CDS

TECHNICAL MEMORANDUM NO. CIT-CDS 97-001
January, 1997

"Finite Receding Horizon Linear Quadratic Control: A Unifying Theory for Stability and Performance Analysis"

Vesna Nevistic and James A. Primbs

Control and Dynamical Systems
California Institute of Technology
Pasadena, California 91125

Technical Memorandum No. CIT–CDS 97-001

Finite Receding Horizon Linear Quadratic Control: A Unifying Theory for Stability and Performance Analysis

Vesna Nevistić* and James A. Primbs†

January,1997

Abstract

We consider a finite horizon based formulation of receding horizon control for linear discrete-time plants with quadratic costs. A framework is developed for analyzing stability and performance of finite receding horizon control for arbitrary terminal weights. Previous stability and performance results, including end constraints, infinite horizon formulations, and the fake algebraic Riccati equation, are all shown to be special cases of the derived results. The unconstrained case is presented, where conditions for finite receding horizon control to be stabilizing and within specified bounds of the optimal infinite horizon performance can be computed from the solution to the Riccati difference equations. Nevertheless, the framework presented is general in that it lays the groundwork for extension to constrained systems.

Keywords: predictive control, optimal control, linear systems.

1 Introduction

Receding horizon control (RHC), also known as model predictive control (MPC), is a discrete-time technique in which the control action is obtained by solving open loop optimization problems at each time step. The flexibility of this type of implementation has been useful in addressing various implementation issues that traditionally have been problematic.

From a practical viewpoint, an attractive feature of RHC is its ability to naturally and explicitly handle both multivariable input and output constraints by direct incorporation into the optimization. The RHC strategy was first exploited and successfully employed on linear plants, especially in the process industries [12, 5], where relatively slow sample times made extensive on-line intersample computation feasible. Recent improvements in computer power have made RHC an attractive and more viable alternative approach in a variety of additional applications as well.

On the other hand, theoretical aspects associated with stability and performance properties of RHC have proven to be a very complicated and difficult issue. The attractive capabilities of RHC and a promising outlook for future applications has motivated an intensive study of the stability and performance of various RHC implementations, providing a challenging research area.

*Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH), CH-8092 Zürich, Switzerland, e-mail: vesna@aut.ee.ethz.ch

†Control and Dynamical Systems, California Institute of Technology, Pasadena, California 91125, e-mail: jprimbs@cds.caltech.edu. Supported by NSF.

Stability results have been established for special choices of the control parameters, and rely basically on two approaches. For finite horizon based RHC, stability results often require the addition of an end constraint that the state be zero at the end of the prediction horizon [9, 10, 3, 7]. These constraints are somewhat artificial since they are not satisfied in the closed-loop, and the state only asymptotically approaches zero. An alternate approach is to use an infinite horizon formulation. This has been explored by Rawlings and Muske [13] for linear plants and quadratic costs, including the presence of linear state and control constraints.

In this paper, we introduce a framework for analyzing both stability and performance of *finite horizon* based RH policies, *without* using *end constraints*, and *without* relying on *monotonicity* arguments. Following the ideas and approach used by Shamma and Xiong for non-quadratic RHC [14], we derive stability and performance results for linear systems with quadratic costs. The stability results presented in this paper are based on Lyapunov arguments that use the finite receding horizon cost as a Lyapunov function, and lead to performance bounds on the RH policies. In particular, finite horizon computations which calculate the tolerance to which the infinite horizon RHC cost is within that of the optimal infinite horizon cost are provided. Furthermore, our approach appears to be a unifying framework which easily explains the majority of previously established stability results, and provides a natural setting for extending stability and performance analysis to constrained and nonlinear systems. Results which fully extend the ideas in this paper to these more general systems will be submitted for publication shortly.

2 Linear Quadratic Optimal Control

Consider a discrete-time linear system of the form

$$x(k+1) = Ax(k) + Bu(k), \quad x(0) = x_0, \tag{1}$$

where $x(k) \in \mathcal{R}^n$ and $u(k) \in \mathcal{R}^m$ denote the state and control, respectively. A popular design paradigm for linear time-invariant systems is linear-quadratic (LQ) optimal control [8]. The LQ optimal control problem may be posed in either an infinite or finite horizon framework.

2.1 Infinite Horizon Formulation

The *infinite horizon* LQ problem is formulated as follows. Minimize the infinite horizon cost:

$$J(x_0) = \inf_{u(\cdot)} \sum_{k=0}^{\infty} \left(x^T(k)Qx(k) + u^T(k)Ru(k) \right), \tag{2}$$

subject to the system dynamics (1).

Under standard technical assumptions ($[A, B]$ stabilizable, $R > 0$, $Q \geq 0$, $[Q^{1/2}, A]$ observable), the solution is given by:

$$u^*(k) = -(B^TPB + R)^{-1}B^TPAx(k) \tag{3}$$

where P is the unique positive definite solution to the algebraic Riccati equation (ARE):

$$P = A^T[P - PB(B^TPB + R)^{-1}B^TP]A + Q \tag{4}$$

In this case the control u^* is guaranteed to be stabilizing. In addition, P provides a method for evaluating the optimal cost in that $J(x_0) = x_0^TPx_0$.

2.2 Finite Horizon Formulation

The corresponding *finite horizon* problem is defined by the objective function:

$$J_N(x_0) = \inf_{u(\cdot)} \left[x^T(N)P_0x(N) + \sum_{k=0}^{N-1} \left(x^T(k)Qx(k) + u^T(k)Ru(k) \right) \right] \tag{5}$$

Note that P_0 weights the final term of the horizon, $x(N)$, and is usually referred to as a *terminal weight.* The solution of this LQ regulator problem may also be given directly in closed loop form as follows. One merely iterates the Riccati difference equation (RDE) from the initial condition P_0, according to:

$$P_{j+1} = A^T[P_j - P_jB(B^TP_jB + R)^{-1}B^TP_j]A + Q. \tag{6}$$

and implements the finite horizon control policy,

$$u_N^*(k) = -(B^TP_{N-(k+1)}B + R)^{-1}B^TP_{N-(k+1)}Ax(k)$$

Note that the initial condition P_0 used for the RDE is also the terminal weight for the finite horizon cost J_N. We will interchangeably refer to P_0 as an initial condition or terminal weight owing to these two interpretations.

Again, the optimal finite horizon cost can be evaluated from knowledge of P_N by $J_N(x_0) = x_0^TP_Nx_0$.

3 Receding Horizon Formulation: Problem Setup

A receding horizon implementation [6] is typically formulated by introducing the following open-loop optimization problem.

$$J_{(p,m)}(x_0) = \inf_{u(\cdot)} \left[x^T(p)P_0x(p) + \sum_{i=0}^{p-1} x^T(i)Qx(i) + \sum_{i=0}^{m-1} u^T(i)Ru(i) \right] \tag{7}$$

($p \geq m$) where p denotes the length of the *prediction horizon*, and m denotes the length of the *control horizon.* (When $p = \infty$, we refer to this as the *infinite horizon problem*, and similarly, when p is finite, we refer to it as a *finite horizon problem.*)

Let $u^*_{(p,m)}(i)$, $i = 0, \cdots, m-1$ be the minimizing control sequence for $J_{(p,m)}(x(k))$ subject to the system dynamics (1). A receding horizon policy proceeds by implementing *only the first* control $\hat{u}_{(p,m)}(x(k)) = u^*_{(p,m)}(0)$ to obtain $x(k+1) = Ax(k) + Bu^*_{(p,m)}(0)$. The rest of the control sequence $u^*_{(p,m)}$ is discarded and $x(k+1)$ is used to update the optimization problem (7) as a new initial condition. This process is repeated, each time using only the first control action to obtain a new initial condition, then shifting the cost ahead one time step and repeating, hence the name receding horizon control.

In particular, if we consider the case $p = m = N$, then $J_{(p,m)} = J_N$ as defined in (5). This receding horizon policy can then be simply characterized as:

$$\hat{u}_N(x(k)) = \arg\min_u \left\{ x^T(k)Qx(k) + u^TRu + J_{N-1}(Ax(k) + Bu) \right\} \tag{8}$$

Note that the receding horizon policy $\hat{u}_N$ is equal to the initial control action of the solution to the finite horizon LQ problem and can be calculated from the RDE (6). i.e.:

$$\hat{u}_N(x(k)) = -(B^TP_{N-1}B + R)^{-1}B^TP_{N-1}Ax(k) \tag{9}$$

For the remainder of this paper, in order to simplify notation, we make the assumption that $p = m = N$. Receding horizon policies will be denoted as $\hat{u}_N$ and optimal LQ policies as u_N^* and u^*.

4 Stability and Performance of Finite Receding Horizon Linear Quadratic Control

4.1 Motivation and Main Ideas: J_N as a Lyapunov function

Our approach to finite horizon based RH control relies on the use of the finite horizon cost J_N as a Lyapunov function, but without imposing restrictive end constraints or relying on monotonicity arguments to prove stability. In fact, we present a general theory for performance and stability analysis of finite receding horizon LQ control for *any* initialization of the RDE, P_0.

The main argument is as follows. We wish to show that J_N is a Lyapunov function, or equivalently that $J_N(x(k)) - J_N(x(k+1)) > 0$ for $x \neq 0$. Rewriting $J_N(x(k)) - J_N(x(k+1))$ by using that $J_N(x(k)) = x^T(k)Qx(k) + u^{*T}(0)Ru^*(0) + J_{N-1}(x(k+1))$ gives:

$$J_N(x(k)) - J_N(x(k+1)) = \underbrace{[x^T(k)Qx(k) + \hat{u}_N(x(k))R\hat{u}_N(x(k))}_{>0} + \underbrace{J_{N-1}(x(k+1))] - J_N(x(k+1))}_{??} \tag{10}$$

(where we have used that $\hat{u}_N(x(k)) = u_N^*(0)$). If it can be shown that the right hand side of (10) is positive, then stability is proven. Assuming $Q > 0$, the first term

$$[x^T(k)Qx(k) + \hat{u}_N(x(k))R\hat{u}_N(x(k))] \tag{11}$$

is positive. In general, it *cannot* be asserted that the second term

$$[J_{N-1}(x(k+1)) - J_N(x(k+1))] \tag{12}$$

is also positive. But, under appropriate technical conditions, it is known from LQ theory that $J_N \to J$ as $N \to \infty$ where J is the optimal cost for the LQ infinite horizon problem (2). This implies that $(J_{N-1} - J_N) \to 0$ as $N \to \infty$. We will show that by choosing N large enough, it is possible to guarantee that the second term in (10), $[J_{N-1}(x(k+1)) - J_N(x(k+1))]$ is always smaller than the first term $[x^T(k)Qx(k) + u^T(k)Ru(k)]$, independent of $x(k)$. In this case, the right hand side of (10) will be positive, implying that J_N is a valid Lyapunov function and proving stability. Furthermore, it is possible to establish a criterion based on finite-horizon computations relying on the RDE (6), which not only establishes the above stability, but ensures that the cost of the RHC will perform within a specified tolerance of the optimal infinite horizon performance.

Finally, we will show that most other known and well established stability methods, such as finite RHC with end constraints ($x(p) = 0$) [10, 3], or the infinite RHC approach ($p = \infty$) [13], rely on choosing the terminal weight P_0, so that the finite horizon costs, J_N, are *monotonically non–increasing*, i.e. $J_{N-1}(x) \geq J_N(x)$. Note that this trivially results in stability by ensuring that the second term (12), is non–negative. The results we present do not require any such monotonicity arguments and represent a general theory that encompasses many previous results as special cases.

4.2 Preliminaries

Let $\mathcal{Z}^+$ denote the set of non-negative integers. We will consider the discrete-time linear system (1) with $x(k) \in \mathcal{R}^n$ and $u(k) \in \mathcal{R}^m$. In addition we require that $P_0 \geq 0$, $Q > 0$ and $R > 0$, and assume that $[A, B]$ is a stabilizable pair. Finally, recall the following standard result concerning the convergence of the solutions of the RDE (6) to that of the ARE (4).

Proposition 4.1 *Let P_N be the solution of the RDE (6), then $P_N \to P > 0$ as $N \to \infty$ where P is the solution of the algebraic Riccati equation (ARE) (4).*

Proof This is a standard result from LQ theory [1]. ■

We will define two parameters, α_N and ρ_N that will play an important role throughout the rest of the paper. These parameters help to characterize the right hand side of our key equation (10).

Definition 4.1

$$\begin{aligned}\alpha_N &= \min\{\alpha : \alpha J_N(x) \geq J_{N+1}(x),\ \forall x\} \\ &= \min\{\alpha : \alpha P_N \geq P_{N+1}\}\end{aligned}$$

Remark: Note that since $P_N = P_N^T > 0$,

$$\begin{aligned}\alpha P_N \geq P_{N+1} &\Leftrightarrow \alpha P_N - P_{N+1} \geq 0 \\ &\Leftrightarrow \alpha I - P_N^{-1/2} P_{N+1} P_N^{-1/2} \geq 0 \\ &\Leftrightarrow \overline{\lambda}(P_N^{-1/2} P_{N+1} P_N^{-1/2}) \leq \alpha \\ &\Leftrightarrow \overline{\lambda}(P_{N+1} P_N^{-1}) \leq \alpha\end{aligned}$$

where $\overline{\lambda}(\cdot)$ denotes the maximum eigenvalue. Hence, $\alpha_N = \overline{\lambda}(P_{N+1}P_N^{-1})$.

Remark: The parameter α_N is a measure of the maximum ratio by which the finite horizon cost may increase with the addition of a single step. That is, another characterization of α_N is $\alpha_N = \max(J_{N+1}(x)/J_N(x))$. This parameter can be used to characterize monotonicity, for if $\alpha_N < 1$, then this implies that the finite horizon costs are decreasing as the horizon N is increasing. By arguments in the previous subsection, stability is immediately implied. On the other hand, if α_N is not less that one, it still provides an appropriate measure of the amount by which the finite horizon costs may increase, which is exactly the information needed to bound the second term (12).

An important property of α_N is stated in the following proposition.

Proposition 4.2

$$\lim_{N\to\infty} \alpha_N = 1.$$

Proof: Immediate since $P_N \to P$. ■

The second parameter related to equation (10) is defined as follows:

Definition 4.2

$$\begin{aligned}\rho_N &= \max\{\rho : x^T Q x \geq \rho J_N(x),\ \forall x\} \\ &= \max\{\rho : Q \geq \rho P_N\}\end{aligned}$$

Remark: Similar considerations to those given above show that $\rho_N = \underline{\lambda}(QP_N^{-1})$ where $\underline{\lambda}(\cdot)$ denotes the minimum eigenvalue.

Remark: Note that ρ_N determines a lower bound for the first term in (10) in terms of the finite horizon cost J_N:

$$x^T Qx + u^T Ru \geq \rho_N J_N(x).$$

It is also easily seen that ρ_N is the minimum ratio of $x^T Qx$ to $J_N(x)$. Finally, note that since $P_N \geq Q$, ρ_N is always less than or equal to 1 and strictly greater than 0.

The main results presented in this paper follow from simple applications of the main idea, as presented in equation (10), using the parameters α_N and ρ_N as defined above, to bound terms.

4.3 Main Results

In the subsections that follow we present the main stability and performance results which provide sufficient conditions for stability as well as performance bounds for the infinite horizon cost of the RH policy, without any special requirement on the terminal weight P_0. As will be shown later, these theorems provide a framework in which the majority of previous results can be viewed as special cases.

4.3.1 Stability

Consider the receding horizon policy (8) based on the finite horizon cost (5).

Theorem 4.1 *Let N be such that*

$$\gamma_N = \alpha_{N-1}(1 - \rho_N) < 1$$

then the receding horizon policy $\hat{u}_N(\cdot)$ is stabilizing, and $J_N(\cdot)$ is a Lyapunov function for the closed-loop system with

$$J_N(x(k+1)) \leq \gamma_N J_N(x(k))$$

Proof: Let $x(k)$ and $u(k)$ be the state and control trajectory, respectively, resulting from the receding horizon policy

$$u(k) = \hat{u}_N(x(k)).$$

For any $k \in \mathcal{Z}^+$,

$$J_N(x(k)) = x^T(k)Qx(k) + u^T(k)Ru(k) + J_{N-1}(x(k+1)). \tag{13}$$

Therefore,

$$\begin{aligned} J_N(x(k)) - J_N(x(k+1)) &= x^T(k)Qx(k) + u^T(k)Ru(k) + J_{N-1}(x(k+1)) - J_N(x(k+1)) \\ &\geq x^T(k)Qx(k) + u^T(k)Ru(k) + \frac{1}{\alpha_{N-1}} J_N(x(k+1)) - J_N(x(k+1)) \\ &\geq \rho_N J_N(x(k)) + \frac{1}{\alpha_{N-1}} J_N(x(k+1)) - J_N(x(k+1)) \end{aligned}$$

Collecting $J_N(x(k))$ terms on the left, and $J_N(x(k+1))$ on the right gives:

$$(1 - \rho_N) J_N(x(k)) \geq \frac{1}{\alpha_{N-1}} J_N(x(k+1))$$

Multiplying by α_{N-1} gives:

$$\alpha_{N-1}(1-\rho_N)J_N(x(k)) \geq J_N(x(k+1))$$

which completes the proof. ∎

Remark: Note that since $\rho_N \to \underline{\lambda}(QP^{-1}) > 0$, then by Proposition 4.2, there always exists a finite N such that $\gamma_N < 1$. Furthermore, γ_N is computable through finite horizon computations only.

Remark: Just as α_N and ρ_N can be interpreted as ratios of costs, it is clear that γ_N is in fact an upper bound for the ratio $J_N(x(k+1))/J_N(x(k))$.

4.3.2 Performance

Now we will derive bounds on the performance achieved by a stabilizing receding horizon policy.

Theorem 4.2 *Let N and γ_N be as in Theorem 4.1. Denote the infinite horizon performance using the receding horizon policy $u(k) = \hat{u}_N(x(k))$ by:*

$$J_{\hat{u}_N}(x(0)) = \sum_{k=0}^{\infty} x^T(k)Qx(k) + \hat{u}_N^T(x(k))R\hat{u}(x(k)) \tag{14}$$

Then a bound for the infinite-horizon performance is given by:

$$J_{\hat{u}_N}(x(0)) \leq \mathcal{P}_N J_N(x(0)). \tag{15}$$

where

$$\mathcal{P}_N = \left(1 + \left(\frac{\max\{0, \alpha_{N-1}-1\}}{\alpha_{N-1}}\right)\frac{\gamma_N}{1-\gamma_N}\right) \tag{16}$$

Proof: Bounding the cost term-by-term gives the following. First,

$$\begin{aligned}
x^T(0)Qx(0) + u^T(0)Ru(0) &= J_N(x(0)) - J_{N-1}(x(1)) \\
&= J_N(x(0)) - J_N(x(1)) + J_N(x(1)) - J_{N-1}(x(1)) \\
&\leq J_N(x(0)) - J_N(x(1)) + J_N(x(1)) - \frac{1}{\alpha_{N-1}}J_N(x(1)) \\
&\leq J_N(x(0)) - J_N(x(1)) + \left(\tfrac{\alpha_{N-1}-1}{\alpha_{N-1}}\right)J_N(x(1)).
\end{aligned}$$

Similarly,

$$x^T(1)Qx(1) + u^T(1)Ru(1) \leq J_N(x(1)) - J_N(x(2)) + \left(\tfrac{\alpha_{N-1}-1}{\alpha_{N-1}}\right)J_N(x(2)).$$

By a summation of corresponding bounds on $x^T(k)Qx(k) + u^T(k)Ru(k)$ we obtain the following:

$$\begin{aligned}
\sum_{k=0}^{\infty} x^T(k)Qx(k) + u^T(k)Ru(k) &\leq J_N(x(0)) + \left(\tfrac{\alpha_{N-1}-1}{\alpha_{N-1}}\right)\sum_{k=1}^{\infty} J_N(x(k)) \\
&\leq J_N(x(0)) + \left(\tfrac{\max\{0,\alpha_{N-1}-1\}}{\alpha_{N-1}}\right)\sum_{k=1}^{\infty} J_N(x(k)) \\
&\leq \left(1 + \left(\tfrac{\max\{0,\alpha_{N-1}-1\}}{\alpha_{N-1}}\right)\textstyle\sum_{k=1}^{\infty}\gamma_N^k\right)J_N(x(0)) \\
&= \left(1 + \left(\tfrac{\max\{0,\alpha_{N-1}-1\}}{\alpha_{N-1}}\right)\tfrac{\gamma_N}{1-\gamma_N}\right)J_N(x(0)),
\end{aligned}$$

■

Remark: Note that if $\alpha_{N-1} \leq 1$, then $\mathcal{P}_N = 1$ and the above bound reduces to $J_{\hat{u}_N}(x) \leq J_N(x)$. This is implied by a monotonically *non–increasing* cost J_N (or equivalently P_N).

4.3.3 Reformulation of Main results

The results presented in Theorems 4.1 and 4.2 provide a general theory for the performance and stability analysis of finite receding horizon LQ control for *any* initialization (terminal weight) of the RDE, P_0. In fact, these results used virtually none of the special structure present in LQ optimal control. They provide the framework for extending the above results to constrained finite RH LQ control, as well as to other more general systems.

Recall that the parameter γ_N provides only a sufficient condition for J_N to be a Lyapunov function. By using the fact that in the LQ case the RH controller can be characterized in closed form as given in (9), we may actually check *exactly* whether J_N is a Lyapunov function. Additionally this allows us to relax the assumption that $Q > 0$, and include *semi-definite* Q with $[Q^{1/2}, A]$ observable.

Theorem 4.3 *Let N be such that $P_N > 0$ and*

$$\zeta_N = \overline{\lambda}\left(P_N^{-1/2}(A+BK_N)^T P_N (A+BK_N) P_N^{-1/2}\right) < 1 \tag{17}$$

where K_N corresponds to

$$\hat{u}_N(k) = -(B^T P_{N-1} B + R)^{-1} B^T P_{N-1} A x(k) = K_N x(k).$$

Then the receding horizon policy $\hat{u}_N(\cdot)$ is stabilizing. Furthermore, $J_N(\cdot)$ is a Lyapunov function for the closed-loop system, so that

$$J_N(x(k+1)) \leq \zeta_N J_N(x(k)).$$

Proof: Let $x(k)$ be the state trajectory resulting from the receding horizon policy

$$\hat{u}_N(x(k)) = K_N x(k).$$

Then,

$$\begin{aligned}
J_N(x(k)) - J_N(x(k+1)) &= x^T(k) P_N x(k) - x^T(k+1) P_N x(k+1) \\
&= x^T(k) P_N x(k) - x^T(k)(A+BK_N)^T P_N (A+BK_N) x(k) \\
&= x^T(k)\left(P_N - (A+BK_N)^T P_N (A+BK_N)\right) x(k) \\
&\geq \underline{\lambda}\left(P_N^{-1/2}\left(P_N - (A+BK_N)^T P_N (A+BK_N)\right) P_N^{-1/2}\right) J_N(x(k)) \\
&= \underline{\lambda}\left(I - P_N^{-1/2}(A+BK_N)^T P_N (A+BK_N) P_N^{-1/2}\right) J_N(x(k)) \\
&= \left[1 - \overline{\lambda}\left(P_N^{-1/2}(A+BK_N)^T P_N (A+BK_N) P_N^{-1/2}\right)\right] J_N(x(k))
\end{aligned}$$

where we have used arguments similar to those in Section 4.2. Rearranging terms completes the proof. ■

Remark: Since ζ_N provides an exact test of whether J_N is a Lyapunov function, ζ_N may be thought of as the following quantity:

$$\zeta_N = \min \{\zeta : \zeta J_N(x(0)) \geq J_N(x(1)), \quad \forall x(0)\}$$

or as $\zeta_N = \max_{x(0)}(J_N(x(1))/J_N(x(0)))$. Showing that ζ_N may be calculated from the eigenvalue formula in Theorem 4.3 is little more than an application of the remark after Definition 4.1 which showed that α_N could also be characterized by an appropriate eigenvalue computation.

Remark: Clearly, Theorem 4.2 is applicable to the above result by replacing γ_N with ζ_N in equation (16). When we need to distinguish between using γ_N and ζ_N in (16), we will write $\mathcal{P}^\gamma$ and $\mathcal{P}^\zeta$, respectively.

4.4 Using J_{N-1} as a Lyapunov function

As was done with J_N, it is also possible to use J_{N-1} as a Lyapunov function and obtain results analogous to those in Theorems 4.1, 4.2, and 4.3. For completeness, we also present this construction.

Consider using J_{N-1} as a Lyapunov function,

$$\begin{aligned} J_{N-1}(x(k)) - J_{N-1}(x(k+1)) &= J_{N-1}(x(k)) - [J_N(x(k)) - (x^T(k)Qx(k) + \hat{u}_N(x(k))R\hat{u}_N(x(k)))] \\ &= \underbrace{[x^T(k)Qx(k) + \hat{u}_N(x(k))R\hat{u}_N(x(k))]}_{>0} + \underbrace{[J_{N-1}(x(k)) - J_N(x(k))]}_{??} \end{aligned} \tag{18}$$

Similar arguments to those given in Section 4.1 apply.

The following theorems paralleling Theorems 4.1, 4.2, and 4.3 are now easy to prove: Consider the receding horizon policy (8) based on the finite horizon cost (5).

Theorem 4.4 *Let N be such that*

$$\tilde{\gamma}_N = \alpha_{N-1} - \rho_{N-1} < 1$$

then the receding horizon policy $\hat{u}_N(\cdot)$ is stabilizing, and $J_{N-1}(\cdot)$ is a Lyapunov function for the closed-loop system with

$$J_{N-1}(x(k+1)) \leq \tilde{\gamma}_N J_{N-1}(x(k))$$

Proof: See Appendix A. ∎

Theorem 4.5 *Let N and $\tilde{\gamma}_N$ be as in Theorem 4.4. Using the receding horizon policy*

$$u(k) = \hat{u}_N(x(k)),$$

a bound for the infinite-horizon performance is given by:

$$J_{\hat{u}_N}(x(0)) \leq \tilde{\mathcal{P}}_N J_{N-1}(x(0))$$

where

$$\tilde{\mathcal{P}}_N = \left(1 + (\max\{0, \alpha_{N-1} - 1\})\frac{1}{1 - \tilde{\gamma}_N}\right) \tag{19}$$

Proof: See Appendix A. ■

Theorem 4.6 *Let N be such that $P_{N-1} > 0$ and*

$$\tilde{\zeta}_N = \overline{\lambda}\left(P_{N-1}^{-1/2}(A+BK_N)^T P_{N-1}(A+BK_N)P_{N-1}^{-1/2}\right) < 1 \tag{20}$$

where K_N corresponds to

$$\hat{u}_N(k) = -(B^T P_{N-1}B + R)^{-1}B^T P_{N-1}Ax(k) = K_N x(k)$$

The receding horizon policy $\hat{u}_N(\cdot)$ is stabilizing. Furthermore, $J_{N-1}(\cdot)$ is a Lyapunov function for the closed-loop system, so that

$$J_{N-1}(x(k+1)) \leq \tilde{\zeta}_N J_{N-1}(x(k)),$$

Proof: See Appendix A. ■

4.5 Discussion of Results

Before proceeding, we take a moment to review the previous theorems. A summary of the notation used in the stability results is given in Table 1. The theorems rely on using either J_N or J_{N-1} as a Lyapunov function. Note that parameters with no tilde correspond to J_N and those with tilde to J_{N-1}. Furthermore we differentiate between results that did not use the specific structure of the control u, which have parameters γ_N or $\tilde{\gamma}_N$, and those that used u to *exactly* determine whether J_N or J_{N-1} was a Lyapunov function and are denoted by ζ_N and $\tilde{\zeta}_N$. In all cases, a sufficient condition for stability occurs when the corresponding parameter achieves a value less than 1.

Table 1

	without u	exact
J_N	Thm. 4.1 γ_N	Thm. 4.3 ζ_N
J_{N-1}	Thm 4.4 $\tilde{\gamma}_N$	Thm. 4.6 $\tilde{\zeta}_N$

Table 1: Stability Theorems

Furthermore, we would like to note the following. Recall that J denotes the optimal infinite horizon cost for the LQ problem (cf. eqn. (2)), $J_{\hat{u}_N}$ is the infinite horizon cost of the RH policy (cf. eqn. (14)), and J_N is the optimal cost of the finite horizon LQ problem (cf. eqn. (5)).

- If the finite horizon costs J_N are monotonically *non-increasing* (i.e. $\alpha_N \leq 1$), then $\mathcal{P}_N = 1$ which leads to the following bound (cf. Thm. 4.2):

$$J(x(0)) \leq J_{\hat{u}_N}(x(0)) \leq J_N(x(0)). \tag{21}$$

- If the finite horizon costs J_N are monotonically *non–decreasing*, then we know that $J_N(x(0)) \leq J(x(0))$. Combining this with the bound from Theorem 4.2, gives:

$$J_N(x(0)) \leq J(x(0)) \leq J_{\hat{u}_N}(x(0)) \leq \mathcal{P}_N J_N(x(0)) \tag{22}$$

  Hence, it is possible to determine that the infinite horizon performance of the finite RHC is within specified bounds of the optimal performance, by performing the finite horizon calculations of $J_N(x(0))$ and $\mathcal{P}_N$ only.

- In general, the approach using J_N as a Lyapunov function gives less conservative stability results and tighter performance bounds than the results of Section 4.4 which relied on J_{N-1}. Results concerning this can be found in Appendix B.

- Recall that ζ_N and $\tilde{\zeta}_N$ provide exact tests of whether J_N and J_{N-1}, respectively, are Lyapunov functions. Clearly these are only sufficient tests for stability and, as will be seen in the examples, a gap may exists between the horizon length at which the RH policy is stabilizing, and when J_N or J_{N-1} is sufficient to act as a Lyapunov function and guarantee stability. On the other hand, one might question the performance properties of a control action obtained by minimizing J_N, but that does not cause J_N to decrease along its trajectories (i.e. J_N is not a Lyapunov function), even if it is stabilizing. Hence, it could be argued that it is only reasonable to use horizons for which J_N is a Lyapunov function, which furthermore allow for explicit bounds on performance through Theorem 4.2.

Finally, it is worthwhile to mention that the ideas presented here can be naturally generalized to constrained linear and nonlinear systems. These results will be presented in a following paper.

5 A Unifying Framework for Stability Analysis

As was clearly shown in the previous section, since it relies neither on monotonicity, nor end constraints, the finite RH LQ approach provides a general framework for stability and performance analysis. Moreover, it can be considered as a unifying approach which presents a general picture for understanding all previously established stability results. In what follows, the relationship between different stability approaches will be explored in detail. It will also be shown that all other stability approaches can be obtained as special cases of finite RH LQ control.

5.1 Previous results: Monotonicity as an underlying principle

The concept of monotonicity as a key to determining the stability of RHC is an idea which underlies the majority of previous results. Both the finite horizon result using end constraints [9, 10, 3, 7] and the infinite horizon approach [13] can be viewed as different methods for choosing the terminal weight P_0 so that J_N (or equivalently P_N) is monotonically non–increasing (i.e. $\alpha_N \leq 1$), which trivially guarantees stability as described in Section 4.1.

5.1.1 Finite Horizon Approach

Results concerning the stability of finite horizon based RH control for linear systems with quadratic costs have, to date, been primarily based upon the use of end constraints. These end constraints typically consist of forcing the entire state to zero at the end of the prediction horizon ($x(k+p) = 0$). Despite the fact that these constraints guarantee asymptotic stability of the finite horizon based

RH control scheme, as was shown by Kwon and Pearson [10], they have been criticized for being somewhat artificial since they are not achieved in the closed loop.

Bitmead et al. showed that end constraints, or equivalently, $P_0^{-1} = 0$, in a finite horizon RH formulation guarantee that P_N is monotonically non–increasing [4]. In this way, stability is assured for every horizon length $N \geq n$ as stated below in Theorem 5.1. [3, 4]

Theorem 5.1 *Consider the system (1) with dimension n, and assume that $R > 0$, $Q \geq 0$, A invertible and $[A, B]$ controllable. Further, consider $P_0^{-1} = 0$ as the initial condition for the RDE (6). Then the receding horizon policy $\hat{u}_N(x)$ is stabilizing for all $N \geq n$.*

In our approach, the argument that end constraints imply non–increasing monotonicity of J_N (or equivalently P_N) is simple. Consider two finite horizon optimization problems, both with initial condition $x(k)$, but one with cost function $J_N(x(k))$ and end constraint $x(k+N) = 0$, and the second with cost function $J_{N+1}(x(k))$ and end constraint $x(k+N+1) = 0$. Assume $u_N^*(\cdot)$ is the solution of the first finite horizon optimal control problem. Then it is clear that $(u_N^*(\cdot), 0)$ is feasible for the second, and produces a cost for the second problem that is equal to the optimal cost $J_N(x(k))$ that it produces for the first. So the *optimal* cost for the second (i.e. $J_{N+1}(x(k))$) must be at least as small as $J_N(x(k))$. Hence we conclude that $J_{N+1} \leq J_N$ or equivalently $P_{N+1} \leq P_N$. Recall that this monotonicity is equivalent to having $\alpha_N \leq 1$, which was shown to guarantee stability by Thm. 4.1.

As shown previously, our finite RH LQ method does not require this monotonicity to achieve stability, and the end constraint argument appears only as a result of a particular choice of initialization P_0.

5.1.2 Infinite Horizon Approach

Rawlings and Muske showed that nominal closed-loop stability for all choices of the control tuning parameters (Q, R, and control horizon m) can be guaranteed by making the output horizon infinite (i.e. $p = \infty$), and *forcing the unstable modes to be zero at the end of the control horizon.* The following theorem summarizes these results for the unconstrained case [13].

Theorem 5.2 *For stabilizable $[A, B]$ with n_u unstable modes and control horizon $m \geq n_u$, the infinite horizon receding horizon control policy with quadratic objective is stabilizing.*

Recall the infinite horizon based RH formulation as given by equation (7) with $p = \infty$. For *stable plants*, from time step m to ∞, the state evolves according to the uncontrolled dynamics (i.e. $u = 0$). Hence, from $i = m \ldots \infty$, the cost is given by

$$\sum_{i=m}^{\infty} x^T(i) Q x(i) = x^T(m) \left(\sum_{i=0}^{\infty} A^{Ti} Q A^i \right) x(m),$$

and the infinite RHC problem can be reformulated into the finite RHC setup according to the following:

$$J_{(\infty,m)}(x_0) = \inf \left[x^T(m) \left(\sum_{i=0}^{\infty} A^{Ti} Q A^i \right) x(m) + \sum_{i=0}^{m-1} (x^T(i) Q x(i) + u^T(i) R u(i)) \right]$$

Hence the infinite horizon setup is equivalent to that of a finite horizon problem with horizon length m and RDE initialization $P_0 = \sum_{i=0}^{\infty} A^{Ti} Q A^i$, which will be referred to as P_0^{∞}. This initialization can be easily calculated from the discrete Lyapunov equation:

$$A^T P_0^{\infty} A - P_0^{\infty} + Q = 0. \tag{23}$$

This choice of $P_0 = P_0^\infty$ also results in a monotonically non-increasing cost, J_N. Viewing this problem in terms of the infinite horizon cost, it is clear that any state trajectory obtained using a control sequence of length N can be replicated with a control sequence of length $N+1$ merely by adjoining the control $u = 0$ to the end of the sequence of length N. This implies that $J_{N+1}(x) \leq J_N(x)$, demonstrating that the key property of an infinite horizon formulation is that it actually guarantees that the finite horizon cost is monotonically non–increasing.

For an unstable plant, an end constraint that the unstable modes are brought to zero is imposed. Considering a Jordan decomposition of A, the infinite horizon formulation is equivalent to a choice of $P_0 = P_0^\infty$ which equals the combination of $\sum_{i=0}^{\infty} A_s^{Ti} Q A_s^i$ where A_s contains the stable modes of A, and ∞ corresponding to the unstable modes of A. This also leads to a monotonically non-increasing cost, J_N.

From the discussion above, it is clear that the infinite horizon approach also fits into our finite RH LQ formulation by choosing the proper initialization, P_0.

5.2 Monotonicity via a Riccati approach

Although the idea of using monotonicity follows easily from our Lyapunov function approach, it is also possible to arrive at this through an appropriate reformulation of the RDE (6). If one defines, $\bar{Q}_j = Q - (P_{j+1} - P_j)$, then the RDE can be rewritten as:

$$P_j = A^T[P_j - P_j B(B^T P_j B + R)^{-1} B^T P_j]A + \bar{Q}_j. \tag{24}$$

This equation is called the Fake Algebraic Riccati Equation (FARE) because it is an ARE (4), except in terms of P_j and $\bar{Q}_j$. Writing the RDE in the form of the FARE allows one to use the following standard result concerning the stability of solutions to the ARE [3, 4, 15].

Theorem 5.3 *Consider the ARE (4) where $[A, B]$ is stabilizable, $[Q, A]$ has no unobservable modes on the unit circle, $Q \geq 0$ and $R \geq 0$. Then there exists a unique maximal, nonnegative definite symmetric solution P, which is stabilizing.*

Stability of the finite RH control $\hat{u}_N$ is equivalent to P_{N-1} being stabilizing, where P_{N-1} is determined through iterations of the RDE (6). Hence, if $[\bar{Q}_{N-1}, A]$ has no unobservable modes on the unit circle and $\bar{Q}_{N-1} \geq 0$, then P_{N-1} is stabilizing. If $\bar{Q}_{N-1} > 0$, then this ensures that these conditions are met. Hence, if

$$\bar{Q}_{N-1} = Q - (P_N - P_{N-1}) > 0 \tag{25}$$

then P_{N-1}, or equivalently $\hat{u}_N$, is stabilizing. Additionally, one can see that if P_j is monotonically *non–increasing*, then $\bar{Q}_j$ will always be positive definite, and stability is guaranteed for *any* choice of the finite horizon N.

5.2.1 A connection with finite RH LQ control

The FARE based stability result, (25) is actually a special case of the results presented in Section 4.4, and is similar to the result given in Theorem 4.4.

As in Section 4.4 consider using J_{N-1} as a Lyapunov function:

$$\begin{aligned} J_{N-1}(x(k)) - J_{N-1}(x(k+1)) &= J_{N-1}(x(k)) - [J_N(x(k)) - (x^T(k)Qx(k) + \hat{u}_N^T(k)R\hat{u}_N(k)] \\ &= [x^T(k)Qx(k) - (J_N(x(k)) - J_{N-1}(x(k)))] + \hat{u}_N^T(k)R\hat{u}_N(k) \\ &= x^T(k)\,[Q - (P_N - P_{N-1})]\,x(k) + \hat{u}_N^T(k)R\hat{u}_N(k) \end{aligned}$$

The FARE merely states that if the first term is positive ($Q - (P_N - P_{N-1}) > 0$), then J_{N-1} is a Lyapunov function and P_{N-1}, or equivalently $\hat{u}_N$, is stabilizing.

To see the connection with Theorem 4.4, note that:

$$\begin{aligned} x^T(k)\left[Q-(P_N-P_{N-1})\right]x(k) &= x^T(k)\left[QP_{N-1}^{-1}-(P_NP_{N-1}^{-1}-I)\right]P_{N-1}x(k) \\ &\geq x^T(k)[\rho_{N-1}-(\alpha_{N-1}-1)]P_{N-1}x(k) \end{aligned}$$

From this consideration, it is apparent that Theorem 4.4 is slightly more conservative than the FARE result due to the fact that α_{N-1} and ρ_{N-1} are independent bounds on individual terms involved in the above equation. The reason for calculating the quantities α_{N-1} and ρ_{N-1} separately is that α_{N-1} is a quantity needed for the performance bound $\tilde{\mathcal{P}}_N$ (eqn. (19)) and hence can serve a dual purpose as a parameter for both stability and performance.

Also note that the term $u^T(k)Ru(k)$ in the FARE result, as well as in Theorem 4.4, is ignored. On the other hand, Theorem 4.6 takes full advantage of the known structure of control $\hat{u}_N(k)$. It is therefore clear that the FARE result and Theorem 4.4 are more conservative than the result given in Theorem 4.6.

5.2.2 Bounds on performance through the FARE

Let us consider the following costs associated with the FARE:

1. $J(Q)$ is the optimal value of the infinite horizon control problem using the cost parameter Q (In our notation, this is simply the optimal infinite horizon cost J).

2. $J(\bar{Q}_{N-1})$ is the optimal cost of the infinite horizon control problem corresponding the the cost parameter $\bar{Q}_{N-1}$ (Note that from the FARE, this is exactly the finite horizon cost J_{N-1}).

3. $J_{\bar{Q}_{N-1}}(Q)$ is the value of the infinite horizon cost for the problem corresponding to the cost parameter Q, but using the controller designed for an infinite horizon problem with cost parameter $\bar{Q}_{N-1}$ (This is in fact the infinite horizon cost of the RH policy with horizon length N, $J_{\hat{u}_N}$).

Then it is possible to show the following result [3], [11].

Theorem 5.4 *$\bar{Q}_{N-1} \geq Q$ implies $J(\bar{Q}_{N-1}) \geq J_{\bar{Q}_{N-1}}(Q) \geq J(Q)$.*

Using the notation established in this paper, the above theorem states that under the condition $\bar{Q}_{N-1} \geq Q$, then $J \leq J_{\hat{u}_N} \leq J_{N-1}$. Furthermore, by noting that the condition $\bar{Q}_{N-1} \geq Q$ is equivalent to $P_N \geq P_{N-1}$ or $\alpha_{N-1} \leq 1$, it is revealed that Thm. 5.4 is nothing more than a special case of Thm. 4.5. When $\alpha_{N-1} \leq 1$, then by Thm. 4.5, $\tilde{\mathcal{P}}_N = 1$ and it reduces to the above result.

Recall that a similar result is given in Section 4.5, equation (21). In fact, due to $\alpha_{N-1} \leq 1$, $J_N \leq J_{N-1}$ and the bound given in (21) ($J \leq J_{\hat{u}_N} \leq J_N$) is *always* a better bound than that given in Theorem 5.4 (equivalently Thm. 4.5).

It should be noted that if one begins an iteration of the RDE with $P_0 = Q$, then the proper monotonicity will not be present, and Theorem 5.4 will *never* apply. For the above result to be of use, some other initialization for the RDE must be chosen which will result in the proper monotonicity. On the other hand, Theorems 4.2 and 4.5 apply for any initialization P_0 of the RDE.

5.2.3 (Non)usefulness of monotonicity based results

Using monotonicity to derive stability results and performance bounds turns out to be quite limited in its application to many well known and commonly used receding horizon related schemes. For example, Generalized Predictive Control (GPC) can be cast in a receding horizon LQ framework, but only with $P_0 = Q$. This *always* implies a monotonically non–decreasing sequence of P_N, which is noted for its difficulty in [3] by Bitmead et al.:

> *"In fact, with monotonicity going the wrong way, it is hard to say anything about the achieved performance of the RH GPC controller as measured against its optimal infinite horizon cost criterion except as N goes to infinity."*

Our finite RH LQ control formulation provides the framework to analyze situations, such as that described above, that are not within the reach of previous results. Furthermore, when $P_0 = Q$, we are able to determine the amount by which the RH policy can exceed that of the optimal infinite horizon policy, through knowledge of the finite horizon cost $J_N(x(0))$ and $\mathcal{P}_N$ only (recall equation (22)).

A unified picture of stability results, including the results discussed in this section, is given in Figure 1.

5.3 The optimal terminal weight

Finally, we remark that the best strategy for choosing a terminal weight would be to approximate the infinite horizon cost $J(x)$. For if one were to use the infinite horizon cost itself as a terminal weight, then a receding horizon policy of *any* horizon length would always recover the optimal infinite horizon controller. To see why this is, consider the finite horizon cost with terminal weight $J(x)$:

$$\begin{aligned} J_N(x_0) &= \inf_{u(\cdot)} \left[J(x(N)) + \sum_{k=0}^{N-1} \left(x^T(k)Qx(k) + u^T(k)Ru(k) \right) \right] \\ &= J(x_0) \end{aligned}$$

Hence the finite horizon problem is actually the infinite horizon problem.

Stability considerations have required that only special choices of terminal weights (end constraints and the infinite horizon approach) be used, which are usually far from the optimal choice of $J(x)$. From this viewpoint, the importance of the ability to analyze the stability of RH policies with arbitrary terminal weights is clear in that it may allow for a serious attempt at choosing a more "optimal" P_0.

6 Examples

Calculations demonstrating the characteristics of the finite RH LQ approach will be applied to two examples taken from [14] and [13]. In order to show the full range of flexibility of this approach, for each example we will consider the following initializations of the RDE:

(i) $P_0 = Q$ (resulting in monotonically non–decreasing P_N ($\alpha_N \geq 1$)).

(ii) $P_0 = P_0^\infty$ ($p = \infty$, resulting in monotonically non–increasing P_N ($\alpha_N \leq 1$), (see eqn. 23)).

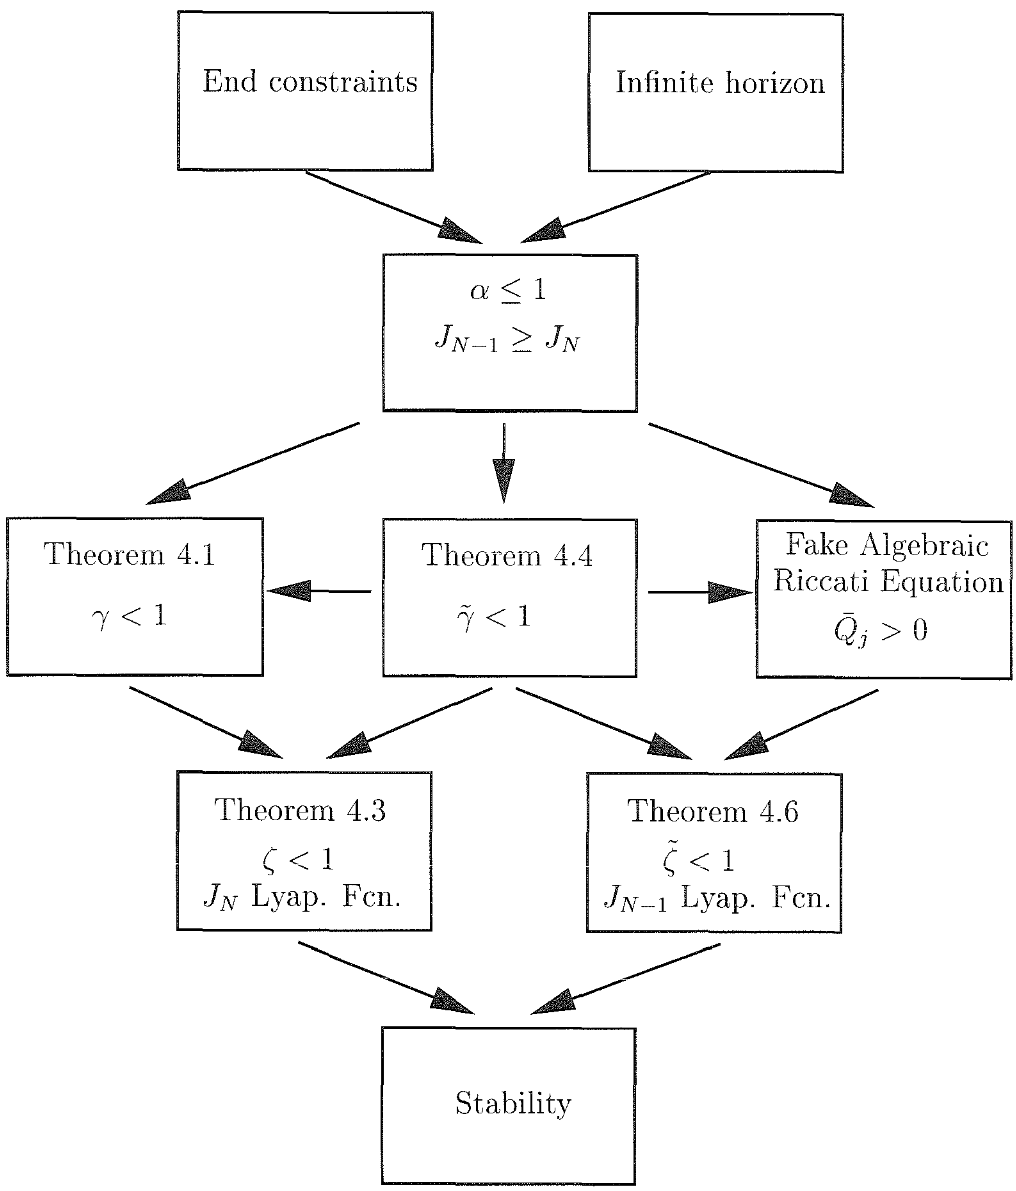

Figure 1: Diagram illustrating connections between stability results. (arrows signify "implies")

(iii) $P_0^{-1} = 0$ ($x(k+N) = 0$, resulting in monotonically non–increasing P_N ($\alpha_N \leq 1$)).

While the infinite horizon approach (ii), and end constraints (iii) fall within previously established monotonicity based results, only the methods presented in this paper are applicable to (i).

Stability will be determined by checking the conditions presented in Theorem 4.1 ($\gamma_N < 1$) and Theorem 4.3 ($\zeta_N < 1$) which correspond to J_N as a Lyapunov function, as well as those in Theorem 4.4 ($\tilde{\gamma}_N < 1$) and Theorem 4.6 ($\tilde{\zeta}_N < 1$) which correspond to J_{N-1}. Due to the fact that a monotonically non–increasing P_N guarantees stability for any N, we do not tabulate these parameters for (ii) and (iii).

Performance results are given in two different forms:

- **Performance w.r.t. optimal:** The upper bound for the infinite horizon performance of the RHC ($\mathcal{P}_N J_N$, cf. Theorem 4.2) is compared to the optimal infinite horizon cost (J), which is obtained by solving the ARE. The tabulated value is the ratio $\mathcal{P}_N J_N / J$.

- **Finite horizon performance bound**: When P_N is monotonically non–decreasing, as in (i), then $\mathcal{P}_N$ bounds the infinite horizon performance of the RH policy (cf. eqn. (22)). Note that this bound is not applicable to (ii) and (iii).

6.1 Example 1

Consider the unstable system

$$x(k+1) = \begin{pmatrix} 1 & 1.1 \\ -1.1 & 1 \end{pmatrix} x(k) + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u(k),$$

and the finite receding horizon control (8) with $Q = I$, $R = 1$.

Since both modes of this system are unstable, the infinite horizon approach (ii) is equivalent to imposing end constraints (iii) (cf. Section 5.1.2).

For comparison, in addition to $P_0 = Q$ and $P_0^{-1} = 0$, we provide the following initialization,

$$P_0^* = \begin{pmatrix} 397.7522 & 423.4626 \\ 423.4626 & 702.2478 \end{pmatrix} \tag{26}$$

which results in monotonically non–increasing P_N. Note that while this initialization is not equivalent to either end constraints or an infinite horizon approach, it is covered by FARE results, which require only monotonically non–increasing P_N to provide stability guarantees and performance bounds. Additionally, it should be mentioned that it was not trivial to find the P_0^* given above so that the proper monotonicity was present.

Table 2 presents stability results corresponding to Theorems 4.1 (γ), 4.4 ($\tilde{\gamma}$), 4.3 (ζ) and 4.6 ($\tilde{\zeta}$). Recall that γ_N and ζ_N correspond to using J_N, and $\tilde{\gamma}_N$ and $\tilde{\zeta}_N$ correspond to J_{N-1}, where in all cases, stability is guaranteed when the parameter is less than 1. For comparison, the first column lists the magnitude of the largest eigenvalue of the exact closed loop system. Using this "exact" stability test, a horizon of $N = 2$ stabilizes the system. The parameter ζ indicates that stability is guaranteed for a horizon length $N = 3$, while all the others require a more conservative estimate of $N = 4$. Note the existence of a "gap" between the horizon at which stability is achieved ($N = 2$) and when J_N (ζ column) becomes a Lyapunov function for the system ($N = 3$).

Table 3 presents guaranteed performance with respect to the optimal infinite horizon cost, where $\mathcal{P}_N^{\gamma}$ indicates that γ was used in equation (16), and similarly $\mathcal{P}_N^{\zeta}$ indicates the use of ζ. For reference, when $P_0 = Q$, the first column tabulates the ratio of the exact infinite horizon cost of

Table 2

N	Stability: $P_0 = Q$				
N	$\overline{\lambda}(A+BK_N)$	γ	$\tilde{\gamma}$	ζ	$\tilde{\zeta}$
1	1.0512	2.2100	2.2100	1.8418	2.2100
2	0.7718	1.7549	1.7566	1.2217	1.7230
3	0.6014	1.1805	1.1816	0.9436	1.1340
4	0.5348	0.9596	0.9597	0.8981	0.9367
5	0.5189	0.9147	0.9147	0.8857	0.8974
6	0.5146	0.9024	0.9024	0.8813	0.8854
7	0.5132	0.8980	0.8980	0.8802	0.8812
8	0.5128	0.8968	0.8968	0.8800	0.8802

Table 2: Stability, Example 1

Table 3

	Performance w.r.t. Optimal				
N	$P_0 = Q$			$P_0 = P_0^*$	$P_0^{-1} = 0$
	$\frac{J_{\hat{u}_N}}{J}$	$\frac{\mathcal{P}_N^{\gamma} J_N}{J}$	$\frac{\mathcal{P}_N^{\zeta} J_N}{J}$	$\frac{\mathcal{P}_N J_N}{J}$	$\frac{\mathcal{P}_N J_N}{J}$
1	—	—	—	33.9195	—
2	1.3346	—	—	1.6358	1.6730
3	1.0208	—	5.0411	1.0637	1.0660
4	1.0012	2.6224	1.5982	1.0070	1.0072
5	1.0001	1.2225	1.1603	1.0057	1.0058
6	1.0000	1.0623	1.0499	1.0018	1.0019
7	1.0000	1.0144	1.0120	1.0002	1.0002
8	1.0000	1.0031	1.0026	1.0000	1.0000

Table 3: Example 1: Ratio of exact RH performance ($J_{\hat{u}_N}$ (first column)) or performance bound of RH policy ($\mathcal{P}_N J_N$ (Thm. 4.2)) to optimal infinite horizon cost (J), for tabulated terminal weights P_0.

the receding horizon controller[1], $J_{\hat{u}_N}$, to that of the optimal infinite horizon cost J. The exact performance ($J_{\hat{u}_N}/J$) using $P_0 = P_0^*$ and $P_0^{-1} = 0$ is not significantly less than the upper bound ($\mathcal{P}_N J_N/J$) and hence is omitted. Also, recall that when non-increasing monotonicity is present (P_0^* and $P_0^{-1} = 0$), the upper bound using ζ and γ coincide (i.e. $J_N/J = \mathcal{P}_N^{\gamma} J_N/J = \mathcal{P}_N^{\zeta} J_N/J$ since $\mathcal{P}_N = 1$), hence no superscript is necessary to differentiate between these.

Both end constraints ($P_0^{-1} = 0$) and $P_0 = P_0^*$ provide better bounds than those using the initial condition $P_0 = Q$, guaranteeing a performance of less than 1% over that of the optimal performance with a horizon of $N = 4$. Using $P_0 = Q$ requires a longer horizon ($N = 8$) to guarantee a similar level of performance, even though as indicated in the first column ($\frac{J_{\hat{u}_N}}{J}$), this controller achieves this level of performance at horizon length $N = 4$. Hence, the performance bounds using $P_0 = Q$ tend to be conservative for short horizon lengths when compared to the "exact" performance bound.

[1]This cost can be calculated from the Lyapunov equation

$$(A + BK_N)^T P(A + BK_N) - P + (Q + K_N^T R K_N) = 0$$

As a final note, we mention that bounds based on using J_{N-1} (i.e. involving $\tilde{\gamma}$ or $\tilde{\zeta}$) are more conservative than those presented in Table 3, and have been omitted due to space considerations.

Table 4

	Finite horizon performance bounds $P_0 = Q$			
N	$\mathcal{P}_N^{\gamma}$	$\tilde{\mathcal{P}}_N^{\tilde{\gamma}}$	$\mathcal{P}_N^{\zeta}$	$\tilde{\mathcal{P}}_N^{\zeta}$
1	—	—	—	—
2	—	—	—	—
3	—	—	5.1760	—
4	2.6393	2.8404	1.6085	2.1707
5	1.2244	1.2505	1.1621	1.2081
6	1.0627	1.0700	1.0504	1.0596
7	1.0146	1.0162	1.0122	1.0139
8	1.0031	1.0034	1.0026	1.0030

Table 4: Factor $\mathcal{P}_N$ establishing bound on infinite horizon performance finite RH controller: $J_N(x(0)) \leq J(x(0)) \leq J_{\hat{u}_N}(x(0)) \leq \mathcal{P}_N J_N(x(0))$ (cf. eqn. (22))

Finally, in Table 4 we calculate the finite horizon performance bounds, given by $\mathcal{P}_N$ [2], and corresponding to $P_0 = Q$ which do not rely on our ability to calculate the optimal infinite horizon cost. Note that these bounds are not available for $P_0 = P_0^*$ and $P_0^{-1} = 0$, where P_N is *not* monotonically non–decreasing. Comparing Table 3 with Table 4, we see that the finite horizon bounds given in Table 4 are quite close to those in Table 2, demonstrating the potential usefulness of this approach.

6.2 Example 2

Consider the following stable dynamics taken from [13].

$$x(k+1) = \begin{pmatrix} 4/3 & -2/3 \\ 1 & 0 \end{pmatrix} x(k) + \begin{pmatrix} 1 \\ 0 \end{pmatrix} u(k) \tag{27}$$

with Q and R chosen as:

$$Q = Q^* = \begin{pmatrix} 4/9 & -2/3 \\ -2/3 & 1 \end{pmatrix}, \qquad R = 1. \tag{28}$$

Due to the semi-definiteness of Q, Theorem 4.1 (γ) and Theorem 4.4 ($\tilde{\gamma}$) are not applicable. Since this system is stable, the infinite horizon approach (ii) and end constraints (iii) result in distinct initializations P_0. The three compared in this example are:

(i) $P_0 = Q$

(ii) $P_0^{\infty} = \begin{pmatrix} 1.8889 & -1.4074 \\ -1.4074 & 1.8395 \end{pmatrix}$, (infinite horizon approach)

(iii) $P_0^{-1} = 0$, (end constraints).

Table 5

	Stability: $P_0 = Q^*$		
N	$\overline{\lambda}(A+BK_N)$	ζ	$\tilde{\zeta}$
1	$\underline{0.8255}$	2.8422	—
2	0.6774	1.6842	2.9207
3	0.6263	1.1995	1.6282
4	0.5720	1.0245	1.1513
5	0.5478	$\underline{0.9890}$	1.0146
6	0.5417	0.9851	$\underline{0.9877}$
7	0.5407	0.9852	0.9851
8	0.5407	0.9853	0.9853

Table 5: Stability: $Q = Q^*$, Example 2

Table 5 presents stability results for $P_0 = Q$. Recall that P_0^∞ and $P_0^{-1} = 0$ are guaranteed to result in stability. Once again the eigenvalues of the closed loop are provided in the first column for reference as an "exact" stability test. Observe that using ζ for $P_0 = Q$ is quite conservative and requires a horizon of $N = 5$ to guarantee stability, while by considering eigenvalues of the closed loop, the system is stable at $N = 1$. This shows that it is possible to have a substantial "stability gap" between when a system is stable, and when J_N will act as a Lyapunov function.

Table 6

	Performance w.r.t. Optimal			
N	$P_0 = Q$		$P_0 = P_0^\infty$	$P_0^{-1} = 0$
	$\frac{J_{\hat{u}_N}}{J}$	$\frac{\mathcal{P}_N^\zeta J_N}{J}$	$\frac{\mathcal{P}_N^\zeta J_N}{J}$	$\frac{\mathcal{P}_N^\zeta J_N}{J}$
1	2.5043	—	1.2736	—
2	1.8928	—	1.1825	3.5904
3	1.1036	—	1.0784	1.5558
4	1.0064	—	1.0245	1.1365
5	1.0002	4.6659	1.0057	1.0284
6	1.0000	1.3999	1.0009	1.0042
7	1.0000	1.0354	1.0001	1.0004
8	1.0000	1.0144	1.0000	1.0002

Table 6: Example 2: Ratio of exact RH performance ($J_{\hat{u}_N}$ (first column)) or performance bound of RH policy ($\mathcal{P}_N J_N$ (Thm. 4.2)) to optimal infinite horizon cost (J), for $Q = Q^*$ (eqn. (28)) and tabulated terminal weights P_0.

Table 6 tabulates performance results. Considering the horizon length $N = 5$ for $P_0 = P_0^\infty$ and $P_0^{-1} = 0$, performance is guaranteed to be very close to the optimal. By horizon length $N = 8$, the performance bound for $P_0 = Q$ has improved to within 2% of the optimal.

For comparison, the above calculations were redone with positive definite $Q = I$. This allows for a comparison with results from Theorem 4.1 (γ) and Theorem 4.4 ($\tilde{\gamma}$). The initializations P_0 now become:

[2] $\tilde{\mathcal{P}}_N$ for $\tilde{\gamma}$ and $\tilde{\zeta}$

Table 7

	Stability: $P_0 = Q = I$				
N	$\overline{\lambda}(A + BK_N)$	γ	$\tilde{\gamma}$	ζ	$\tilde{\zeta}$
1	<u>0.5774</u>	2.0000	2.0000	<u>0.6862</u>	1.4805
2	0.4140	<u>0.7990</u>	<u>0.8048</u>	0.4678	<u>0.4946</u>
3	0.3990	0.7165	0.7177	0.4683	0.4697
4	0.3970	0.7109	0.7109	0.4641	0.4661
5	0.3962	0.7077	0.7077	0.4634	0.4636
6	0.3961	0.7073	0.7073	0.4634	0.4634
7	0.3961	0.7072	0.7072	0.4634	0.4634
8	0.3961	0.7072	0.7072	0.4634	0.4634

Table 7: Stability: $Q = I$, Example 2

Table 8

	Performance w.r.t. Optimal				
N	$P_0 = Q$			$P_0 = P_0^{\infty}$	$P_0^{-1} = 0$
	$\frac{J_{\hat{u}_N}}{J}$	$\frac{\mathcal{P}_N^{\gamma} J_N}{J}$	$\frac{\mathcal{P}_N^{\zeta} J_N}{J}$	$\frac{\mathcal{P}_N J_N}{J}$	$\frac{\mathcal{P}_N J_N}{J}$
1	1.1418	—	2.3567	1.3594	—
2	1.0007	1.4717	1.0987	1.0912	1.3129
3	1.0001	1.0384	1.0130	1.0076	1.0164
4	1.0000	1.0131	1.0045	1.0008	1.0034
5	1.0000	1.0015	1.0005	1.0003	1.0010
6	1.0000	1.0001	1.0000	1.0000	1.0001
7	1.0000	1.0000	1.0000	1.0000	1.0000
8	1.0000	1.0000	1.0000	1.0000	1.0000

Table 8: Example 2: Ratio of exact RH performance ($J_{\hat{u}_N}$ (first column)) or performance bound of RH policy ($\mathcal{P}_N J_N$ (Thm. 4.2)) to optimal infinite horizon cost (J), for $Q = I$ and tabulated terminal weights P_0.

(i) $P_0 = Q$

(ii) $P_0^{\infty} = \begin{pmatrix} 10.00 & -5.33 \\ -5.33 & 5.44 \end{pmatrix}$

(iii) $P_0^{-1} = 0$

The results are supplied in Table 7 and 8. In this case, the stability and performance results corresponding to $P_0 = Q = I$ (Table 8) are far more competitive with those obtained by using $P_0 = P_0^{\infty}$ and $P_0^{-1} = 0$, than they were for the semi-definite $Q = Q^*$ considered previously. Both γ and ζ yield bounds that quickly converge to the optimal infinite horizon performance.

For both choices of Q used in this example, the finite horizon performance bounds were very close to the performance with respect to the optimal (given in Table 6 and 8), as was the case in Example 1.

7 Conclusion

Performance bounds and conditions for guaranteed stability were presented for a general formulation of finite receding horizon control applied to unconstrained linear systems with quadratic costs. In this case, necessary parameters were efficiently calculated using the Riccati difference equation. Different from previously established stability results, the approach did not rely on end constraints, an infinite horizon, or monotonicity arguments, but rather was shown to be valid for any terminal weight and encompass the majority of previous results as special cases. Although not explored in this paper, these results offer a general framework that can be naturally extended to nonlinear and constrained systems, further demonstrating the power of this finite receding horizon approach.

8 Acknowledgment

The authors would like to thank Jeff Shamma for providing the motivation behind exploring this work, and Manfred Morari, whose initial suspicion and later enthusiasm was inspiring and helpful. Additionally, the authors would like to acknowledge John Doyle for his constructive criticisms. The second author would like to thank NSF for their support.

A Proofs of Theorems 4.4, 4.5 and 4.6

Consider the receding horizon policy (8) based on the finite horizon cost (5).

Theorem A.1 (4.4) *Let N be such that*

$$\tilde{\gamma}_N = \alpha_{N-1} - \rho_{N-1} < 1$$

then the receding horizon policy $\tilde{u}_N(\cdot)$ is stabilizing, and $J_{N-1}(\cdot)$ is a Lyapunov function for the closed-loop system with

$$J_{N-1}(x(k+1)) \leq \tilde{\gamma}_N J_{N-1}(x(k))$$

Proof: Let $x(k)$ and $u(k)$ be the state and control trajectory, respectively, resulting from the receding horizon policy

$$u(k) = \hat{u}_N(x(k)).$$

For any $k \in \mathcal{Z}^+$,

$$J_N(x(k)) = x^T(k)Qx(k) + u^T(k)Ru(k) + J_{N-1}(x(k+1)).$$

Therefore,

$$\begin{aligned}
J_{N-1}(x(k)) - J_{N-1}(x(k+1)) &= J_{N-1}(x(k)) - [J_N(x(k)) - (x^T(k)Qx(k) + u^T(k)Ru(k))] \\
&= x^T(k)Qx(k) + u^T(k)Ru(k) + J_{N-1}(x(k)) - J_N(x(k)) \\
&\geq x^T(k)Qx(k) + u^T(k)Ru(k) + J_{N-1}(x(k)) - \alpha_{N-1}J_{N-1}(x(k)) \\
&\geq x^T(k)Qx(k) + u^T(k)Ru(k) + (1 - \alpha_{N-1})J_{N-1}(x(k)) \\
&\geq (1 - (\alpha_{N-1} - \rho_{N-1}))J_{N-1}(x(k))
\end{aligned}$$

which completes the proof. ∎

Theorem A.2 (4.5) *Let N and $\tilde{\gamma}_N$ be as in Theorem 4.4. Using the receding horizon policy*

$$u(k) = \hat{u}_N(x(k)),$$

a bound for the infinite-horizon performance is given by:

$$J_{\hat{u}_N}(x(0)) \leq \tilde{\mathcal{P}}_N J_{N-1}(x(0))$$

where

$$\tilde{\mathcal{P}}_N = \left(1 + (\max\{0, \alpha_{N-1} - 1\})\frac{1}{1-\tilde{\gamma}_N}\right) \tag{29}$$

Proof: Bounding the cost term-by-term gives the following. First,

$$\begin{aligned}
x^T(0)Qx(0) + u^T(0)Ru(0) &= J_N(x(0)) - J_{N-1}(x(1)) \\
&= J_N(x(0)) - J_{N-1}(x(0)) + J_{N-1}(x(0)) - J_{N-1}(x(1)) \\
&\leq \alpha_{N-1}J_{N-1}(x(0)) - J_{N-1}(x(0)) + J_{N-1}(x(0)) - J_{N-1}(x(1)) \\
&\leq (\alpha_{N-1} - 1)J_{N-1}(x(0)) + J_{N-1}(x(0)) + J_{N-1}(x(1)).
\end{aligned}$$

Similarly,

$$x^T(1)Qx(1) + u^T(1)Ru(1) \leq (\alpha_{N-1} - 1)J_{N-1}(x(1)) - J_{N-1}(x(1)) + J_{N-1}(x(2)).$$

By a summation of corresponding bounds on $x^T(k)Qx(k) + u^T(k)Ru(k)$ we obtain the following:

$$\begin{aligned}
\sum_{k=0}^{\infty} x^T(k)Qx(k) + u^T(k)Ru(k) &\leq J_{N-1}(x(0)) + (\alpha_{N-1} - 1)\sum_{k=0}^{\infty} J_{N-1}(x(k)) \\
&\leq J_{N-1} + (\max\{0, \alpha_{N-1} - 1\})\sum_{k=0}^{\infty} \tilde{\gamma}_N^k J_{N-1}(x(0)) \\
&= \left(1 + (\max\{0, \alpha_{N-1} - 1\})\frac{1}{1+\tilde{\gamma}_N}\right) J_{N-1}(x(0)),
\end{aligned}$$

■

Theorem A.3 *Let N be such that $P_{N-1} > 0$ and*

$$\tilde{\zeta}_N = \overline{\lambda}\left(P_{N-1}^{-1/2}(A + BK_N)^T P_{N-1}(A + BK_N)P_{N-1}^{-1/2}\right) < 1 \tag{30}$$

where K_N corresponds to

$$\hat{u}_N(k) = -(B^T P_{N-1}B + R)^{-1}B^T P_{N-1}Ax(k) = K_N x(k)$$

The receding horizon policy $\hat{u}_N(\cdot)$ is stabilizing. Furthermore, $J_{N-1}(\cdot)$ is a Lyapunov function for the closed-loop system, so that

$$J_{N-1}(x(k+1)) \leq \tilde{\zeta}_N J_{N-1}(x(k)),$$

Proof: Let $x(k)$ be the state trajectory resulting from the receding horizon policy

$$\hat{u}_N(x(k)) = K_N x(k).$$

Therefore,

$$\begin{aligned}
J_{N-1}(x(k)) - J_{N-1}(x(k+1)) &= x^T(k)P_{N-1}x(k) - x^T(k+1)P_{N-1}x(k+1) \\
&= x^T(k)P_{N-1}x(k) - x^T(k)(A+BK_N)^T P_{N-1}(A+BK_N)x(k) \\
&\geq \underline{\lambda}\left(I - P_{N-1}^{-1/2}(A+BK_N)^T P_{N-1}(A+BK_N)P_{N-1}^{-1/2}\right) J_{N-1}(x(k)) \\
&= \left[1 - \overline{\lambda}\left(P_{N-1}^{-1/2}(A+BK_N)^T P_{N-1}(A+BK_N)P_{N-1}^{-1/2}\right)\right] J_{N-1}(x(k))
\end{aligned}$$

which completes the proof. ∎

B Proof that γ is less conservative than $\tilde{\gamma}$

We show that $\tilde{\gamma}$ (cf. Thm 4.4) always results in a *more conservative* stabilty estimate than γ (cf. Thm 4.1). Since stability is guaranteed when either γ or $\tilde{\gamma}$ is less than 1, if we can show that $\gamma \leq \tilde{\gamma}$ then this will give the desired result.

Theorem B.1 $\gamma \leq \tilde{\gamma}$.

Proof: Note that $\rho_{N-1} \leq \alpha_{N-1}\rho_N$. This is true because:

$$Q \geq \rho_{N-1}P_{N-1} \geq \frac{\rho_{N-1}}{\alpha_{N-1}}P_N$$

which implies that $\rho_N \geq \frac{\rho_{N-1}}{\alpha_{N-1}}$ or $\rho_{N-1} \leq \alpha_{N-1}\rho_N$.

So:

$$\begin{aligned}
\gamma_N &= \alpha_{N-1}(1-\rho_N) \\
&= \alpha_{N-1} - \alpha_{N-1}\rho_N \\
&\leq \alpha_{N-1} - \rho_{N-1} \\
&= \tilde{\gamma}_N
\end{aligned}$$

∎

The next theorem shows that the performance bound obtained by using γ (i.e. $J_{\hat{u}_N} \leq \mathcal{P}_N^{\gamma} J_N$) is always tighter than that obtained from $\tilde{\gamma}$ (i.e. $J_{\hat{u}_N} \leq \tilde{\mathcal{P}}_N^{\tilde{\gamma}} J_{N-1}$). This is equivalent to showing that $\mathcal{P}_N^{\gamma} J_N \leq \tilde{\mathcal{P}}_N^{\tilde{\gamma}} J_{N-1}$

Theorem B.2 $\mathcal{P}_N^{\gamma} J_N \leq \tilde{\mathcal{P}}_N^{\tilde{\gamma}} J_{N-1}$

Proof: If $\alpha_{N-1} \leq 1$ then it is obvious since $\mathcal{P}_N^{\gamma} = \tilde{\mathcal{P}}_N^{\tilde{\gamma}} = 1$ and $J_N \leq \alpha_{N-1}J_{N-1} \leq J_{N-1}$. So assume that $\alpha_{N-1} > 1$. Then by using the bound in Theorem 4.2, we get:

$$\mathcal{P}_N^{\gamma} J_N = \left[1 + \left(\frac{\alpha_{N-1}-1}{\alpha_{N-1}}\right)\left(\frac{\gamma_N}{1-\gamma_N}\right)\right] J_N$$

$$
\begin{aligned}
&= \left[1 + \left(\frac{\alpha_{N-1} - 1}{\alpha_{N-1}}\right)\left(\frac{\alpha_{N-1}(1-\rho_N)}{1-\alpha_{N-1}(1-\rho_N)}\right)\right] J_N \\
&= \left[\alpha_{N-1} + (\alpha_{N-1} - 1)\left(\frac{\alpha_{N-1}(1-\rho_N)}{1-\alpha_{N-1}(1-\rho_N)}\right)\right] \frac{J_N}{\alpha_{N-1}} \\
&\leq \left[1 + (\alpha_{N-1} - 1) + (\alpha_{N-1} - 1)\left(\frac{\alpha_{N-1}(1-\rho_N)}{1-\alpha_{N-1}(1-\rho_N)}\right)\right] J_{N-1} \\
&= \left[1 + (\alpha_{N-1} - 1)\left(1 + \frac{\alpha_{N-1}(1-\rho_N)}{1-\alpha_{N-1}(1-\rho_N)}\right)\right] J_{N-1} \\
&= \left[1 + (\alpha_{N-1} - 1)\left(\frac{1}{1-\alpha_{N-1}(1-\rho_N)}\right)\right] J_{N-1} \\
&\leq \left[1 + (\alpha_{N-1} - 1)\left(\frac{1}{1-(\alpha_{N-1}-\rho_{N-1})}\right)\right] J_{N-1} \\
&= \left[1 + (\alpha_{N-1} - 1)\left(\frac{1}{1-\tilde{\gamma}_N}\right)\right] J_{N-1} \\
&= \tilde{\mathcal{P}}_N^{\tilde{\gamma}} J_{N-1}
\end{aligned}
$$

∎

Careful analysis of the above proof also reveals that $\mathcal{P}_N^{\gamma} \leq \tilde{\mathcal{P}}_N^{\tilde{\gamma}}$.

References

[1] D.P. Bertsekas, "Dynamic Programming and Stochastic Control", In *Academic Press* New York, 1976.

[2] R.R. Bitmead, M. Gevers, I.R. Petersen, and R.J. Kaye, "Monotonicity and stabilizability properties of solutions of the Riccati difference equation: propositions, lemmas, theorems, fallacious conjectures and counterexamples", In *Systems and Control Letters*, 5:309–315, 1985.

[3] R.R. Bitmead, M. Gevers, and V. Wertz, "Optimal control redesign of generalized predictive control", In *Proc. IFAC Symp. on Adaptive Systems in Control and Signal Proc.*, 129–134, 1989

[4] R.R. Bitmead, M. Gevers, and V. Wertz, *Adaptive Optimal Control*, Prentice Hall, Englewood Cliffs, N.J., 1990.

[5] C.R. Cutler and B.L. Ramaker, "Dynamic matrix control–a computer algorithm", In *AIChE National Meeting, Houston, TX*, Apr. 1979.

[6] C.E. García, D.M. Prett, and M. Morari, "Model Predictive Control: Theory and Practice – A Survey", In *Automatica*, Vol 25., No. 3, May 1989.

[7] 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", In *Journal of Opt. Theory and Appl.*, Vol 57., No. 2, May 1988.

[8] H. Kwakernaak and R. Sivan, In *Linear Optimal Control Systems*, Wiley-Interscience, New York, 1972.

[9] W.H. Kwon and A.E. Pearson, "A modified quadratic cost problem and feedback stabilization of a linear system", In *IEEE Trans. Auto. Control*, AC–22:838–842, 1977.

[10] W.H. Kwon and A.E. Pearson, "On feedback stabilization of time-varying discrete linear systems", In *IEEE Trans. Auto. Control*, AC–23:479–481, 1978.

[11] T. Nishimura, "On the a priori information in sequential estimation problems", In *IEEE Trans. Automatic Control*, AC-12: 123–125, 1967.

[12] J. Richalet, A. Rault, J.L. Testud, and J. Papon, "Model predictive heuristic control: Applications to industrial processes", In *Automatica*, vol. 14, pp. 413–428, 1978.

[13] J.B. Rawlings and K.R. Muske, "The stability of constrained receding horizon control", In *IEEE Trans. Auto. Control*, Vol 38, No. 10:1512–1516, 1993.

[14] J.S. Shamma and D. Xiong, "Linear Non-Quadratic Optimal Control", To appear in *IEEE Trans. Auto. Control*, 1997.

[15] C.E. de Souza, M. Gevers, and G.C. Goodwin, "Riccati equations in optimal filtering of nonstabilizable systems having singular state transition matrices", In *IEEE Trans. Automatic Control*, AC-31: 831–838, 1986.