}$ and $\mathcal{X}_f$ are polytopic, then (22) can be written as a tractable second order cone program (SOCP) when the set W is ellipsoidal or 2−norm bounded, or as a tractable quadratic program (QP) when W is polytopic [1].

Remark 3. A similar result can be derived for a broad class of alternative cost functions; for example [1] employs a quadratic function of the undisturbed state and control sequences, and [8] employs a cost function akin to that employed in $\mathcal{H}_\infty$ control. The assumption may also be satisfied for various min-max formulations as in [3, 9, 10].

5 Preliminary Results

We wish to find conditions under which the closed-loop system (18) is input-to-state stable (ISS). In order to do this, we first develop some preliminary results related to the convexity of the value function $V_N^*(\cdot)$ in (19), and to input-to-state stability for systems with convex Lyapunov functions.

5.1 Continuity and Convexity of the Value Function

We first demonstrate that the value function $V_N^*(\cdot)$ in (19) is convex and continuous on the interior of its domain, by exploiting the relation (22); this property will prove useful in our subsequent proof of stability for the closed loop system (18). Note that the proof presented here requires *only* convexity of the state and input constraints, and does *not* make the usual assumption (as in [1, 11, 12]) that the constraint set $\mathcal{Y}$ and disturbance set W are polytopic, leading to a piecewise-quadratic value function. We instead exploit several results from variational analysis to establish convexity of $V_N^*(\cdot)$ directly.

Proposition 3 (Continuity and convexity of $V_N^*(\cdot)$ and $\mu_N(\cdot)$). *If X_N^{sf} is non-empty, then the receding horizon control law $\mu_N(\cdot)$ is unique on X_N^{sf} and continuous on $\operatorname{int}X_N^{sf}$. The value function $V_N^*(\cdot)$ is convex on X_N^{sf}, continuous on $\operatorname{int}X_N^{sf}$ and lower semicontinuous everywhere on X_N^{sf}.*

Proof. See the Appendix. □

5.2 Input-to-State Stability

We next develop a result on the input-to-state stability of systems with convex value functions. We can then exploit the convexity of the value function $V_N^*(\cdot)$ to provide conditions in which the closed-loop system (18) is input-to-state stable (ISS) when implemented in a receding horizon fashion.

Consider a nonlinear, time-invariant, discrete-time system of the form

$$x^+ = f(x, w), \tag{26}$$

where $x \in \mathbb{R}^n$ is the state and $w \in \mathbb{R}^l$ is a disturbance that takes on values in a compact set $W \subset \mathbb{R}^l$ containing the origin. It is assumed that the state is measured at each time instant, that $f : \mathbb{R}^n \times \mathbb{R}^l \to \mathbb{R}^n$ is continuous at the origin and that $f(0, 0) = 0$. Given an initial state x and a disturbance sequence $w(\cdot)$, where $w(k) \in W$ for all $k \in \mathbb{Z}_{[0,\infty)}$, let the solution to (26) at time k be denoted by $\phi(k, x, w(\cdot))$. For systems of this type, a useful definition of stability is input-to-state stability:

Definition 1 (ISS). *For system* (26)*, the origin is* input-to-state stable (ISS) *with region of attraction $\mathcal{X} \subseteq \mathbb{R}^n$, which contains the origin in its interior, if there exist a $\mathcal{KL}$-function $\beta(\cdot)$ and a $\mathcal{K}$-function $\gamma(\cdot)$ such that for all initial states $x \in \mathcal{X}$ and disturbance sequences $w(\cdot)$, where $w(k) \in W$ for all $k \in \mathbb{Z}_{[0,\infty)}$, the solution of the system satisfies $\phi(k, x, w(\cdot)) \in \mathcal{X}$ and*

$$\|\phi(k, x, w(\cdot))\| \leq \beta(\|x\|, k) + \gamma\left(\sup\left\{\|w(\tau)\| \;\middle|\; \tau \in \mathbb{Z}_{[0,k-1]}\right\}\right) \tag{27}$$

for all $k \in \mathbb{N}$.

Lemma 1 (ISS-Lyapunov function [13, Lem. 3.5]). *For the system* (26)*, the origin is ISS with region of attraction* $\mathcal{X} \subseteq \mathbb{R}^n$ *if the following conditions are satisfied:*

- $\mathcal{X}$ *contains the origin in its interior and is robust positively invariant for* (26)*, i.e.* $f(x, w) \in \mathcal{X}$ *for all* $x \in \mathcal{X}$ *and all* $w \in W$.
- *There exist* $\mathcal{K}_\infty$ *functions* $\alpha_1(\cdot)$, $\alpha_2(\cdot)$ *and* $\alpha_3(\cdot)$, *a* $\mathcal{K}$*-function* $\sigma(\cdot)$, *and a function* $V : \mathcal{X} \to \mathbb{R}_{\geq 0}$ *such that for all* $x \in \mathcal{X}$,

$$\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|) \tag{28a}$$
$$V(f(x, w)) - V(x) \leq -\alpha_3(\|x\|) + \sigma(\|w\|) \tag{28b}$$

Remark 4. A function $V(\cdot)$ that satisfies the conditions in Lemma 1 is called an *ISS-Lyapunov function.* It is important to note that continuity of the function V is *not required* in the proof of [13, Lem. 3.5], though condition (28a) does imply continuity at the origin.

Proposition 4 (Convex Lyapunov function for undisturbed system). *Let* $\mathcal{X} \subseteq \mathbb{R}^n$ *be a compact robust positively invariant set for* (26) *containing the origin in its interior. Furthermore, let there exist* $\mathcal{K}_\infty$*-functions* $\alpha_1(\cdot)$, $\alpha_2(\cdot)$ *and* $\alpha_3(\cdot)$ *and a function* $V : \mathcal{X} \to \mathbb{R}_{\geq 0}$ *that is convex on* $\mathcal{X}$ *such that for all* $x \in \mathcal{X}$,

$$\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|) \tag{29a}$$
$$V(f(x, 0)) - V(x) \leq -\alpha_3(\|x\|) \tag{29b}$$

The function $V(\cdot)$ *is an ISS-Lyapunov function and the origin is ISS for the system* (26) *with region of attraction* $\mathcal{X}$ *if* $f(\cdot)$ *can be written as*

$$f(x, w) := g(x) + w, \tag{29c}$$

and W *is compact and convex, containing the origin in its relative interior.*

Proof. See the Appendix. □

Remark 5. Note that unlike in [1], which requires Lipschitz continuity of the function $V(\cdot)$, the proof of Proposition 4 requires only the weaker condition that $V(\cdot)$ be convex on $\mathcal{X}$. This allows application of the result to a broader class of systems with arbitrary convex constraints, since in these cases one can often only find functions, such as the value function (22), which are convex and lower semicontinuous on their domains.

6 Minimum Expected Value Control Law

Given the results of the previous sections, we can now provide conditions which allow for the synthesis of a control law that guarantees that the closed-loop system (18) is (ISS). We first make the following assumption:

A1 (Terminal Cost and Constraint) The terminal constraint set X_f is chosen to be both constraint admissible and robust positively invariant under the control $u = Kx$, i.e.

$$X_f \subseteq \{x \mid (x, Kx) \in \mathcal{Z}\} \tag{30a}$$
$$(A + BK)x + w \in X_f,\ \forall x \in X_f, \forall w \in W \tag{30b}$$

We further assume that the feedback matrix K and terminal cost function P are derived from the solution to the discrete algebraic Riccati equation

$$P := Q + A^T PA - A^T PB(R + B^T PB)^{-1} B^T PA \tag{30c}$$

$$K := -(R + B^T PB)^{-1} B^T PA \tag{30d}$$

Remark 6. The reader is referred to [5, 14, 15] and the references therein for details on how to compute a set X_f that satisfies (30). Note that the terminal cost $F(x) := x^T Px$ is a Lyapunov function in the terminal set X_f for the undisturbed closed loop system $x^+ = (A + BK)x$ in the sense that

$$F((A + BK)x) - F(x) \leq -x^T(Q + K^T RK)x, \quad \forall x \in X_f. \tag{31}$$

Remark 7. Note that, when the constraint sets $\mathcal{Z}$ and X_f are $\mathbb{R}^n \times \mathbb{R}^m$ and $\mathbb{R}^n$, respectively, the control policy $u = Kx$ minimizes *both* the expected value of $\Phi(x, \cdot)$ (assuming $\mathbb{E}[\mathbf{w}] = 0$), *and* the value of the deterministic or certainty-equivalent cost one would compute by setting $\mathbf{w} = \{0\}$ [16]. It should be noted that this certainty equivalence property does *not* hold in the general constrained case considered here; i.e.

$$\underset{(\mathbf{L},\mathbf{g})\in\Pi_N^{sf}(x)}{\text{argmin}} \ \mathbb{E}[V_N(x, \bar{\mathbf{u}})] \neq \underset{(\mathbf{L},\mathbf{g})\in\Pi_N^{sf}(x)}{\text{argmin}} \ V_N(x, \mathbb{E}[\bar{\mathbf{u}}]).$$

However, it is still true that $v_0^*(x) = Kx$ for all $x \in X_f$, since in this case the conditions (30) guarantee that the optimal *unconstrained* state feedback gain K is also constraint admissible.

Theorem 3 (ISS for RHC). *If **A1** holds, then the origin is ISS for the closed-loop system* (18) *with region of attraction X_N^{sf}. Furthermore, the input and state constraints* (3) *are satisfied for all time and for all allowable disturbance sequences if and only if the initial state $x(0) \in X_N^{sf}$.*

Proof. For the system of interest, we select $V(\cdot) = V_N^*(\cdot) - \text{tr}(\mathbf{E}^T \mathcal{Q} \mathbf{E} \mathbf{C}_w)$, and let $f(x, w) := Ax + B\mu_N(x) + w$. The set X_N^{sf} is robust positively invariant for system (18), with $0 \in \text{int}X_N^{sf}$ [1]. X_N^{sf} is compact since it is closed (cf. Remark 1) and bounded because $\mathcal{Z}$ is assumed bounded. Since $0 \in X_f$, it is easy to show that $f(0, 0) = 0$ if **A1** holds. By the principle of optimality, V is lower bounded by $\alpha_1(\|x\|) := x^T Px$. Since $0 \in \text{int}X_N^{sf}$, one can also construct a function $\alpha_2(\cdot)$ to satisfy (29a), using arguments similar to those in the proof of Proposition 3.

Using standard techniques [2], it is easy to show that $V(\cdot) := V_N^*(\cdot)$ is a Lyapunov function for the *undisturbed* system $x^+ = Ax + B\mu_N(x)$. More precisely, the methods in [2] can be employed to show that (29b) holds with $\alpha_3(z) := (1/2)\lambda_{\min}(Q)z^2$.

Finally, recall from Proposition 3 that $V_N^*(\cdot)$ is convex and continuous on $\text{int}X_N^{sf}$. By combining all of the above, it follows from Proposition 4 that $V_N^*(\cdot)$ is an ISS-Lyapunov function for system (18). $\square$

6.1 Relationship to LQ control

We next consider the relationship between the optimal *unconstrained* linear quadratic (LQ) control law $u = Kx$, and the optimal *constrained* control policy $(\mathbf{M}^*(x), \mathbf{v}^*(x)$ in (21). We will demonstrate that this optimal pair may be characterized as a weighted projection of the optimal unconstrained policy onto a convex set.

6.1.1 Optimal Unconstrained Disturbance Feedback Policy

It can be shown [17, Sec. 11.2] that when the initial state is known exactly and the disturbances w_i are zero mean and independent of the states x_i and controls u_i, (14) can be written as

$$\mathbb{E}\left[\Phi(x,\mathbf{u})\right] = \mathbb{E}\left[x_0 P x_0 + \sum_{i=0}^{N-1} \|(u_i - K_i x_i)\|^2_{(B^T PB+R)} + \sum_{i=0}^{N-1} w_i^T P w_i\right],$$

where K and P are defined in (30). Note that, in the unconstrained case, it follows immediately from the above that this expected value can be minimized by selecting $u_i = Kx_i$, and that this is true *regardless of the covariance of w*. The above can be written in stacked or vectorized form as

$$\mathbb{E}\left[\Phi(x,\mathbf{u})\right] = \mathbb{E}\left[x_0^T P x_0 + \|\mathbf{u} - \mathbf{K}\mathbf{x}\|^2_{\mathcal{T}} + \mathbf{w}^T \mathcal{P} \mathbf{w}\right],$$

where $\mathbf{K} := I_N \otimes K$, $\mathcal{T} := I_N \otimes (B^T PB + R)$, and $\mathcal{P} := I_N \otimes P$. Recalling (8), it then follows that (21) can be rewritten (dropping constant terms) as

$$(\mathbf{M}^*(x), \mathbf{v}^*(x)) = \underset{(\mathbf{M},\mathbf{v})\in\Pi_N^{df}(x)}{\operatorname{argmin}} \ \mathbb{E}\left[\|(\mathbf{M}\mathbf{w} + \mathbf{v} - \mathbf{K}\hat{\mathbf{x}})\|^2_{\mathcal{T}}\right],$$

where $\hat{\mathbf{x}} = \mathbf{A}x + \mathbf{B}\mathbf{v} + (\mathbf{E} + \mathbf{B}\mathbf{M})\mathbf{w}$. Eliminating the states $\hat{\mathbf{x}}$, it follows that

$$(\mathbf{M}^*(x), \mathbf{v}^*(x)) = \underset{(\mathbf{M},\mathbf{v})\in\Pi_N^{df}(x)}{\operatorname{argmin}} \ \mathbb{E}\left[\|(I-\mathbf{KB})\mathbf{M}\mathbf{w} - \mathbf{KE}\mathbf{w} + (I-\mathbf{KB})\mathbf{v} - \mathbf{KA}x\|^2_{\mathcal{T}}\right]. \tag{32}$$

It is then easy to see that, in the *unconstrained* case, an optimal policy pair $(\bar{\mathbf{M}}, \bar{\mathbf{v}})$ for the *unconstrained* system can be defined as

$$\bar{\mathbf{M}}(x) := (I - \mathbf{KB})^{-1}\mathbf{KE} \tag{33a}$$

$$\bar{\mathbf{v}}(x) := (I - \mathbf{KB})^{-1}\mathbf{KA}x, \tag{33b}$$

where $(I - \mathbf{KB})^{-1}$ always exists since $\mathbf{KB}$ is strictly lower triangular. By comparing this result to that in [1], it is easy to show that this pair matches the control from the unconstrained LQ control problem.

6.1.2 Optimal Constrained Disturbance Feedback Policy

We are of course more interested in characterizing the optimal control policy in the *constrained* case, i.e. where the optimal unconstrained LQ control policy $(\bar{\mathbf{M}}(x), \bar{\mathbf{v}}(x)) \notin \Pi_N^{df}(x)$. This leads to the following result:

Proposition 5. *If C_w is positive definite, the solution to the optimization problem* (21) *may be found as a weighted projection of the origin onto the convex set $\Pi_N^{df}(x) - (\bar{\mathbf{M}}(x), \bar{\mathbf{v}}(x))$.*

Proof. Define $\delta\mathbf{M}(x) \in \mathbb{R}^{mN\times nN}$ and $\delta\mathbf{v}(x) \in \mathbb{R}^{mN}$ as

$$\delta\mathbf{M}(x) := \mathbf{M}(x) - \bar{\mathbf{M}}(x), \quad \delta\mathbf{v}(x) := \mathbf{v}(x) - \bar{\mathbf{v}}(x) \tag{34}$$

with corresponding feasible set

$$\delta\Pi_N^{df}(x) := \Pi_N^{df}(x) + \{(-\bar{\mathbf{M}}(x), -\bar{\mathbf{v}}(x))\}.$$

This set is closed and convex, since it is the translation of a closed and convex set. In the cases of interest, it will *not* contain the origin. The optimization problem (32) is equivalent to

$$(\delta\mathbf{M}^*(x), \delta\mathbf{v}^*(x)) := \operatorname*{argmin}_{(\delta\mathbf{M},\delta\mathbf{v})\in\delta\Pi_N^{df}(x)} \mathbb{E}\left[\|(I-\mathbf{KB})\delta\mathbf{M}\mathbf{w} + (I-\mathbf{KB})\delta\mathbf{v}\|_{\mathcal{T}}^2\right] \tag{35}$$

Once again exploiting the independence of $\mathbf{w}$, and defining the positive definite matrix $\mathcal{H} := (I-\mathbf{KB})^T\mathcal{T}(I-\mathbf{KB})$, the above can be written as

$$\begin{aligned}(\delta\mathbf{M}^*(x), \delta\mathbf{v}^*(x)) &= \operatorname*{argmin}_{(\delta\mathbf{M},\delta\mathbf{v})\in\delta\Pi_N^{df}(x)} \mathbb{E}\left[\|\delta\mathbf{M}\mathbf{w}\|_{\mathcal{H}}^2 + \|\delta\mathbf{v}\|_{\mathcal{H}}^2\right] \\ &= \operatorname*{argmin}_{(\delta\mathbf{M},\delta\mathbf{v})\in\delta\Pi_N^{df}(x)} \operatorname{vec}(\delta\mathbf{M})^T(\mathbf{C}_w\otimes\mathcal{H})\operatorname{vec}(\delta\mathbf{M}) + \|\delta\mathbf{v}\|_{\mathcal{H}}^2 \end{aligned} \tag{36}$$

Finally, defining the set $\delta\hat{\Pi}_N^{df}(x)$ as

$$\delta\hat{\Pi}_N^{df}(x) := \left\{(\delta\hat{\mathbf{M}}, \delta\hat{\mathbf{v}}) \;\middle|\; \begin{array}{l} \operatorname{vec}(\delta\hat{\mathbf{M}}) = (\mathbf{C}_w\otimes\mathcal{H})^{\frac{1}{2}}\operatorname{vec}(\delta\mathbf{M}) \\ \delta\hat{\mathbf{v}} = \mathcal{H}^{\frac{1}{2}}\delta\mathbf{v},\ (\delta\mathbf{M},\delta\mathbf{v})\in\delta\Pi_N^{df}(x) \end{array}\right\}, \tag{37}$$

the solution to (35) may be found by solving the Euclidean projection problem

$$(\delta\hat{\mathbf{M}}^*(x), \delta\hat{\mathbf{v}}^*(x)) := \operatorname*{argmin}_{(\delta\hat{\mathbf{M}},\delta\hat{\mathbf{v}})\in\delta\hat{\Pi}_N^{df}(x)} \Big(\|\operatorname{vec}(\delta\hat{\mathbf{M}})\| + \|\delta\hat{\mathbf{v}}\|\Big) \tag{38}$$

and then setting

$$\operatorname{vec}(\delta\mathbf{M}^*(x)) = (\mathbf{C}_w\otimes\mathcal{H})^{-\frac{1}{2}}\operatorname{vec}(\delta\hat{\mathbf{M}}^*(x)) \tag{39a}$$

$$\delta\mathbf{v}^*(x) = \mathcal{H}^{-\frac{1}{2}}\delta\hat{\mathbf{v}}^*(x) \tag{39b}$$

The optimal pair $(\delta\mathbf{M}^*(x), \delta\mathbf{v}^*(x))$ is thus one which minimizes the weighted Euclidean distance between the origin and the set $\delta\Pi_N^{df}(x)$. The optimal policy $(\mathbf{M}^*(x), \mathbf{v}^*(x))$ is therefore one which is the minimum weighted distance from the optimal unconstrained solution to the set of constraint admissible policies $\Pi_N^{df}(x)$. □

7 Conclusions

Using an affine state feedback policy parameterization and exploiting the results in [1] in the calculation of optimal receding horizon control laws, we have shown that input-to-state stability of the closed-loop system can be established for problems with general convex state and input constraints using the expected value of a quadratic cost, given appropriate terminal conditions.

The keys to this result are proving the existence of minimizers and convexity of the value function in the underlying optimal control problem using results from variational analysis, as well as providing conditions under which input-to-state stability may be established using convex Lyapunov functions.

This result represents an important generalization of the results in [1], as it establishes stability for a broad class of optimal control problems using this framework with non-polytopic but convex disturbance sets (e.g. ellipsoidal or 2-norm bounded disturbances), or for problems with general convex constraints on the states and inputs. The results presented here may also prove useful for problems with alternative convex cost functions, including min-max problems [3, 9, 10] and $\mathcal{H}_\infty$ formulations [8].

We further demonstrated that the resulting control law has an interesting interpretation in terms of the projection of the optimal unconstrained control policy onto the set of constraint admissible feedback policies.

Acknowledgements

Paul Goulart would like to thank the Gates Cambridge Trust for their support. Eric Kerrigan is a Royal Academy of Engineering Research Fellow and would like to thank the Royal Academy of Engineering for supporting this research. The authors wish to thank Professor Zhong-Ping Jiang for his useful feedback during the preparation of this paper.

A Appendix

A.1 Supporting Results and Proofs

Closedness of projections of $\mathcal{C}_N$

We consider the general case where the set $\mathcal{C}_N$ defined in (13) is unbounded, which can happen when the interior of the disturbance set W is empty, and demonstrate that projections of this set (in particular the sets $\Pi_N^{df}(x)$ and X_N^{df}) are closed. We define the set $\mathcal{M}$ and its orthogonal complement $\mathcal{M}_\perp$ to be

$$\mathcal{M} := \{\mathbf{M} \mid \mathbf{M} \text{ satisfies (10)}, \ \mathbf{M}\mathbf{w} = 0, \ \forall \mathbf{w} \perp \text{lin}\mathcal{W}\} \tag{40}$$

$$\mathcal{M}_\perp := \{\mathbf{M} \mid \mathbf{M} \text{ satisfies (10)}, \ \mathbf{M}\mathbf{w} = 0, \ \forall \mathbf{w} \in \text{lin}\mathcal{W}\} \tag{41}$$

Note that both of these sets are actually subspaces, with $\mathcal{M} \cup \mathcal{M}_\perp$ equal to the set of all matrices satisfying (10). Additionally note that if the set W (and thus $\mathcal{W}$) has non-empty interior, then $\mathcal{M}_\perp = \{0\}$, and $\mathcal{C}_N$ is easily shown to be compact. We can now state the following:

Lemma 2. *Any projection of the set $\mathcal{C}_N$ is closed.*

Proof. We define the set

$$\tilde{\mathcal{C}}_N := \mathcal{C}_N \cap (\mathcal{M} \times \mathbb{R}^{mN} \times \mathbb{R}^n) \tag{42}$$

which is closed, since it is the intersection of a closed set and a subspace, and bounded, since the set $\mathcal{Y}$ is compact and $\max_{\mathbf{w}\in\mathcal{W}} \|\mathbf{M}\mathbf{w}\| > 0$ for any non-zero $\mathbf{M} \in \mathcal{M}$. From the definition of $\mathcal{M}_\perp$ in (41), it immediately follows that $\mathcal{C}_N$ in (13) can be written as

$$\mathcal{C}_N = \tilde{\mathcal{C}}_N + (\mathcal{M}_\perp \times \{0\} \times \{0\})$$

which is the sum of a compact set and a subspace. Since projection is distributive with respect to set addition, any projection of $\mathcal{C}_N$ is the sum of the projections of these sets. Since any projection of the compact set $\tilde{\mathcal{C}}_N$ is compact, and any projection of the subspace $\mathcal{M}_\perp \times \{0\} \times \{0\}$ is a subspace, projections of the set $\mathcal{C}_N$ are closed [18, Ex. 3.12]. □

Using arguments similar to those above, we can prove Proposition 2 by decomposing the set $\Pi_N^{df}(x)$ into the sum of a compact set and a subspace.

Proof of Proposition 2

Proof. Consider the set

$$\tilde{\Pi}_N^{df}(x) := \Pi_N^{df}(x) \cap (\mathcal{M} \times \mathbb{R}^{mN}),$$

which is compact for the same reasons that $\tilde{\mathcal{C}}_N$ in (42) is compact. The set $\Pi_N^{df}(x)$ can then be written as

$$\Pi_N^{df}(x) = \tilde{\Pi}_N^{df}(x) + (\mathcal{M}_\perp \times \{0\})$$

Since a continuous function always attains its minimum on a compact set, $J_N(x, \cdot, \cdot)$ attains its minimum on $\tilde{\Pi}_N^{df}(x)$. We denote this minimizer

$$(\tilde{\mathbf{M}}^*(x), \tilde{\mathbf{v}}^*(x)) := \underset{(\mathbf{M},\mathbf{v}) \in \tilde{\Pi}_N^{df}(x)}{\operatorname{argmin}} \; J_N(x, \mathbf{M}, \mathbf{v}),$$

and will show that this pair also minimizes $J_N(x, \cdot, \cdot)$ over the set $\Pi_N^{df}(x)$ in (21). We assume the contrary, so that there exists some $\mathbf{M}_\perp \in \mathcal{M}_\perp$ such that

$$J_N(x, \tilde{\mathbf{M}}^*(x) + \mathbf{M}_\perp, \tilde{\mathbf{v}}^*(x)) < J_N(x, \tilde{\mathbf{M}}^*(x), \tilde{\mathbf{v}}^*(x))$$

(Note that it is obvious by inspection of $J_N(x, \cdot, \cdot)$ in (24) that modification of $\tilde{\mathbf{v}}^*(x)$ cannot produce a better result). From (24), this implies that

$$\operatorname{tr}(\mathbf{C}_w H_M \mathbf{M}_\perp) + \operatorname{tr}\left((\mathbf{M}^T \mathcal{S} \mathbf{M}_\perp + \mathbf{M}_\perp^T \mathcal{S} \mathbf{M} + \mathbf{M}_\perp^T \mathcal{S} \mathbf{M}_\perp)\mathbf{C}_w\right) < 0 \tag{43}$$

Since $\mathbf{C}_w$ is positive semidefinite, it can be factored as $U \Lambda U^T := \mathbf{C}_w$, where U is a matrix whose columns are the eigenvectors of $\mathbf{C}_w$ with corresponding non-zero eigenvalues. Thus (43) may be rewritten as

$$\operatorname{tr}(H_M \mathbf{M}_\perp U \Lambda U^T) + \operatorname{tr}\left(U^T(\mathbf{M}^T \mathcal{S} \mathbf{M}_\perp + \mathbf{M}_\perp^T \mathcal{S} \mathbf{M} + \mathbf{M}_\perp^T \mathcal{S} \mathbf{M}_\perp) U \Lambda\right) < 0.$$

Since $\mathbf{C}_w$ is assumed well-defined with respect to $\mathcal{W}$ (i.e. the columns of U span $\operatorname{lin}\mathcal{W}$), $\mathbf{M}_\perp U = 0$ for all $\mathbf{M}_\perp \in \mathcal{M}_\perp$, a contradiction. Thus the pair $(\tilde{\mathbf{M}}^*(x), \tilde{\mathbf{v}}^*(x))$ minimizes $J_N(x, \cdot, \cdot)$ over $\Pi_N^{df}(x)$ in (21). ◻

Proof of Proposition 3

Proof. First, define the extended real function $\tilde{J}_N : \mathbb{R}^n \times \mathbb{R}^{mN \times nN} \times \mathbb{R}^{mN} \to \bar{\mathbb{R}}_{\geq 0}$ as

$$\tilde{J}_N(x, \mathbf{M}, \mathbf{v}) := \begin{cases} J_N(x, \mathbf{M}, \mathbf{v}) & (\mathbf{M}, \mathbf{v}) \in \Pi_N^{df}(x) \\ \infty & \text{otherwise} \end{cases} \tag{44}$$

Consider the set

$$\mathcal{P}_N := \{(x, \mathbf{v}) \mid \exists \mathbf{M} \text{ s.t. } (\mathbf{M}, \mathbf{v}, x) \in \mathcal{C}_N\} \tag{45}$$

which is convex and closed, since it is a projection of the convex set $\mathcal{C}_N$ (cf. Remark 1 and Lemma 2), and bounded since $\mathcal{Y}$ is compact. Thus $\mathcal{P}_N$ is compact. Consider also the function

$$p(x, \mathbf{v}) := \inf_{\mathbf{M}} \tilde{J}_N(x, \mathbf{M}, \mathbf{v}) \tag{46}$$

with domain $\mathcal{P}_N$, which is convex on $\mathbb{R}^{mN}$ [18, Prop. 2.22], and which is easily shown to be strictly convex in $\mathbf{v}$ (cf. $J_N(x, \cdot, \cdot)$ in (25), which can be separated into components which are *independent* convex functions of the parameters $\mathbf{M}$ and $\mathbf{v}$, and which is strictly convex in $\mathbf{v}$). The value function $V_N^*(\cdot)$ in (22) can then be written as

$$V_N^*(x) = \inf_{\mathbf{v}} p(x, \mathbf{v}), \tag{47}$$

with compact domain X_N^{sf}, assumed non-empty. This function is convex and lower semi-continuous on its domain [18, Cor. 3.32], and is thus strictly continuous on $\mathrm{int}X_N^{sf}$ [18, Thm. 2.35]. The optimal feedback policy parameter $\mathbf{v}^*(x)$, defined in (21), can likewise be written as

$$\mathbf{v}^*(x) = \operatorname*{argmin}_{\mathbf{v}} p(x, \mathbf{v}) \tag{48}$$

with $\mathrm{dom}\,\mathbf{v}^* = X_N^{sf}$. This function is single-valued on X_N^{sf} and continuous on $\mathrm{int}X_N^{sf}$ [18, Cor. 7.43 and Thm. 3.31]. The uniqueness and continuity properties of $\mu_N(\cdot) = v_0^*(\cdot)$ then follow directly. □

Proof of Proposition 4

Proof. We assume throughout that the condition $W = \{0\}$ does not hold; if $W = \{0\}$, the proof is trivial. It is sufficient to show that there exists a constant γ such that

$$V(f(x, w)) - V(f(x, 0)) \leq \gamma \|w\| \tag{49}$$

for all $x \in \mathcal{X}$ and all $w \in W$. It then follows that $V(f(x,w)) - V(x) = V(f(x,0)) - V(x) + V(f(x,w)) - V(f(x,0)) \leq -\alpha_3(\|x\|) + \gamma \|w\|$, and the conditions of Lemma 1 are satisfied with $\sigma(s) := \gamma \|s\|$.

When the disturbance set W is compact and contains the origin in its (relative) interior, there exists a constant $\rho > 0$ such that

$$\rho := \max \{\epsilon \mid (\mathcal{B}_\epsilon \cap \mathrm{lin}W) \subseteq W\}, \tag{50}$$

where $\mathcal{B}_\epsilon := \{x \mid \|x\| \leq \epsilon\}$. Thus ρ is the size of the smallest vector on the (relative) boundary of W. Note that when W has a non-empty interior, this simplifies to

$$\rho = \min \{\|w\| \mid w \in \partial W\}. \tag{51}$$

Since the set $\mathcal{X}$ is compact, (29a) implies that V is upper bounded by a constant $\bar{\beta}$ and lower bounded by 0. Since the set $\mathcal{X}$ is robust positively invariant, it follows that

$$g(x) \in \tilde{\mathcal{X}} := \mathcal{X} \sim W, \tag{52}$$

where $\mathcal{X} \sim W$ denotes the Pontryagin difference, i.e.

$$\mathcal{X} \sim W := \{x \in \mathbb{R}^n \mid x + w \in \mathcal{X}, \forall w \in W\}. \tag{53}$$

Finding a suitable γ in (49) is equivalent to finding one which satisfies

$$V(\tilde{x} + w) - V(\tilde{x}) \leq \gamma \|w\|, \quad \forall \tilde{x} \in \mathcal{X}, \ \forall w \in W. \tag{54}$$

Since W is convex and compact, for any given $w \in W$ there exists a $\tilde{w}$ on the (relative) boundary of W such that $w = \tau\tilde{w}$ with $0 \leq \tau \leq 1$. Note also that $\tau = \|w\| / \|\tilde{w}\| \leq \|w\| / \rho$. Since $\mathcal{X}$ is robust positively invariant, $\tilde{x} + \tilde{w} \in \mathcal{X}$ $\forall \tilde{x} \in \tilde{\mathcal{X}}$. Since V is convex, it follows that $V(\tilde{x} + w) \leq (1 - \tau)V(\tilde{x}) + \tau V(\tilde{x} + \tilde{w})$, or

$$V(\tilde{x} + w) - V(\tilde{x}) \leq \tau(V(\tilde{x} + \tilde{w}) - V(\tilde{x})) \leq (\bar{\beta}/\rho) \|w\|. \tag{55}$$

The proof is completed by selecting $\gamma := \bar{\beta}/\rho$. □

A.2 Matrix Definitions

Define $\mathbf{A} \in \mathbb{R}^{n(N+1)\times n}$ and $\mathbf{E} \in \mathbb{R}^{n(N+1)\times nN}$ as

$$\mathbf{A} := \begin{bmatrix} I_n \\ A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \quad \mathbf{E} := \begin{bmatrix} 0 & 0 & \cdots & 0 \\ I_n & 0 & \cdots & 0 \\ A & I_n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1} & A^{N-2} & \cdots & I_n \end{bmatrix}.$$

The matrix $\mathbf{B} \in \mathbb{R}^{n(N+1)\times mN}$ is defined as $\mathbf{B} := \mathbf{E}(I_N \otimes B)$.

References

[1] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski, “Optimization over state feedback policies for robust control with constraints,” Cambridge University Engineering Department, Technical Report CUED/F-INFENG/TR.494, Aug. 2005, Accepted to *Automatica*. [Online]. http://www-control.eng.cam.ac.uk/

[2] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000, Survey paper.

[3] P. O. M. Scokaert and D. Q. Mayne, “Min-max feedback model predictive control for constrained linear systems,” *IEEE Trans. Automatic Control*, vol. 43, pp. 1136–1142, 1998.

[4] L. Chisci, J. A. Rossiter, and G. Zappa, “Systems with persistent state disturbances: predictive control with restricted constraints,” *Automatica*, vol. 37, pp. 1019–1028, 2001.

[5] Y. I. Lee and B. Kouvaritakis, “Constrained receding horizon predictive control for systems with disturbances,” *International Journal of Control*, vol. 72, no. 11, pp. 1027–1032, Aug. 1999.

[6] D. Q. Mayne, M. M. Seron, and S. V. Raković, “Robust model predictive control of constrained linear systems with bounded disturbances,” *Automatica*, vol. 41, pp. 219–24, 2005.

[7] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis.* Cambridge University Press, 1991.

[8] E. C. Kerrigan and T. Alamo, “A convex parameterization for solving constrained min-max problems with a quadratic cost,” in *Proc. 2004 American Control Conference*, Boston, MA, USA, June 2004.

[9] J. Löfberg, “Minimax approaches to robust model predictive control,” Ph.D. dissertation, Linköping University, Apr 2003.

[10] A. Bemporad, F. Borrelli, and M. Morari, “Min-max control of constrained uncertain discrete-time linear systems,” *IEEE Transactions on Automatic Control*, vol. 48, no. 9, pp. 1600–1606, September 2003.

[11] A. Bemporad, M. Morari, V. Dua, and E. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, pp. 3–20, 2002.

[12] F. Borrelli, *Constrained Optimal Control of Linear and Hybrid Systems*, ser. Lecture Notes in Control and Information Sciences. Germany: Springer, 2003, vol. 290.

[13] Z. Jiang and Y. Wang, "Input-to-state stability for discrete-time non-linear systems," *Automatica*, vol. 37, pp. 857–869, 2001.

[14] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, pp. 1747–1767, 1999.

[15] I. Kolmanovsky and E. G. Gilbert, "Theory and computations of disturbance invariant sets for discrete-time linear systems," *Mathematical Problems in Engineering*, vol. 4, pp. 317–363, 1998.

[16] H. van de Water and J. C. Willems, "The certainty equivalence property in stochastic control theory," *IEEE Trans. Automatic Control*, vol. 26, no. 5, pp. 1080–1086, Oct. 1981.

[17] K. J. Åström and B. Wittenmark, *Computer-Controlled Systems: Theory and Design*, 3rd ed. Prentice Hall, 1997.

[18] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*. Springer-Verlag, 1998.