

On receding-horizon approximation in time-varying optimal control

Jintao Sun and Michael Cantoni

Abstract—The closed-loop stability and infinite-horizon performance of receding-horizon approximations are studied for non-stationary linear-quadratic regulator (LQR) problems. The approach is based on a lifted reformulation of the optimal control problem, under assumed uniform controllability and observability, leading to a strict contraction property of the corresponding Riccati operator. Leveraging this contraction property, a stabilizing linear time-varying state-feedback approximation of the infinite-horizon optimal control policy is constructed to meet a performance-loss specification. Its synthesis involves only finite preview of the time-varying problem data at each time step, over a sufficiently long prediction horizon.

Index Terms—Non-stationary discrete-time systems, LQR, model predictive control, Riccati difference equations

I. INTRODUCTION

Consider the following infinite-horizon linear-quadratic regulator (LQR) problem:

$$\min_u \sum_{k=0}^{\infty} x_k' Q_k x_k + u_k' R_k u_k, \quad (1a)$$

where $x_0 = \xi$, and

$$x_{k+1} = A_k x_k + B_k u_k, \quad k \in \{0, 1, 2, \dots\}, \quad (1b)$$

for given problem data (A_k, B_k, Q_k, R_k) and initial state ξ . The task is to determine the cost minimizing infinite-horizon control input sequence $u = (u_0, u_1, \dots)$ and corresponding state sequence $x = (x_0, x_1, \dots)$.

Under stabilizability and detectability conditions, a linear state-feedback characterization of the optimal control policy is well-known [1], [2]. It involves the stabilizing solution of a corresponding non-stationary infinite-horizon Riccati recursion. As such, without ab initio knowledge of the problem data over all time, one can only approximate the optimal policy at each time step.

In this paper, a lifted reformulation of (1) is employed to construct a receding-horizon approximation of the infinite-horizon optimal control policy. The key feature of the approximation is the use of only finite preview of the problem data over a prediction horizon at each time step. It is established that the receding-horizon policy is exponentially stabilizing. Further, explicit bounds are given for setting the prediction horizon length to achieve specified tolerance

This work was supported by a Melbourne Research Scholarship and the Australian Research Council (DP210103272).

J. Sun is a PhD student at the Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, Australia jintaos@student.unimelb.edu.au

M. Cantoni is with the Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, Australia cantoni@unimelb.edu.au

of infinite-horizon performance degradation. This is the main contribution. The receding-horizon approximation is ultimately a linear time-varying state-feedback controller.

Related time-varying work appears in [3] and [4]. In [3], the analysis of infinite-horizon performance degradation is somewhat implicit. By contrast, the explicit bounds provided here can be used for direct synthesis of a receding-horizon approximation that achieves a performance-loss specification. In [4], a Smoothed Online Convex Optimization (SOCO) sensitivity analysis is used to study the performance of a receding-horizon approximation of the optimal policy for a possibly long but finite horizon problem in terms of dynamic regret. The distinguishing feature of the developments presented below relates to the way prediction horizon length is linked to *infinite-horizon* performance loss. In particular, the link is established via a strict contraction property of the time-varying Riccati operator in the lifted domain, as recently established in [5], which builds upon a foundation result from [6]. Recent Riccati contraction based analysis of receding-horizon schemes can be found in [7], [8] for linear time-invariant (i.e., stationary) problems. Earlier related work for stationary nonlinear problems appears in [9].

The paper is structured as follows. Notation and preliminary results are presented next. The lifted reformulation of problem (1) is developed in Section II. Approximation by receding-horizon control policies is discussed in Section III. Results pertaining to exponential stability under receding-horizon control are given in Section IV. Infinite-horizon performance bounds and receding-horizon controller synthesis are then considered in Sections V and VI. Concluding remarks are given in Section VII.

A. Preliminaries

Notation: \mathbb{N} denotes the set of natural numbers, $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$, and for $i \leq j \in \mathbb{N}_0$, $[i : j] := \{k \in \mathbb{N}_0 \mid i \leq k \leq j\}$. The field of real numbers is denoted by \mathbb{R} , and $\mathbb{R}_{>0} := \{\gamma \in \mathbb{R} \mid \gamma > 0\}$. For $m, n \in \mathbb{N}$, the Euclidean space of real-valued n -vectors, with norm $|\cdot|$, and the space of real $m \times n$ matrices, are denoted by \mathbb{R}^n , and $\mathbb{R}^{m \times n}$, respectively. The space \mathbb{R}^n is implicitly associated with $\mathbb{R}^{n \times 1}$. The $n \times n$ identity matrix is denoted by I_n . The $m \times n$ matrix of zeros is denoted by $0_{m,n}$. The transpose of $M \in \mathbb{R}^{n \times n}$ is denoted by $M' \in \mathbb{R}^{n \times m}$; note, $|x|^2 = x'x$ for $x \in \mathbb{R}^n$. When it exists, the inverse of square $M \in \mathbb{R}^{n \times n}$ is denoted by $M^{-1} \in \mathbb{R}^{n \times n}$ (i.e., $M^{-1}M = MM^{-1} = I_n$); when the relevant inverses exist, the Woodbury matrix identity $(M_1 + M_2 M_3 M_4)^{-1} = M_1^{-1} - M_1^{-1} M_2 (M_3^{-1} + M_4 M_1^{-1} M_2)^{-1} M_4 M_1^{-1}$ holds. The induced norm of $M \in \mathbb{R}^{m \times n}$ is $\|M\|_2 := \max_{|x|=1} |Mx|$. The respective subsets of symmetric, positive semi-definite, and

positive definite matrices, are denoted by $\mathbb{S}^n := \{M \in \mathbb{R}^{n \times n} \mid M = M'\}$, $\mathbb{S}_+^n := \{M \in \mathbb{S}^n \mid (\forall x \in \mathbb{R}^n) x'Mx \geq 0\}$, and $\mathbb{S}_{++}^n := \{M \in \mathbb{S}_+^n \mid (\exists \gamma \in \mathbb{R}_{>0})(\forall x \in \mathbb{R}^n) x'Mx \geq \gamma|x|^2\}$. For $M \in \mathbb{S}^n$, all eigenvalues are real; the minimum value is denoted by $\lambda_{\min}(M) = \min_{|x|=1} x'Mx$, and the maximum value by $\lambda_{\max}(M) = \max_{|x|=1} x'Mx$. Also, given $M_1, M_2 \in \mathbb{S}^n$, the notation $M_1 \prec M_2$ (resp. $\preceq M_2$) means $(M_2 - M_1) \in \mathbb{S}_{++}^n$ (resp. $(M_2 - M_1) \in \mathbb{S}_+^n$). For $M \in \mathbb{S}_+^n$, $\lambda_{\min}(M) \geq 0$ and the matrix square-root is denoted by $M^{1/2} \in \mathbb{S}_+^n$ (i.e., $M^{1/2}M^{1/2} = M$). The vector space of \mathbb{R}^n -valued sequences is denoted by $\ell(\mathbb{R}^n)$, and $(\mathbb{R}^n)^m$ denotes the m -times Cartesian product $\mathbb{R}^n \times \dots \times \mathbb{R}^n$. For $w = (w_0, w_1, \dots) \in \ell(\mathbb{R}^n)$ and $a \leq b \in \mathbb{N}_0$, the vector $w_{[a:b]} := (w_a, \dots, w_b) \in (\mathbb{R}^n)^{(b-a+1)}$; note, $w_{[a:a]} = w_a$.

Definition 1. Given $Y, Z \in \mathbb{S}_{++}^n$, the Riemannian distance is

$$\delta(Y, Z) := \sqrt{\sum_{i=1}^n \log^2 \lambda_i}, \quad (2)$$

where $\{\lambda_1, \dots, \lambda_n\} = \text{spec}\{YZ^{-1}\} \subset \mathbb{R}_{>0}$ is the spectrum (i.e., set of eigenvalues) of $YZ^{-1} \in \mathbb{R}^{n \times n}$.

II. A LIFTED REFORMULATION

A lifting approach is used to transform (1) into an LQR problem that is 1-step controllable and observable; see [5]. Given problem data $A_k \in \mathbb{R}^{n \times n}$, $B_k \in \mathbb{R}^{n \times m}$, $Q_k \in \mathbb{S}_+^n$, $R_k \in \mathbb{S}_{++}^m$, for $k \in \mathbb{N}_0$, define the d -step state transition matrix, $d \in \mathbb{N}$, by

$$\Theta_{k,d} := A_{k+d-1}A_{k+d-2} \cdots A_k, \quad (3)$$

with $\Theta_{k,0} := I_n$. The corresponding d -step controllability matrix is defined by

$$C_{k,d} := [\Theta_{k+1,d-1}B_k \quad \Theta_{k+2,d-2}B_{k+1} \quad \dots \quad \Theta_{k+d,0}B_{k+d-1}]. \quad (4)$$

Further, the d -step observability matrix is defined by

$$\mathcal{O}_{k,d} := \begin{bmatrix} C_k \Theta_{k,0} \\ C_{k+1} \Theta_{k,1} \\ \vdots \\ C_{k+d} \Theta_{k,d} \end{bmatrix}, \quad (5)$$

where $C_k := Q_k^{\frac{1}{2}}$.

Assumption 1. There exist $\bar{a}, \bar{b}, \bar{q}, \bar{r} \in \mathbb{R}_{>0}$ such that for all $k \in \mathbb{N}_0$, $\|A_k\|_2 \leq \bar{a}$, $\|B_k\|_2 \leq \bar{b}$, $\|Q_k\|_2 \leq \bar{q}$, and $\|R_k\|_2 \leq \bar{r}$.

Assumption 2. A_k in (1b) is non-singular for all $k \in \mathbb{N}_0$. This is a standard assumption, which holds when (1b) arises from the discretization of continuous-time dynamics, for example.

Assumption 3. There exists $d \in \mathbb{N}$ such that for all $k \in \mathbb{N}_0$, $C_{k,d}$ has full row rank, and $\mathcal{O}_{k,d}$ has full column rank.

With $d \in \mathbb{N}$ as per Assumption 3, define

$$\hat{A}_t := I_{n(d+1)} - \begin{bmatrix} 0_{n,nd} & 0_{n,n} \\ \text{diag}(A_{dt}, \dots, A_{d(t+1)-1}) & 0_{nd,n} \end{bmatrix}, \quad (6a)$$

$$\hat{B}_t := \begin{bmatrix} 0_{n,md} \\ \text{diag}(B_{dt}, \dots, B_{d(t+1)-1}) \end{bmatrix}, \quad (6b)$$

$$\hat{C}_t := [\text{diag}(C_{dt}, \dots, C_{d(t+1)-1}) \quad 0_{nd,n}], \quad (6c)$$

for $t \in \mathbb{N}_0$, where $C_t := Q_t^{\frac{1}{2}}$. With

$$\Phi_t := [0_{n,nd} \quad I_n] \hat{A}_t^{-1} \begin{bmatrix} I_n \\ 0_{nd,n} \end{bmatrix}, \quad (7a)$$

$$\Gamma_t := [0_{n,nd} \quad I_n] \hat{A}_t^{-1} \hat{B}_t, \quad (7b)$$

$$\Xi_t := \hat{C}_t \hat{A}_t^{-1} \begin{bmatrix} I_n \\ 0_{nd,n} \end{bmatrix}, \quad (7c)$$

$$\Delta_t := \hat{C}_t \hat{A}_t^{-1} \hat{B}_t, \quad (7d)$$

the lifted problem data comprises

$$\tilde{Q}_t := \Xi_t' \Xi_t - \Xi_t' \Delta_t \tilde{R}_t^{-1} \Delta_t' \Xi_t, \quad (8a)$$

$$\tilde{R}_t := \text{diag}(R_{dt}, \dots, R_{d(t+1)-1}) + \Delta_t' \Delta_t, \quad (8b)$$

$$\tilde{A}_t := \Phi_t - \Gamma_t \tilde{R}_t^{-1} \Delta_t' \Xi_t, \quad (8c)$$

$$\tilde{B}_t := \Gamma_t. \quad (8d)$$

Remark 1. Under Assumption 2, and with $d \in \mathbb{N}$ as per Assumption 3, \tilde{A}_t is non-singular [5, Lem. 6], \tilde{B}_t is full row rank [5, Rem. 1], and $\tilde{Q}_t \in \mathbb{S}_{++}^n$ [5, Lem. 5], for all $t \in \mathbb{N}_0$.

Assumption 4. The observability and controllability property in Assumption 3 holds uniformly in the sense that $\inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{Q}_t) > 0$, and $\inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{B}_t \tilde{B}_t') > 0$. It is also assumed that $\inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{A}_t \tilde{A}_t') > 0$.

Given $T \in \mathbb{N}$, for $t \in \mathbb{N}_0$, $\chi \in \mathbb{R}^n$, and $w = (w_t, \dots, w_{t+T-1}) \in ((\mathbb{R}^m)^d)^T$, let

$$J_T(t, \chi, w) := \sum_{j=t}^{t+T-1} z_j' \tilde{Q}_j z_j + w_j' \tilde{R}_j w_j \quad (9a)$$

where

$$z_{j+1} = \tilde{A}_j z_j + \tilde{B}_j w_j, \quad j \in [t : t+T-1], \quad (9b)$$

with $z_t = \chi$, and $\tilde{Q}_t, \tilde{R}_t, \tilde{A}_t, \tilde{B}_t$, as given in (8). For $\tilde{u} \in \ell((\mathbb{R}^m)^d)$ define the infinite-horizon cost at time $t \in \mathbb{N}_0$ as

$$J(t, \chi, \tilde{u}) := \lim_{T \rightarrow \infty} J_T(t, \chi, \tilde{u}_{[t:t+T-1]}), \quad (10)$$

and the optimal cost-to-go as

$$V_\infty(t, \chi) := \min_{\tilde{u} \in \ell((\mathbb{R}^m)^d)} J(t, \chi, \tilde{u}). \quad (11)$$

It is well-known (e.g., see [1, Section 3.3] and [2]) that

$$V_\infty(t, \chi) = \chi' P_t \chi, \quad (12)$$

where $(P_t)_{t \in \mathbb{N}_0} \subset \mathbb{S}_+$ is the bounded positive semi-definite solution of the (backward) recursion

$$P_t = \mathcal{R}_t(P_{t+1}), \quad (13)$$

where the Riccati operator is defined by

$$\mathcal{R}_t(P) := \tilde{Q}_t + \tilde{A}'_t(P - P\tilde{B}_t(\tilde{R}_t + \tilde{B}'_tP\tilde{B}_t)^{-1}\tilde{B}'_tP)\tilde{A}_t \quad (14)$$

for $P \in \mathbb{S}^n$; n.b., the recursion does not have a boundary condition. Existence and uniqueness of a bounded positive semi-definite solution $(P_t)_{t \in \mathbb{N}_0}$ is guaranteed under Assumption 3, which implies uniform stabilizability and detectability [10]. The optimal control input $\tilde{u}^* = \arg \min_{\tilde{u} \in \ell((\mathbb{R}^m)^d)} J(0, \xi, \tilde{u})$ corresponds to the linear state-feedback control policy

$$\tilde{u}_t = \mu^*(t, \tilde{x}_t) := -(\tilde{R}_t + \tilde{B}'_tP_{t+1}\tilde{B}_t)^{-1}\tilde{B}'_tP_{t+1}\tilde{A}_t\tilde{x}_t \quad (15)$$

for the dynamics of the lifted system state given by

$$\tilde{x}_{t+1} = \tilde{A}_t\tilde{x}_t + \tilde{B}_t\tilde{u}_t, \quad t \in \mathbb{N}_0, \quad (16)$$

with initial condition $\tilde{x}_0 = \xi$. In particular, $\tilde{u}_t^* = \mu^*(t, \tilde{x}_t)$ [1], [2]. The following result is taken from [5, Lemma 4].

Lemma 1. *The optimal cost associated with the original infinite-horizon problem (1) is equal to $V_\infty(0, \xi) = \xi'P_0\xi$.*

Remark 2. *As shown in the proof of [5, Lemma 4], the solution $u^* \in \ell(\mathbb{R}^m)$ of the original problem (1) can be recovered from (15) as follows:*

$$u_{[dt:d(t+1)-1]}^* = \mu^*(t, \tilde{x}_t) - \tilde{R}_t^{-1}\Delta'_t\Xi_t\tilde{x}_t.$$

Remark 3. *Since $\tilde{R}_t, \tilde{R}_t^{-1} \in \mathbb{S}_{++}^{md}$,*

$$\begin{aligned} D(P) &:= (P - P\tilde{B}_t(\tilde{R}_t + \tilde{B}'_tP\tilde{B}_t)^{-1}\tilde{B}'_tP) \\ &= P^{1/2}(I + P^{1/2}\tilde{B}_t\tilde{R}_t^{-1}\tilde{B}'_tP^{1/2})^{-1}P^{1/2} \in \mathbb{S}_{++}^n \end{aligned}$$

for all $P \in \mathbb{S}_+$; the equality holds by the Woodbury matrix identity. So $\mathcal{R}_t(P) = \tilde{Q}_t + \tilde{A}'_tD(P)\tilde{A}_t \in \mathbb{S}_{++}^n$, because $\tilde{Q}_t \in \mathbb{S}_{++}^n$ as noted in Remark 1. Therefore, the unique positive semi-definite solution $(P_t)_{t \in \mathbb{N}}$ of (13) is positive definite, with $\lambda_{\min}(P_t) \geq \lambda_{\min}(\tilde{Q}_t) > 0$.

Within the time-varying context of this work, it is important to note that implementation of (15) requires knowledge of the lifted problem data $(\tilde{A}_j, \tilde{B}_j, \tilde{Q}_j, \tilde{R}_j)$ for all $j \geq t$, as needed to determine P_{t+1} according to (13), and thus, u_t^* at time $t \in \mathbb{N}_0$. To overcome this impediment, a so-called receding-horizon approximation can be employed. The impact of such an approximation on stability and performance is investigated in the subsequent developments.

III. RECEDING-HORIZON APPROXIMATION

A receding-horizon scheme is presented here to approximate the optimal policy (15) for the lifted reformulation of the original infinite-horizon LQR problem (1); i.e., for (11) with $t = 0$. With $T \in \mathbb{N}$, and $d \in \mathbb{N}$ as per Assumption 3, define $L_T : \mathbb{N}_0 \times \mathbb{R}^n \times ((\mathbb{R}^m)^d)^T \times \mathbb{S}_{++}^n \rightarrow \mathbb{R}_{\geq 0}$ by

$$L_T(t, \chi, w, X) := J_T(t, \chi, w) + z'_{t+T}Xz_{t+T} \quad (17)$$

with J_T and z_{t+T} as per (9) for the given $w = (w_t, \dots, w_{t+T-1})$. Given bounded terminal penalty matrix sequence $(X_{t+T})_{t \in \mathbb{N}_0} \subset \mathbb{S}_{++}^n$, the T -step receding-horizon scheme is defined by the state-feedback control policy

$$\tilde{u}_t = \mu_0^{RH}(t, \tilde{x}_t), \quad (18a)$$

where

$$(\mu_0^{RH}(t, \tilde{x}_t), \dots, \mu_{T-1}^{RH}(t, \tilde{x}_t)) = \arg \min_{w \in ((\mathbb{R}^m)^d)^T} L_T(t, \tilde{x}_t, w, X_{t+T}). \quad (18b)$$

Remark 4. *The receding-horizon policy (18) corresponds to a sampled-data policy in the domain of the original problem:*

$$u_{[td:(t+1)d-1]} = \mu_0^{RH}(t, x_{td}) - \tilde{R}_t^{-1}\Delta'_t\Xi_t x_{td}, \quad (19)$$

where x_k evolves according to (1b) from $x_0 = \xi$.

The quality of this receding-horizon policy, as an approximation of the optimal policy (15), is investigated subsequently in terms of both closed-loop stability and infinite-horizon performance degradation. First, it is noted that the receding-horizon policy (18) recovers the least infinite-horizon cost with particular terminal penalty matrices; however, this requires ab initio knowledge of the problem data over the infinite horizon. Given $T \in \mathbb{N}$, define $W_T : \mathbb{N}_0 \times \mathbb{R}^n \times \mathbb{S}_{++}^n \rightarrow \mathbb{R}_{\geq 0}$ by

$$W_T(t, \chi, X) := \min_{w \in ((\mathbb{R}^m)^d)^T} L_T(t, \chi, w, X), \quad (20)$$

where L_T is defined in (17).

Lemma 2. *For all $t \in \mathbb{N}_0$, and $\chi \in \mathbb{R}^n$,*

$$W_T(t, \chi, X) = \chi'(\mathcal{R}_t \circ \dots \circ \mathcal{R}_{t+T-1}(X))\chi, \quad (21)$$

with $\mathcal{R}_\bullet(\cdot)$ as per (14). In particular,

$$W_T(t, \chi, P_{t+T}) = \chi'P_t\chi = V_\infty(t, \chi), \quad (22)$$

where $(P_t)_{t \in \mathbb{N}_0}$ is the bounded positive definite solution of (13).

Proof. See the proof of [11, Lemma 7]. \square

For given initial state $\xi \in \mathbb{R}^n$, and control input $\tilde{u} \in \ell((\mathbb{R}^m)^d)$, the performance loss with respect to the optimal control input $\tilde{u}^* \in \ell((\mathbb{R}^m)^d)$ as per (15), is defined by

$$\beta(\xi) := J(0, \xi, \tilde{u}) - J(0, \xi, \tilde{u}^*), \quad (23)$$

where the performance index J is given in (10). Note that $\beta(\xi) = J(0, \xi, \tilde{u}) - \xi'P_0\xi$. It quantifies the infinite-horizon performance degradation. The aim here is to obtain an upper bound on (23) for the control input generated by (18), and a method for selecting the prediction horizon $T \in \mathbb{N}$ and penalty matrix sequence to achieve specified performance-loss tolerance.

IV. CLOSED-LOOP STABILITY

For suitable sequences of terminal penalty matrices, the receding-horizon policy (18) is exponentially stabilizing.

Theorem 1. *Given $T \in \mathbb{N}$, let $(X_{t+T})_{t \in \mathbb{N}_0} \subset \mathbb{S}_{++}^n$ be any bounded terminal penalty matrix sequence such that*

$$X_{t+T} \succeq \mathcal{R}_{t+T}(X_{t+T+1}) \quad (24)$$

for all $t \in \mathbb{N}_0$, with $\mathcal{R}_{t+T}(\cdot)$ as per (14). Then, the origin is exponentially stable for (16) under the corresponding receding-horizon state-feedback control policy (18).

Proof. With reference to (20), note that $W_T(t, \tilde{x}_t, X_{t+T}) = \tilde{x}'_t (\mathcal{R}_t \circ \dots \circ \mathcal{R}_{t+T-1}(X_{t+T})) \tilde{x}_t$ by Lemma 2. Thus, in view of Remark 3, and the hypothesis $X_{t+T} \in \mathbb{S}_{++}$, W_T is positive and radially unbounded as a function of \tilde{x}_t . Further, it can be shown that $t \mapsto W_T(t, \tilde{x}_t, X_{t+T})$ is strictly decreasing for the evolution of \tilde{x}_t according to (16) with \tilde{u}_t as per (18). Indeed, given (24),

$$\begin{aligned} & W_T(t+1, \tilde{x}_{t+1}, X_{t+T+1}) - W_T(t, \tilde{x}_t, X_{t+T}) \\ & \leq -\tilde{x}'_t \tilde{Q}_t \tilde{x}_t - \tilde{u}'_t \tilde{R}_t \tilde{u}_t; \end{aligned} \quad (25)$$

see the proof of [11, Thm. 1] for details. This bound is strictly negative for non-zero \tilde{x}_t since $\inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{Q}_t) > 0$ by Assumption 4, and $\tilde{R}_t \succ 0$. Therefore, W_T is a Lyapunov function for the closed-loop dynamics, and the claimed exponential stability property follows from [12, Thm. 5.7]. \square

Next, Theorem 1 is specialized to a one-step policy, which facilitates the presentation of an explicit exponential bound on the closed-loop state trajectory. By the dynamic programming principle [2], for given $T \in \mathbb{N}$, and bounded terminal penalty matrix sequence $(X_{t+T})_{t \in \mathbb{N}_0} \subset \mathbb{S}_{++}^n$, the T -step policy (18) is equivalent to the 1-step policy

$$\begin{aligned} \tilde{u}_t &= \arg \min_{w_t \in (\mathbb{R}^m)^d} L_1(t, \tilde{x}_t, w_t, \hat{X}_{t+1}) \\ &= -(\tilde{R}_t + \tilde{B}'_t \hat{X}_{t+1} \tilde{B}_t)^{-1} \tilde{B}'_t \hat{X}_{t+1} \tilde{A}_t \tilde{x}_t, \end{aligned} \quad (26)$$

where $\hat{X}_{t+1} := \mathcal{R}_{t+1} \circ \dots \circ \mathcal{R}_{t+T-1}(X_{t+T})$ is the terminal penalty matrix for each $t \in \mathbb{N}_0$. Note, in particular, that

$$L_1(t, \tilde{x}_t, w_t, \hat{X}_{t+1}) = (K_t \tilde{x}_t + w_t)' M_t (K_t \tilde{x}_t + w_t) + \tilde{x}'_t \mathcal{R}_t(\hat{X}_{t+1}) \tilde{x}_t,$$

where $K_t := M_t^{-1} \tilde{B}'_t \hat{X}_{t+1} \tilde{A}_t$ and $M_t = \tilde{R}_t + \tilde{B}'_t \hat{X}_{t+1} \tilde{B}_t$. Also note that, with Assumptions 1 and 4,

$$\sup_{t \in \mathbb{N}_0} \|\tilde{Q}_t + \tilde{A}'_t (\tilde{B}_t \tilde{B}'_t)^{-1} \tilde{B}_t \tilde{R}_t \tilde{B}'_t (\tilde{B}_t \tilde{B}'_t)^{-1} \tilde{A}_t\|_2 < +\infty,$$

which implies $\sup_{t \in \mathbb{N}_0} \|\mathcal{R}_t(X)\|_2 < +\infty$ for any $X \in \mathbb{S}_+^n$, since $L_1(t, \chi, -K_t \chi, X) \leq L_1(t, \chi, -\tilde{B}'_t (\tilde{B}_t \tilde{B}'_t)^{-1} \tilde{A}_t \chi, X)$ for all $\chi \in \mathbb{R}^n$, and thus,

$$\mathcal{R}_t(X) \preceq \tilde{Q}_t + \tilde{A}'_t (\tilde{B}_t \tilde{B}'_t)^{-1} \tilde{B}_t \tilde{R}_t \tilde{B}'_t (\tilde{B}_t \tilde{B}'_t)^{-1} \tilde{A}_t. \quad (27)$$

Further, $\inf_{t \in \mathbb{N}_0} \lambda_{\min}(\mathcal{R}_t(X)) > 0$ in view of Remark 3 and Assumption 4.

Theorem 2. Consider the state-feedback control policy (26) for given bounded sequence $(\hat{X}_{t+1})_{t \in \mathbb{N}_0} \subset \mathbb{S}_{++}^n$. Suppose $\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1})$ for all $t \in \mathbb{N}$, with $\mathcal{R}_t(\cdot)$ as per (14). Then, the evolution of (16) with \tilde{u}_t as per (26) satisfies

$$|\tilde{x}_t|^2 \leq \frac{\bar{\omega}}{\underline{\omega}} \left(1 - \frac{\underline{\lambda}}{\bar{\omega}}\right)^t |\tilde{x}_0|^2 \quad (28)$$

for all $t \in \mathbb{N}_0$, with

$$\bar{\omega} := \sup_{t \in \mathbb{N}_0} \|\mathcal{R}_t(\hat{X}_{t+1})\|_2 < +\infty, \quad (29)$$

$$\underline{\omega} := \inf_{t \in \mathbb{N}_0} \lambda_{\min}(\mathcal{R}_t(\hat{X}_{t+1})) > 0, \quad (30)$$

$$\underline{\lambda} := \inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{Q}_t + K'_t \tilde{R}_t K_t) \leq \bar{\omega}, \quad (31)$$

and $K_t := (\tilde{R}_t + \tilde{B}'_t \hat{X}_{t+1} \tilde{B}_t)^{-1} \tilde{B}'_t \hat{X}_{t+1} \tilde{A}_t$. Further, $\underline{\lambda} > 0$.

Proof. In view of the hypothesis $\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1})$, (25), and (26),

$$\begin{aligned} & W_1(t+1, \tilde{x}_{t+1}, \hat{X}_{t+2}) - W_1(t, \tilde{x}_t, \hat{X}_{t+1}) \\ & \leq -\tilde{x}'_t \tilde{Q}_t \tilde{x}_t - \tilde{u}'_t \tilde{R}_t \tilde{u}_t \leq -\underline{\lambda} |\tilde{x}_t|^2, \end{aligned} \quad (32)$$

with $\underline{\lambda}$ as per (31). Note that $\underline{\lambda} \geq \inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{Q}_t) > 0$ by Assumption 4, and that $\underline{\lambda} \leq \bar{\omega}$, because

$$\begin{aligned} \tilde{x}'_t \tilde{Q}_t \tilde{x}_t + \tilde{u}'_t \tilde{R}_t \tilde{u}_t &= \tilde{x}'_t \mathcal{R}_t(\hat{X}_{t+1}) \tilde{x}_t - \tilde{x}'_{t+1} \hat{X}_{t+1} \tilde{x}_{t+1} \\ &\leq \tilde{x}'_t \mathcal{R}_t(\hat{X}_{t+1}) \tilde{x}_t \end{aligned}$$

for $t \in \mathbb{N}_0$. Now, $W_1(t, \tilde{x}_t, \hat{X}_{t+1}) = \tilde{x}'_t \mathcal{R}_t(\hat{X}_{t+1}) \tilde{x}_t$ by Lemma 2, and as such,

$$\underline{\omega} |\tilde{x}_t|^2 \leq W_1(t, \tilde{x}_t, \hat{X}_{t+1}) \leq \bar{\omega} |\tilde{x}_t|^2. \quad (33)$$

Combining (32) and (33) yields

$$W_1(t+1, \tilde{x}_{t+1}, \hat{X}_{t+2}) \leq (1 - \underline{\lambda}/\bar{\omega}) W_1(t, \tilde{x}_t, \hat{X}_{t+1}),$$

for all $t \in \mathbb{N}_0$. Therefore,

$$W_1(t, \tilde{x}_t, \hat{X}_{t+1}) \leq (1 - \underline{\lambda}/\bar{\omega})^t W_1(0, \tilde{x}_0, \hat{X}_1). \quad (34)$$

Finally, using (33) in (34), it follows that

$$\begin{aligned} |\tilde{x}_t|^2 &\leq \frac{1}{\underline{\omega}} W_1(t, \tilde{x}_t, \hat{X}_{t+1}) \leq \frac{1}{\underline{\omega}} (1 - \underline{\lambda}/\bar{\omega})^t W_1(0, \tilde{x}_0, \hat{X}_1) \\ &\leq \frac{\bar{\omega}}{\underline{\omega}} (1 - \underline{\lambda}/\bar{\omega})^t |\tilde{x}_0|^2, \end{aligned}$$

as claimed. \square

V. BOUNDED PERFORMANCE LOSS WITH $T = 1$

As elaborated below, given any bounded 1-step terminal penalty matrix sequence $(\hat{X}_{t+1})_{t \in \mathbb{N}_0}$ that (i) bounds the positive definite solution of (13), and (ii) satisfies $\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1})$, it is possible to bound the infinite-horizon performance loss associated with the corresponding state-feedback control policy (26), relative to the optimal policy (15). The synthesis of such a sequence is considered in Section VI.

Lemma 3. Given $X, P \in \mathbb{S}_{++}^n$, if $X \succeq P$, then $\mathcal{R}_t(X) \succeq \mathcal{R}_t(P)$ for all $t \in \mathbb{N}_0$, with $\mathcal{R}_t(\cdot)$ as per (14).

Proof. This result is taken from [13, Lemma 10.1]. \square

Lemma 4. Given $X, P \in \mathbb{S}_{++}^n$, if $X \succeq P$, then

$$W_1(t, \chi, X) - W_1(t, \chi, P) \leq \|P\|_2 |\chi|^2 (\exp(\delta(X, P)) - 1)$$

for all $t \in \mathbb{N}_0$, and $\chi \in \mathbb{R}^n$, with $\delta(\cdot, \cdot)$ as per (2), and $W_1(\cdot, \cdot, \cdot)$ as per (20).

Proof. See the proof of [11, Lemma 9]. \square

Theorem 3. Consider the state-feedback control policy (26) for given bounded sequence $(\hat{X}_{t+1})_{t \in \mathbb{N}_0} \subset \mathbb{S}_{++}^n$. Suppose

$$\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1}), \quad \hat{X}_t \succeq P_t, \quad \text{and} \quad \delta(\hat{X}_t, P_t) \leq \eta,$$

for all $t \in \mathbb{N}$, with $\eta \in \mathbb{R}_{>0}$, $\delta(\cdot, \cdot)$ as per (2), $(P_t)_{t \in \mathbb{N}_0}$ as the bounded positive definite solution of (13), and $\mathcal{R}_t(\cdot)$ as per

(14). Then, for initial state $\xi \in \mathbb{R}^n$, the performance loss as defined in (23) satisfies

$$\beta(\xi) \leq \frac{\bar{\lambda} \bar{\omega}}{\underline{\lambda} \underline{\omega}} \bar{\omega} (\exp(\eta) - 1) |\xi|^2,$$

with

$$\bar{\lambda} := \sup_{t \in \mathbb{N}_0} \|P_t\|_2 < +\infty, \quad (35)$$

and $\bar{\omega}, \underline{\omega}, \underline{\lambda} \in \mathbb{R}_{>0}$ as per (29), (30), and (31).

Proof. Let \tilde{x}_t evolve from $\tilde{x}_0 = \xi$ according to (16) with the input \tilde{u}_t as per (26). Further, let $l(t, \tilde{x}_t) := \tilde{x}_t' \tilde{Q}_t \tilde{x}_t + \tilde{u}_t' \tilde{R}_t \tilde{u}_t$, noting that

$$l(t, \tilde{x}_t) = W_1(t, \tilde{x}_t, \hat{X}_{t+1}) - \tilde{x}_{t+1}' \hat{X}_{t+1} \tilde{x}_{t+1}, \quad (36)$$

where W_1 is given in (20). Since $W_1(t, \tilde{x}_t, P_{t+1}) = \tilde{x}_t' P_t \tilde{x}_t$ by Lemma 2, and since $\hat{X}_{t+1} \succeq P_{t+1}$ by hypothesis,

$$\begin{aligned} l(t, \tilde{x}_t) &\leq W_1(t, \tilde{x}_t, \hat{X}_{t+1}) - W_1(t, \tilde{x}_t, P_{t+1}) \\ &\quad + \tilde{x}_t' P_t \tilde{x}_t - \tilde{x}_{t+1}' P_{t+1} \tilde{x}_{t+1} \\ &\leq \|P_{t+1}\|_2 |\tilde{x}_t|^2 (\exp(\delta(\hat{X}_{t+1}, P_{t+1})) - 1) \\ &\quad + \tilde{x}_t' P_t \tilde{x}_t - \tilde{x}_{t+1}' P_{t+1} \tilde{x}_{t+1} \\ &\leq \bar{\lambda} |\tilde{x}_t|^2 (\exp(\eta) - 1) + \tilde{x}_t' P_t \tilde{x}_t - \tilde{x}_{t+1}' P_{t+1} \tilde{x}_{t+1}, \end{aligned}$$

where $\bar{\lambda} < +\infty$ as $(P_t)_{t \in \mathbb{N}_0}$ is bounded, the second inequality holds by Lemma 4, and the last follows from (35) and the hypothesis $\delta(\hat{X}_{t+1}, P_{t+1}) \leq \eta$. As such,

$$\sum_{t=0}^N l(t, \tilde{x}_t) \leq \bar{\lambda} (\exp(\eta) - 1) \sum_{t=0}^N |\tilde{x}_t|^2 + \xi' P_0 \xi - \tilde{x}_{N+1}' P_{N+1} \tilde{x}_{N+1},$$

for all $N \in \mathbb{N}$. Since $\tilde{x}_{N+1} \rightarrow 0$ as $N \rightarrow \infty$ by Theorem 2, it follows from (28) that

$$\begin{aligned} J(0, \xi, \tilde{u}) &= \lim_{N \rightarrow \infty} \sum_{t=0}^N l_t(\tilde{x}_t) \\ &\leq \bar{\lambda} (\exp(\eta) - 1) \lim_{N \rightarrow \infty} \sum_{t=0}^N |\tilde{x}_t|^2 + \xi' P_0 \xi \\ &\leq \bar{\lambda} (\exp(\eta) - 1) \frac{\bar{\omega} \bar{\omega}}{\underline{\omega} \underline{\lambda}} |\xi|^2 + \xi' P_0 \xi. \end{aligned}$$

Therefore,

$$\beta(\xi) = J(0, \xi, \tilde{u}) - \xi' P_0 \xi \leq (\exp(\eta) - 1) \frac{\bar{\lambda} \bar{\omega}}{\underline{\lambda} \underline{\omega}} |\xi|^2,$$

as claimed. \square

VI. RECEDING-HORIZON POLICY SYNTHESIS

In this section, it is shown how to set the prediction horizon $T \in \mathbb{N}$, and construct a sequence of terminal penalty matrices, to achieve a performance loss specification for the receding-horizon policy (18). The approach is based on a result, taken from [5], that establishes a strict contraction property of the Riccati operator given in (14).

Given $T \in \mathbb{N}$, the proposed terminal penalty matrix sequence $(X_{t+T})_{t \in \mathbb{N}_0}$ is given by

$$\begin{aligned} X_{t+T} &= \tilde{Q}_{t+T} + \tilde{A}'_{t+T} (\tilde{B}_{t+T} \tilde{B}'_{t+T})^{-1} \tilde{B}_{t+T} \tilde{R}_{t+T} \\ &\quad \times \tilde{B}'_{t+T} (\tilde{B}_{t+T} \tilde{B}'_{t+T})^{-1} \tilde{A}_{t+T} \in \mathbb{S}_{++}^n. \quad (37) \end{aligned}$$

Remark 5. In view of Assumptions 1 and 4, with X_{t+T} as per (37), $\sup_{t \in \mathbb{N}_0} \|X_{t+T}\|_2 < +\infty$.

Lemma 5. Given $T \in \mathbb{N}$, and $(X_{t+T})_{t \in \mathbb{N}_0}$ as per (37), let

$$\hat{X}_{t+1} := \mathcal{R}_{t+1} \circ \dots \circ \mathcal{R}_{t+T-1}(X_{t+T}), \quad (38)$$

for $t \in \mathbb{N}_0$. Then, $\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1})$, and $\hat{X}_t \succeq P_t$, where $(P_t)_{t \in \mathbb{N}_0}$ is the bounded positive definite solution of (13). Further, $\sup_{t \in \mathbb{N}_0} \|\hat{X}_{t+1}\|_2 < +\infty$.

Proof. Given (37), it follows from (27) that $\mathcal{R}_{t+T-1}(X_{t+T}) \preceq X_{t+T-1}$ for all $t \in \mathbb{N}$, and thus, $\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1})$ by Lemma 3. Similarly, $X_{t+T-1} \succeq \mathcal{R}_{t+T-1}(P_{t+T}) = P_{t+T-1}$ for all $t \in \mathbb{N}$, and thus, $\hat{X}_t = \mathcal{R}_t \circ \dots \circ \mathcal{R}_{t+T-2}(X_{t+T-1}) \succeq \mathcal{R}_t \circ \dots \circ \mathcal{R}_{t+T-2}(P_{t+T-1}) = P_t$. Boundedness of $(\hat{X}_{t+1})_{t \in \mathbb{N}_0}$ also follows from (27) and Assumptions 1 and 4. \square

The following is taken from [5, Thm. 1]. It characterizes a strict contraction property of the Riccati operator (14) with respect to the Riemannian metric in Definition 1.

Lemma 6. For all $t \in \mathbb{N}$, and $Y, Z \in \mathbb{S}_{++}^n$,

$$\delta(\mathcal{R}_t(Y), \mathcal{R}_t(Z)) \leq \rho_t \cdot \delta(Y, Z), \quad (39)$$

where $\rho_t = \zeta_t / (\zeta_t + \varepsilon_t) < 1$,

$$\zeta_t = \|(\tilde{Q}_t + \tilde{Q}_t \tilde{A}_t^{-1} \tilde{B}_t \tilde{R}_t^{-1} \tilde{B}_t' (\tilde{A}_t')^{-1} \tilde{Q}_t)^{-1}\|_2, \quad (40a)$$

$$\varepsilon_t = \lambda_{\min}(\tilde{A}_t^{-1} \tilde{B}_t (\tilde{R}_t + \tilde{B}_t' (\tilde{A}_t')^{-1} \tilde{Q}_t \tilde{A}_t^{-1} \tilde{B}_t)^{-1} \tilde{B}_t' (\tilde{A}_t')^{-1}). \quad (40b)$$

Lemma 7. Given $T \in \mathbb{N}$, with $(X_{t+T})_{t \in \mathbb{N}_0}$ as per (37), and $(P_t)_{t \in \mathbb{N}_0}$ as the bounded positive definite solution of (13), the Riemannian distance $\delta(X_{t+T}, P_{t+T}) \leq \bar{\delta}$ for all $t \in \mathbb{N}_0$, where

$$\begin{aligned} \bar{\delta} &:= \sqrt{n} \log \left(\sup_{t \in \mathbb{N}_0} \frac{\|X_{t+T}\|_2}{\lambda_{\min}(P_{t+T})} \right) \\ &\leq \sqrt{n} \log \left(\sup_{t \in \mathbb{N}_0} \frac{\|X_{t+T}\|_2}{\lambda_{\min}(\tilde{Q}_{t+T})} \right) < +\infty. \end{aligned} \quad (41)$$

Proof. See the proof of [11, Lemma 12]. \square

The following result builds upon Theorem 3. For given prediction horizon $T \in \mathbb{N}$, and with the terminal penalty matrix sequence set according to (37), the closed-loop is exponentially stable under the T -step receding-horizon policy (18), with bounded performance loss.

Theorem 4. Given $T \in \mathbb{N}$, and $(X_{t+T})_{t \in \mathbb{N}_0}$ as per (37), consider the receding-horizon state-feedback control policy (18). For all initial states $\xi \in \mathbb{R}^n$, the performance loss as defined in (23) satisfies

$$\beta(\xi) \leq \frac{\bar{\lambda} \bar{\omega}}{\underline{\lambda} \underline{\omega}} \bar{\omega} \left(\exp \left(\left(\frac{\bar{\zeta}}{\bar{\zeta} + \underline{\varepsilon}} \right)^{T-1} \bar{\delta} \right) - 1 \right) |\xi|^2 \quad (42)$$

with

$$\bar{\zeta} := \sup_{t \in \mathbb{N}_0} \zeta_t < +\infty, \quad \underline{\varepsilon} := \inf_{t \in \mathbb{N}_0} \varepsilon_t > 0, \quad (43)$$

and $\bar{\omega}, \underline{\omega}, \underline{\lambda}, \bar{\lambda}, \zeta_t, \varepsilon_t, \bar{\delta} \in \mathbb{R}_{>0}$ as per (29), (30), (31), (35), (40), and (41), using $(\hat{X}_{t+1})_{t \in \mathbb{N}_0}$ as per (38).

Proof. As observed in the preamble to Theorem 2, with $(\hat{X}_{t+1})_{t \in \mathbb{N}_0}$ as per (38), the T -step receding-horizon control

policy (18) is equivalent to the 1-step policy (26). Now, by Lemma 5, $\hat{X}_t \succeq P_t$, and $\hat{X}_t \succeq \mathcal{R}_t(\hat{X}_{t+1})$, for all $t \in \mathbb{N}$. Further,

$$\begin{aligned} & \delta(\hat{X}_t, P_t) \\ &= \delta(\mathcal{R}_t \circ \dots \circ \mathcal{R}_{t+T-2}(X_{t+T-1}), \mathcal{R}_t \circ \dots \circ \mathcal{R}_{t+T-2}(P_{t+T-1})) \\ &\leq \left(\frac{\bar{\zeta}}{\bar{\zeta} + \underline{\varepsilon}}\right)^{T-1} \delta(X_{t+T-1}, P_{t+T-1}) \leq \left(\frac{\bar{\zeta}}{\bar{\zeta} + \underline{\varepsilon}}\right)^{T-1} \bar{\delta}, \end{aligned}$$

by repeated application of Lemma 6. In particular, Theorem 3 applies, which yields the bound (42).

It remains to show that $\bar{\zeta} < +\infty$ and $\underline{\varepsilon} > 0$. First, observe that for all $Z \in \mathbb{S}_{++}^q$, and $Y \in \mathbb{R}^{p \times q}$ with full row rank,

$$\begin{aligned} \lambda_{\min}(YZY') &= \min_{|x|=1} x'YZY'x \\ &= \min_{|x|=1} \frac{(Y'x)'Z(Y'x)}{|Y'x|^2} \frac{x'YY'x}{|x|^2} \geq \lambda_{\min}(Z)\lambda_{\min}(YY'), \end{aligned}$$

and $\lambda_{\max}(YZY') = \max_{|x|=1} x'YZY'x \leq \lambda_{\max}(Z)\lambda_{\max}(YY')$. Further, $\lambda_{\max}(Z^{-1}) = 1/\lambda_{\min}(Z)$, and when $p = q$, $\lambda_{\min}(Z) \leq \lambda_{\min}(Z + YY')$ and $\lambda_{\max}(Z + YY') \leq \lambda_{\max}(Z) + \lambda_{\max}(YY')$. Using these inequalities, it follows from (43), and (40), that $\bar{\zeta} = \sup_{t \in \mathbb{N}_0} 1/\lambda_{\min}(\tilde{Q}_t + \tilde{Q}_t \tilde{A}_t^{-1} \tilde{B}_t \tilde{R}_t^{-1} \tilde{B}_t' (\tilde{A}_t')^{-1} \tilde{Q}_t) \leq 1/\inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{Q}_t)$, and

$$\begin{aligned} \underline{\varepsilon} &\geq \inf_{t \in \mathbb{N}_0} \lambda_{\min}((\tilde{R}_t + \tilde{B}_t' (\tilde{A}_t')^{-1} \tilde{Q}_t \tilde{A}_t^{-1} \tilde{B}_t)^{-1}) \\ &\quad \times \lambda_{\min}(\tilde{B}_t \tilde{B}_t') \lambda_{\min}(\tilde{A}_t^{-1} (\tilde{A}_t')^{-1}) \\ &= \inf_{t \in \mathbb{N}_0} \lambda_{\min}(\tilde{B}_t \tilde{B}_t') / (\lambda_{\max}(\tilde{R}_t + \tilde{B}_t' (\tilde{A}_t')^{-1} \tilde{Q}_t \tilde{A}_t^{-1} \tilde{B}_t) \lambda_{\max}(\tilde{A}_t' \tilde{A}_t)) \\ &\geq \inf_{t \in \mathbb{N}_0} \frac{\lambda_{\min}(\tilde{B}_t \tilde{B}_t')}{\lambda_{\max}(\tilde{A}_t' \tilde{A}_t)} 1 / (\lambda_{\max}(\tilde{R}_t) + \lambda_{\max}(\tilde{Q}_t) \frac{\lambda_{\max}(\tilde{B}_t \tilde{B}_t')}{\lambda_{\min}(\tilde{A}_t \tilde{A}_t)}). \end{aligned}$$

As such, with Assumptions 1 and 4, $\bar{\zeta} < +\infty$ and $\underline{\varepsilon} > 0$. \square

Remark 6. In Theorem 4, $\bar{\omega}$, $\bar{\lambda}$, $\bar{\delta}$, and $\bar{\zeta}$ can be replaced by any corresponding upper bounds for these quantities. Similarly, $\underline{\omega}$, $\underline{\lambda}$, and $\underline{\varepsilon}$ can be replaced by any corresponding lower bounds.

Theorem 4 can be used to set the prediction horizon $T \in \mathbb{N}$ to achieve specified infinite-horizon performance degradation with the terminal penalty sequence given by (37).

Corollary 1. Given $T \in \mathbb{N}$, and $(X_{t+T})_{t \in \mathbb{N}_0}$ as per (37), consider the receding-horizon state-feedback control policy (18). Given $\bar{\beta} \in \mathbb{R}_{>0}$, the performance loss as defined in (23) satisfies $\beta(\xi) \leq \bar{\beta} |\xi|^2$ for all initial states $\xi \in \mathbb{R}^n$, if

$$T \geq \left(\log \left(\log \left(\frac{\bar{\beta} \cdot \bar{\lambda} \cdot \bar{\omega}}{\bar{\omega} \cdot \bar{\lambda} \cdot \bar{\omega}} + 1 \right)^{1/\bar{\delta}} \right) / \log \left(\frac{\bar{\zeta}}{\bar{\zeta} + \underline{\varepsilon}} \right) \right) + 1$$

with $\bar{\omega}$, $\underline{\omega}$, $\bar{\lambda}$, $\underline{\lambda}$, $\bar{\delta}$, $\underline{\delta}$, $\bar{\zeta}$, $\underline{\varepsilon} \in \mathbb{R}_{>0}$ as per (29), (30), (31), (35), (41), and (43), using $(\hat{X}_{t+1})_{t \in \mathbb{N}_0}$ as per (38).

VII. CONCLUSION

A link is established between prediction horizon length and the infinite-horizon performance loss of a receding-horizon approximation of the optimal policy for a lifted reformulation of (1). This is achieved via a strict contraction property of the corresponding Riccati operator, under an

assumed uniform controllability and uniform observability property of the dynamics and stage cost in the original domain. Ongoing work is focused on extending the approach to accommodate cross-terms in the stage cost, uncertainty in the problem data, and constraints on the input and state variables.

REFERENCES

- [1] B. D. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*. Courier Corporation, 2007.
- [2] D. Bertsekas, *Dynamic Programming and Optimal Control: Volume I*. Athena scientific, 2012, vol. 1.
- [3] S. S. Keerthi and E. G. Gilbert, "Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations," *Journal of Optimization Theory and Applications*, vol. 57, pp. 265–293, 1988.
- [4] Y. Lin, Y. Hu, G. Shi, H. Sun, G. Qu, and A. Wierman, "Perturbation-based regret analysis of predictive control in linear time varying systems," *Advances in Neural Information Processing Systems*, vol. 34, pp. 5174–5185, 2021.
- [5] J. Sun and M. Cantoni, "On Riccati contraction in time-varying linear-quadratic control," *arXiv preprint arXiv:2305.06003*, 2023.
- [6] P. Bougerol, "Kalman filtering with random coefficients and contractions," *SIAM Journal on Control and Optimization*, vol. 31, no. 4, pp. 942–959, 1993.
- [7] R. Zhang, Y. Li, and N. Li, "On the regret analysis of online LQR control with predictions," in *2021 American Control Conference (ACC)*, IEEE, 2021, pp. 697–703.
- [8] Y. Li, A. Karapetyan, J. Lygeros, K. H. Johansson, and J. Mårtensson, "Performance bounds of model predictive control for unconstrained and constrained linear quadratic problems and beyond," *arXiv preprint arXiv:2211.06187*, 2022.
- [9] L. Grune and A. Rantzer, "On the infinite horizon performance of receding horizon controllers," *IEEE Transactions on Automatic Control*, vol. 53, no. 9, pp. 2100–2111, 2008.
- [10] G. De Nicolao, "On the time-varying Riccati difference equation of optimal filtering," *SIAM Journal on Control and Optimization*, vol. 30, no. 6, pp. 1251–1269, 1992.
- [11] J. Sun and M. Cantoni, "On receding-horizon approximation in time-varying optimal control," *arXiv preprint arXiv:2305.06010*, 2023.
- [12] N. Bof, R. Carli, and L. Schenato, "Lyapunov theory for discrete time systems," *arXiv preprint arXiv:1809.05289*, 2018.
- [13] R. R. Bitmead and M. Gevers, "Riccati difference and differential equations: Convergence, monotonicity and stability," in *The Riccati Equation*, Springer, 1991, pp. 263–291.