

Performance Bounds of Model Predictive Control for Unconstrained and Constrained Linear Quadratic Problems and Beyond[★]

Yuchao Li^{*} Aren Karapetyan^{**} John Lygeros^{**}
Karl H. Johansson^{*} Jonas Mårtensson^{*}

^{*} *Division of Decision and Control Systems, KTH Royal Institute of Technology, Sweden, (e-mail: yuchao,kallej,jonas1@kth.se).*

^{**} *Automatic Control Laboratory, Swiss Federal Institute of Technology in Zürich, Switzerland, (e-mail: akarapetyan,jlygeros@ethz.ch).*

Abstract: We study unconstrained and constrained linear quadratic problems and investigate the suboptimality of the model predictive control (MPC) method applied to such problems. Considering MPC as an approximate scheme for solving the related fixed point equations, we derive performance bounds for the closed-loop system under MPC. Our analysis, as well as numerical examples, suggests new ways of choosing the terminal cost and terminal constraints, which are *not* related to the solution of the Riccati equation of the original problem. The resulting method can have a larger feasible region, and cause hardly any loss of performance in terms of the closed-loop cost over an infinite horizon.

Keywords: Model predictive control, optimal control theory.

1. INTRODUCTION

Model predictive control (MPC) is a well-established scheme for constrained optimal control problems with continuous state and control spaces. In its most widely adopted form (Sokaert and Rawlings, 1998) a system model is used to predict its performance over some finite number of stages, and a terminal cost is added to account for the trajectory beyond this prediction horizon. Moreover, an additional constraint is imposed on the state at which the prediction ends. Suitably designed terminal cost and constraint lead to desirable properties (Mayne et al., 2000). A recent monograph (Bertsekas, 2022b) shows that this structure of MPC is quite similar to the algorithms applied in the high-profile successes in the field of reinforcement learning (Silver et al., 2017). It has also developed a conceptual framework centered around dynamic programming (DP), which is useful for analyzing the performance of MPC.

In this work, we apply the tools introduced in (Bertsekas, 2022b), and investigate the impact of the terminal cost on the performance of MPC measured by the closed-loop cost accumulated over infinite stages. Using unconstrained and constrained linear quadratic regulation (LQR) problems as a vehicle, we derive performance bounds of MPC applied to those problems compared against optimal control. The insights gained from our analysis suggest new designs of those terminal ingredients, which likely make MPC

feasible for a larger set of states, while cause little to no degradation in its performance.

For the LQR problem, a performance bound of MPC is given (Bitmead et al., 1985). It relies on the monotonicity property, which holds well beyond the problem considered here, as discussed in (Bertsekas, 1975, 1977). A suboptimality analysis that is similar in purpose to our study is reported in (Grune and Rantzer, 2008). It relies in part on a relaxed form of DP introduced in (Rantzer, 2006). Our work makes different assumptions, and leads to new design choices of the terminal ingredients, and may be regarded as complementary to that of (Grune and Rantzer, 2008).

Unlike MPC where the performance bound analysis is relatively rare, bounds are pervasive in the study of Markovian decision problems (MDP), with some classical results presented in (Denardo, 1967). (Bertsekas, 2022b) takes an abstract approach that unifies the MDP and the problems addressed by MPC. It regards MPC as *approximation in value space*, where the *approximation* refers to the fact that terminal ingredients act together to approximate the real optimal cost. It also points out that the typical bounds in MDP may be too conservative to reflect the true performance, which we confirm by deriving better bounds for MPC.

Another concept that is related to our study is the investigation of *regret* in dynamical systems. Regret is a non-asymptotic performance metric for optimal control problems involving uncertainty. Given a controller, regret is measured by comparing its associated cost over a finite number of stages with that of a known policy with full or partial knowledge of the uncertainty. The regret of an MPC controller for unconstrained problems with process noise has been analyzed in (Yu et al., 2020), for cases where

[★] This work was supported by the Swedish Foundation for Strategic Research, the Swedish Research Council, and the Knut and Alice Wallenberg Foundation. It is also supported by the Swiss National Science Foundation under NCCR Automation (grant agreement 51NF40.180545).

the terminal cost is related to the solution of the Riccati equation. Other related work includes (Muthirayan et al., 2021) and (Wabersich and Zeilinger, 2020). Although the major challenge addressed by these works comes from uncertainty or partial information, the results echo our conclusion that much of the credit for good performance goes to the MPC structure itself. This provides an alternative perspective for our work and may point to a direction of extensions of the results reported here.

In summary, this paper makes the following contributions:

- (1) We derive performance bounds for MPC applied to unconstrained and constrained LQR problems, as well as nonlinear systems;
- (2) We propose a new design for the terminal ingredients, which likely leads to a larger region where MPC is feasible;
- (3) We conduct numerical studies that verify our theoretical analysis.

2. PERFORMANCE BOUNDS FOR THE LINEAR QUADRATIC PROBLEM

2.1 Preliminaries

We consider discrete-time optimal control problems with system dynamics

$$x_{k+1} = Ax_k + Bu_k, \quad k = 0, 1, \dots, \quad (1)$$

and k -th stage cost defined

$$x'_k Q x_k + u'_k R u_k, \quad (2)$$

where prime denotes transposition, $R > 0$,¹ $Q \geq 0$, (A, B) stabilizable, and (A, \sqrt{Q}) detectable. The optimal feedback law [with respect to cost (2) accumulated over an infinite number of stages] is then given by

$$L^* x = -(B' K^* B + R)^{-1} B' K^* A x, \quad (3)$$

where K^* is the unique positive definite solution fulfilling the algebraic equation

$$K^* = A'(K^* - K^* B(B' K^* B + R)^{-1} B' K^*) A + Q. \quad (4)$$

For this problem, we introduce the Bellman operator F defined for some symmetric matrices as

$$F(K) = A'(K - KB(B'KB + R)^{-1}B'K)A + Q. \quad (5)$$

The ℓ -fold composition of the operator F is denoted by F^ℓ . In addition, given a policy $\mu(x) = Lx$, we define the L -Bellman operator as

$$F_L(K) = (A + BL)'K(A + BL) + Q + L'RL. \quad (6)$$

Similarly, the ℓ -fold composition of the operator F_L is denoted by F_L^ℓ . A matrix A will be called a *stability matrix* if all of its eigenvalues are less than one in absolute value. Let \mathbb{S}^n denote the space of $n \times n$ real symmetric matrices. For the operators F and F_L , the following classical results hold.

- Proposition 1.* (a) Let $K_1, K_2 \in \mathbb{S}^n$ so that $0 \leq K_1 \leq K_2$. Then we have that $F_L(K_1) \leq F_L(K_2)$ and $F(K_1) \leq F(K_2)$.
- (b) The equation $K = F(K)$ admits a unique positive definite solution K^* . Moreover, $F^k(K) \rightarrow K^*$ as $k \rightarrow \infty$ for $K \in \mathbb{S}^n$ so that $K \geq 0$.

¹ If X and Y are symmetric positive semidefinite matrices, then the notation $X > Y$ ($X \geq Y$) means that the matrix $X - Y$ is positive (semi)definite.

- (c) Let L be an m by n matrix such that $(A + BL)$ is a stability matrix. Then the equation $K = F_L(K)$ admits a unique positive definite solution K_L . Moreover, $F_L^k(K) \rightarrow K_L$ as $k \rightarrow \infty$ for $K \in \mathbb{S}^n$ so that $K \geq 0$.

2.2 Region of Decreasing

We will focus on a special set of symmetric matrices, which we call the *region of decreasing*. It is defined as

$$\mathcal{D} = \{K \in \mathbb{S}^n \mid F(K) \leq K\}. \quad (7)$$

This set of matrices plays an important role in the design of model predictive control, which is a consequence of the definition. The following result can be proved by applying Theorems 10.4 and 10.6 from (Bitmead and Gevers, 1991).

Proposition 2. For all $K \in \mathcal{D}$, $F^k(K) \in \mathcal{D}$ for all k ; moreover, $K \geq K^* \in \mathcal{D}$.

Given a matrix $K \in \mathcal{D}$ and a positive integer ℓ , we are interested in solving the problem via the following scheme:

$$\begin{aligned} \min_{\{u_k\}_{k=0}^{\ell-1}} \quad & x'_\ell K x_\ell + \sum_{k=0}^{\ell-1} x'_k Q x_k + u'_k R u_k \\ \text{s. t.} \quad & x_{k+1} = Ax_k + Bu_k, \quad k = 0, \dots, \ell - 1, \\ & x_0 = x \end{aligned} \quad (8)$$

for all x . If the minimum is attained at $(\tilde{u}_0, \tilde{u}_1, \dots, \tilde{u}_{\ell-1})$, the scheme defines a policy $\tilde{\mu}$ by setting $\tilde{\mu}(x) = \tilde{u}_0$; finite horizon LQR intuition shows that there exists a \tilde{L} such that $\tilde{\mu}(x) = \tilde{L}x$. If we apply the Bellman operator and L -Bellman operator introduced in (5) and (6) respectively, the policy $\tilde{\mu}(x) = \tilde{L}x$ can be defined as

$$F_{\tilde{L}}(F^{\ell-1}(K)) = F^\ell(K). \quad (9)$$

Under our assumption, for $K \in \mathcal{D}$, the resulting closed-loop system $(A + B\tilde{L})$ is stable. This is proved in (Bitmead and Gevers, 1991, Theorem 10.19), and its generalization and connection to performance bound is given in (Bertsekas, 2022b, Section 3.3). Then there exists a unique positive semidefinite solution to the equation $F_{\tilde{L}}(K) = K$, which we denote as $K_{\tilde{L}}$. In the rest of this section, we investigate the suboptimality of the scheme (8), which is measured through the difference between the infinite horizon cost obtained by the resulting policy \tilde{L} (encoded through the matrix $K_{\tilde{L}}$) and the optimal infinite horizon cost (encoded through the matrix K^*).

To investigate the properties of K , we introduce some additional notation. $\|M\|$ stands for the induced 2-norm of the matrix M . Following (Bertsekas, 2022a, p. 382) we denote by $\|M\|_s$ a special weighted Euclidean norm such that $\|A + B\tilde{L}\|_s = \sqrt{\rho} < 1$ for some $\rho \in (0, 1)$. For those two norms, there exist some positive constants c_1 and c_2 such that

$$c_1 \|M\| \leq \|M\|_s \leq c_2 \|M\| \quad (10)$$

holds for all n by n matrices M . An analytical formula for computing c_1 and c_2 is given in the Appendix A of the extended version of this work (Li et al., 2022). Both norms have the submultiplicative property. In addition, we define positive integers α and β_ℓ by $\alpha = \min\{\|A + BL^*\|^2, 1\}$, and $\beta_\ell = \min\{\|(A + BL^*)^{\ell-1}\|^2, 1\}$, where ℓ is the horizon length in (8). Then by the submultiplicative property of the norm, we have that $\beta_\ell \leq \alpha^{\ell-1}$. The norm $\|\cdot\|$ also

connects to positive semidefiniteness through the following lemma.

Lemma 3. Let K_1 and K_2 be symmetric matrices such that $K_1 \geq K_2 \geq 0$. Then $\|K_1\| \geq \|K_2\|$.

Now we are ready to state our first result.²

Proposition 4. For all $K \in \mathcal{D}$ and all integers i ,

$$\|F^i(K) - K^*\| \leq \alpha^i \|K - K^*\|. \quad (11)$$

Moreover, if j satisfies $i = (\ell - 1)j$, then

$$\|F^i(K) - K^*\| \leq \beta_\ell^j \|K - K^*\| \leq \alpha^i \|K - K^*\|. \quad (12)$$

2.3 Performance Bound via Contraction

From the above results we can obtain a performance bound for $K_{\bar{L}}$ by exploiting the contraction properties of the operators F and $F_{\bar{L}}$.

Proposition 5. Let $K \in \mathcal{D}$ and $\tilde{\mu}(x) = \tilde{L}x$ defined as in (9). Then we have that

$$\|K_{\bar{L}} - K^*\| \leq \frac{c_2}{c_1(1-\rho)} \left(\rho + \frac{c_2}{c_1} \alpha \right) \beta_\ell \|K - K^*\|. \quad (13)$$

Remark 6. Our proof has followed closely the approach for establishing the classical error bound for the MDP, cf. (Bertsekas, 2022a, Prop.2.2,1). In particular, if $c_1 = c_2$ and $\alpha = \rho$, then the bound established above can be obtained via applying Prop. 2.1.1(e) and Prop. 2.2.1 in (Bertsekas, 2022a, Prop.2.2,1), and the inequality (13) reduces to the known result in Prop. 2.2.1. This bound indicates that by selecting K close to K^* , the MPC policy obtained by solving (8) will perform close to the optimal policy in the infinite horizon.

2.4 Performance Bounds via Monotonicity and Newton Step Interpretation

Alternative bounds can be obtained based on monotonicity and the Newton step interpretation. We first derive the bound that relies solely on the monotonicity, and then provide an additional bound that is obtained by combining the two properties.

Proposition 7. Let $K \in \mathcal{D}$ and $\tilde{\mu}(x) = \tilde{L}x$ defined as in (9). Then we have that

$$\|K_{\bar{L}} - K^*\| \leq \alpha \beta_\ell \|K - K^*\|. \quad (14)$$

Remark 8. Note that by definition $c_2 \geq c_1$, and focusing on the last term of (13), we see that the bound in Prop. 7 is always tighter than the one in Prop. 5. This is consistent with what is noted in (Bertsekas, 2022b, Appendix A.3). Still, the result given in Prop. 5 resembles the classical performance bound in MDP, and holds true for K where $K \notin \mathcal{D}$, while the bound in Prop. 7 only holds for $K \in \mathcal{D}$.

Apart from the monotonicity properties of the operators F and $F_{\bar{L}}$, the Newton step interpretation of computing $K_{\bar{L}}$ can also be brought to bear for establishing another bound on the suboptimality of $K_{\bar{L}}$. For completeness, a slight generalization of the following lemma adopted from (Hewer, 1971, Theorem 2) and its proof are provided in the Appendix B of the extended version (Li et al., 2022).

² In the interest of space, all the proofs are neglected and can be found in the extended version of this work (Li et al., 2022).

Lemma 9. Let $K \in \mathcal{D}$ and $\tilde{\mu}(x) = \tilde{L}x$ defined as in (9). Then there exists $\gamma > 0$ such that

$$\|K_{\bar{L}} - K^*\| \leq \gamma \|F^{\ell-1}(K) - K^*\|^2. \quad (15)$$

An explicit formula for γ is provided in Appendix B of (Li et al., 2022). Combining Lemma 9 and Prop. 4 leads to the following bound.

Proposition 10. Let $K \in \mathcal{D}$ and $\tilde{\mu}(x) = \tilde{L}x$ defined as in (9). Then there exists $\gamma > 0$ such that

$$\|K_{\bar{L}} - K^*\| \leq \gamma \beta_\ell^2 \|K - K^*\|^2. \quad (16)$$

Remark 11. The quadratic term $\|K - K^*\|^2$ in (16) manifests the nature of scheme (8) as a step of the Newton's method for solving (4). In addition, if $\gamma \beta_\ell \|K - K^*\| / \alpha < 1$, then the bound based on Newton step interpretation is tighter than the bound provided by Prop. 7. This is likely to happen if the value ℓ is large, which in turn makes the value β_ℓ small.

Remark 12. Since the bound is built upon Lemma 9 and Prop. 4, it is more conservative than the bound given in Lemma 9. Still, it connects the matrix K (which is a design parameter), the optimal policy μ^* (through the parameter β_ℓ) and the performance of $K_{\bar{L}}$.

3. PERFORMANCE BOUND FOR CONSTRAINED PROBLEMS

Let us now consider problems involving both state and control constraints. For the same stationary dynamics (1) and stage cost (2), there are state constraint $\hat{X} \subset X = \mathbb{R}^n$ and control constraint $U \subset \mathbb{R}^m$. We assume that both \hat{X} and U are compact and convex, and contain the origin in their interior.

The problem can be modeled as an optimal control problem involving stationary dynamics (1) and stage cost

$$g(x_k, u_k) = x_k' Q x_k + u_k' R u_k + \delta_{\hat{X}}(x_k), \quad (17)$$

where $\delta_{\hat{X}}$ is an indicator function that maps x_k to 0 if $x_k \in \hat{X}$, and ∞ otherwise. We consider *stationary* policies μ , which are functions mapping X to U .

The *cost function* of a policy μ , denoted by J_μ , maps X to $[0, \infty]$, and is defined at any initial state $x_0 \in X$, as

$$J_\mu(x_0) = \sum_{k=0}^{\infty} g(x_k, \mu(x_k)), \quad (18)$$

subject to $x_{k+1} = Ax_k + B\mu(x_k)$, $k = 0, 1, \dots$. The optimal cost function J^* is defined pointwise as

$$J^*(x_0) = \inf_{\substack{u_k \in U, k=0,1,\dots \\ x_{k+1} = Ax_k + Bu_k, k=0,1,\dots}} \sum_{k=0}^{\infty} g(x_k, u_k). \quad (19)$$

A stationary policy μ^* is called *optimal* if $J_{\mu^*}(x) = J^*(x)$, $\forall x \in X$. It has the property that

$$\mu^*(x) \in \min_{u \in U} \{g(x, u) + J^*(Ax + Bu)\}, \quad (20)$$

if the minimum in (20) can be attained. For the problem considered here, it may be shown that there is a stationary optimal policy. A brief discussion regarding the existence of a stationary optimal policy is provided in Appendix C of (Li et al., 2022).

Similar to the scheme in (8), we may apply MPC of the following form to solve the problem

$$\begin{aligned} \min_{\{u_k\}_{k=0}^{\ell-1}} \quad & x'_\ell K x_\ell + \sum_{k=0}^{\ell-1} x'_k Q x_k + u'_k R u_k \\ \text{s.t.} \quad & x_{k+1} = Ax_k + Bu_k, \quad k = 0, \dots, \ell - 1, \\ & x_k \in \hat{X}, \quad k = 0, \dots, \ell - 1, \\ & u_k \in U, \quad k = 0, \dots, \ell - 1, \\ & x_\ell \in S \subset \hat{X}, \quad x_0 = x, \end{aligned} \quad (21)$$

where K is a suitably designed matrix, and S is a suitably designed set. If the minimum of (21) is attained at $(\tilde{u}_0, \tilde{u}_1, \dots, \tilde{u}_{\ell-1})$, then the scheme defines a suboptimal policy by setting $\tilde{\mu}(x) = \tilde{u}_0$. As in the unconstrained case, we would like to investigate the suboptimality of $\tilde{\mu}$, depending on the choice of K and S , where suboptimality of the policy $\tilde{\mu}$ is measured by $J_{\tilde{\mu}} - J^*$.

Due the presence of state constraint \hat{X} and control constraint U , the analysis given in Section 2 does not apply as it is. Still, by suitably modifying the relevant definitions, a similar qualitative analysis remains valid. In particular, we denote by $\mathcal{E}^+(X)$ the set of all functions $J : X \mapsto [0, \infty]$. A mapping that plays the key role in our development is the *Bellman operator* $T : \mathcal{E}^+(X) \mapsto \mathcal{E}^+(X)$, defined pointwise as

$$(TJ)(x) = \inf_{u \in U} \{g(x, u) + J(Ax + Bu)\}. \quad (22)$$

This operator is well-posed in view of the nonnegativity (17) of the stage cost. In addition, we denote as T^ℓ the ℓ -fold composition of T with itself, with the convention that $T^0 J = J$. For every fixed policy μ , we also introduce the μ -operator $T_\mu : \mathcal{E}^+(X) \mapsto \mathcal{E}^+(X)$, which is defined pointwise as

$$(T_\mu J)(x) = g(x, \mu(x)) + J(Ax + B\mu(x)). \quad (23)$$

The operators T and T_μ are generalizations of F and F_L . Their relations are extensively discussed in (Bertsekas, 2022b, Chapter 4).

Given a positive semidefinite matrix K and a set S , we can define a function $J \in \mathcal{E}^+(X)$ associated with K and S as

$$J(x) = x'Kx + \delta_S(x). \quad (24)$$

Using the operator notations (22), (23), as well as the function J defined in (24), the policy $\tilde{\mu}$ defined by the scheme (21) can be written succinctly as

$$(T_{\tilde{\mu}}(T^{\ell-1}J))(x) = (T^\ell J)(x), \quad x \in X, \quad (25)$$

which is a generalization of (9).

For the constrained problem, we can also define the region of decreasing, denoted as $\mathcal{D}(X)$. It is given as a set of functions as

$$\mathcal{D}(X) = \{J \in \mathcal{E}^+(X) \mid (TJ)(x) \leq J(x), \forall x \in X\}. \quad (26)$$

3.1 Performance Bounds via Monotonicity and Newton Step Interpretation

For the constrained problem, a result based on contraction, as in Section 2.3 may not be applicable. However, for suitably designed function J , results that parallel Prop. 7 can be established. The proof of the result is given in Appendix D of the extended version of this work (Li et al., 2022).

Proposition 13. Let $J \in \mathcal{D}(X)$ and $\tilde{\mu}$ defined via (25). Then

$$J_{\tilde{\mu}}(x_0) - J^*(x_0) \leq (T^\ell J)(x_0) - J^*(x_0) \leq J(x_\ell) - J^*(x_\ell), \quad (27)$$

for all $x_0 \in X$ with $x_{k+1} = f(x_k, \mu^*(x_k))$, $k = 0, \dots, \ell - 1$.

In addition, if $\hat{J} \in \mathcal{E}^+(X)$ such that $\hat{J}(x_0) \leq J^*(x_0)$ for all $x_0 \in X$, then

$$J_{\tilde{\mu}}(x_0) - J^*(x_0) \leq (T_{\mu^*}^\ell J)(x_0) - (T_{\mu^*}^\ell \hat{J})(x_0) = J(x_\ell) - \hat{J}(x_\ell), \quad (28)$$

for all $x_0 \in X$ with $x_{k+1} = f(x_k, \mu^*(x_k))$, $k = 0, \dots, \ell - 1$.

Remark 14. A close examination of the proof shows that the above bounds hold well beyond the constrained linear systems discussed here. If a stationary optimal policy μ^* exists, the performance bounds hold for the case where control constraints are state dependent, i.e., $u \in U(x)$, or problems involving nonlinear, as well as hybrid systems. In fact, the tools applied here have been used to analyze MPC for hybrid systems in (Baoti et al., 2006).

A Newton's step interpretation of $J_{\tilde{\mu}}$ similar to that of Prop. 10 can also be established. We refer to (Bertsekas, 2022b, Chapter 5) for details.

3.2 New Designs for the Terminal Costs and Constraints

For the MPC scheme of the form (21), a common choice for K has been K^* given in (4), i.e., the optimal cost function $x'K^*x$ for the unconstrained LQR problem. The rationale for such a choice is that the corresponding $\tilde{\mu}$ would be optimal when x is near the origin, thus the performance of $\tilde{\mu}$ over all feasible x should also be near optimal. However, in the presence of control constraints and if the matrix L^* defined in (3) is large, the associated terminal constraint S in (21) can be small, which means the subset of X from where the scheme (21) is feasible (or equivalently, the support of $J_{\tilde{\mu}}$) is small.

On the other hand, our analysis suggests that a choice of K that is near K^* may not be necessary. In fact, one may prefer a much larger K if its associated set S can be larger. A plausible choice that may lead to larger S is the solution of the equation

$$K = A'(K - KB(B'KB + \zeta R)^{-1}B'K)A + Q, \quad (29)$$

where $\zeta \gg 1$. Such a matrix K belongs to the region of decreasing, and it is likely that the corresponding S is enlarged. This is confirmed in our numerical examples.

We note that a large terminal cost is also advocated in (Limón et al., 2006), where the terminal constraint is removed. Our study here has deployed different analytical tools and focuses on the classical form with both terminal cost and constraint. Still, it suggested that a terminal cost that corresponds to a smaller control would be preferred.

4. NUMERICAL STUDIES

In this section, we provide numerical examples to demonstrate the differences between the bounds established in the preceding sections and the actual performance. It is observed that the actual performance is often much better than the bounds suggest, further corroborating our main point that it is *not* necessary to design the terminal cost close to optimal at the expense of reduced support of $J_{\tilde{\mu}}$.

4.1 Numerical Study Results for LQR

Example 15. (Bellman curve for a scalar system). We consider a scalar system. The purpose is to illustrate the performance bound (14) based on monotonicity and the Newton step interpretation, as shown in Fig. 15. The system parameters are $A = 2$, $B = 0.5$, $Q = 1$, $R = 10$. We set $\ell = 1$ so that $K = \bar{K}$. In this case, we have $\|K_{\bar{L}} - K^*\| \approx 3.3$, the bounds (13), (14) and (16) are given as 534.5, 14.4 and 43.0 respectively. The interpretation of Newton's step is evident by noting that $K_{\bar{L}}$ is obtained by constructing a tangent line of $F(K)$ at $F(\bar{K})$.




Fig. 1. The Bellman function F , the L^* -Bellman function F_{L^*} , the performance bound (14), as well as the Newton step interpretation of MPC applied to the scalar LQR problem with $\ell = 1$. It can be seen that $K^* \leq K_{\bar{L}} \leq F(\bar{K}) \leq F_{L^*}(\bar{K})$.

Example 16. (Conservativeness of bounds). We consider a two dimensional double integrator and a four dimensional example from (Kapasouris et al., 1988). The problem data of those examples are listed in Appendix E of (Li et al., 2022). The matrices K in (8) are computed by solving (29) with $\zeta = 50$. Their magnitudes as well as their relations to the corresponding K^* are listed in Table 1.

Table 1. Matrix K

Problem	$\ K\ /\ K^*\ $	$\ K - K^*\ $
2-D	2.5	9.9
4-D	4.3	486

The actual optimality gap $\|K_{\bar{L}} - K^*\|$ and various performance bounds for those two systems with different ℓ are listed in Table 2. It is clear that the actual performance is much better than what the performance bounds suggest. In addition, the bound (13) is always inferior than (14) and (16), which is consistent with our analysis.

Table 2. Performance bounds with different ℓ

Problem	ℓ	$\ K_{\bar{L}} - K^*\ $	(13)	(14)	(16)
2-D	3	10^{-3}	$< 10^{10}$	9.8	553
2-D	10	$< 10^{-13}$	$< 10^5$	$< 10^{-4}$	$< 10^{-7}$
4-D	3	2.8	10^{52}	486	$< 10^{10}$
4-D	10	$< 10^{-3}$	$> 10^{53}$	404	$< 10^{10}$
4-D	20	$< 10^{-7}$	$< 10^{53}$	248	$< 10^9$

4.2 Numerical Study Results for Constrained LQR

Here we investigate the performance of the scheme (21) when both state and control constraints are present. We

introduce constraints to the two- and four-dimensional problems studied in Example 16. In particular, in Example 17, we demonstrate that with our choice of K , the support of $J_{\tilde{\mu}}$ can be enlarged, while there is hardly any loss of optimality.

Example 17. (Enlarging the feasible region). We consider the constrained version of the two-dimensional system investigated in Example 16. We use the same K as in Example 16 and set $\ell = 3$. The sets S designed according to K and K^* are shown in Fig. 2. The feasible regions of the MPC with those terminal constraints are illustrated in the extended version (Li et al., 2022).




Fig. 2. Illustration of terminal sets S designed for K and K^* respectively. It can be seen that for our choice of K , which is larger than K^* , its corresponding terminal set S is much larger.

If we vary the parameter ζ in (29), the resulting K would change as well. The following table shows the ratio of the respective terminal set volumes V_K and V_{K^*} for K and K^* with different values of ζ in (29).

Table 3. Ratio of terminal set volumes with different ζ

ζ	5	15	25	35
V_K/V_{K^*}	1.23	1.56	1.65	1.63

The suboptimality of the MPC scheme with our choice of K (with $\zeta = 50$) and the optimal (approximated in the study by setting $\ell = 100$) is given in Fig. 3. It can be seen that despite a K that is much larger-than K^* , the closed loop system hardly loses any performance.

Example 18. (Performance loss). Finally, we investigate the constrained version of the system studied in Example 16. Using the terminal matrix K applied there and setting $\ell = 2$, we show the optimality gap $J_{\tilde{\mu}}(x) - J^*(x)$ along a trajectory, which is driven by the policy $\tilde{\mu}$, starting from x_0 . The result is shown in Fig. 4. To put the values in context, the optimal cost from x_0 is $J^*(x_0) \approx 293$. Thus, the performance is practically optimal.

5. CONCLUSION

We considered the performance of MPC applied to unconstrained and constrained LQR problems measured by the closed-loop cost accumulated over an infinite horizon. We derived performance bounds that connect the performance




Fig. 3. Illustration of the value $|J_{\tilde{\mu}}(x) - J^*(x)|/|J^*(x)|$ (in percent) with our choice of K . It can be seen that there is hardly any loss of optimality. Consistent with our analysis, the largest loss occurs when x is near the origin. Still, a mere 0.35% loss of cost in comparison with optimal is occurred.




Fig. 4. Illustration of the value $J_{\tilde{\mu}}(x_k) - J^*(x_k)$ where the system is driven under the policy $\tilde{\mu}$. The optimal cost from x_0 is about 293, which means the policy $\tilde{\mu}$ is practically optimal.

of MPC with its terminal cost as well as the true optimal policy of the problems. The derived bounds apply to problems beyond the scope of linear systems and suggest new design of terminal cost and constraint that is not related to the optimal cost of the problem. Numerical studies showed that the new design leads to larger region from which MPC is feasible, while cost little to none performance.

REFERENCES

- Baoti, M., Christophersen, F.J., and Morari, M. (2006). Constrained optimal control of hybrid systems with a linear performance index. *IEEE Transactions on Automatic Control*, 51(12), 1903–1919.
- Bertsekas, D. (2022a). *Abstract dynamic programming*. Athena Scientific.
- Bertsekas, D. (2022b). *Lessons from AlphaZero for Optimal, Model Predictive, and Adaptive Control*. Athena Scientific.
- Bertsekas, D.P. (1975). Monotone mappings in dynamic programming. In *1975 IEEE Conference on Decision and Control including the 14th Symposium on Adaptive Processes*, 20–25. IEEE.
- Bertsekas, D.P. (1977). Monotone mappings with application in dynamic programming. *SIAM Journal on Control and Optimization*, 15(3), 438–464.
- Bitmead, R.R. and Gevers, M. (1991). Riccati difference and differential equations: Convergence, monotonicity and stability. In *The Riccati Equation*, 263–291. Springer.
- Bitmead, R.R., Gevers, M.R., Petersen, I.R., and Kaye, R.J. (1985). Monotonicity and stabilizability-properties of solutions of the riccati difference equation: Propositions, lemmas, theorems, fallacious conjectures and counterexamples. *Systems & Control Letters*, 5(5), 309–315.
- Denardo, E.V. (1967). Contraction mappings in the theory underlying dynamic programming. *Siam Review*, 9(2), 165–177.
- Grune, L. and Rantzer, A. (2008). On the infinite horizon performance of receding horizon controllers. *IEEE Transactions on Automatic Control*, 53(9), 2100–2111.
- Hewer, G. (1971). An iterative technique for the computation of the steady state gains for the discrete optimal regulator. *IEEE Transactions on Automatic Control*, 16(4), 382–384.
- Kapasouris, P., Athans, M., and Stein, G. (1988). Design of feedback control systems for stable plants with saturating actuators. Technical report, Massachusetts Institute of Technology, Laboratory for Information.
- Li, Y., Karapetyan, A., Lygeros, J., Johansson, K.H., and Mårtensson, J. (2022). Performance bounds of model predictive control for unconstrained and constrained linear quadratic problems and beyond. *arXiv preprint arXiv:2211.06187*.
- Limón, D., Alamo, T., Salas, F., and Camacho, E.F. (2006). On the stability of constrained mpc without terminal constraint. *IEEE transactions on automatic control*, 51(5), 832–836.
- Mayne, D.Q., Rawlings, J.B., Rao, C.V., and Sckaert, P.O. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 789–814.
- Muthirayan, D., Yuan, J., Kalathil, D., and Khargonekar, P.P. (2021). Online learning for receding horizon control with provable regret guarantees. *arXiv preprint arXiv:2111.15041*.
- Rantzer, A. (2006). Relaxed dynamic programming in switching systems. *IEE Proceedings-Control Theory and Applications*, 153(5), 567–574.
- Sckaert, P.O. and Rawlings, J.B. (1998). Constrained linear quadratic regulation. *IEEE Transactions on automatic control*, 43(8), 1163–1169.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al. (2017). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.
- Wabersich, K.P. and Zeilinger, M. (2020). Bayesian model predictive control: Efficient model exploration and regret bounds using posterior sampling. In *Learning for Dynamics and Control*, 455–464. PMLR.
- Yu, C., Shi, G., Chung, S.J., Yue, Y., and Wierman, A. (2020). The power of predictions in online control. *Advances in Neural Information Processing Systems*, 33, 1994–2004.