Worst-Case Finite-Horizon Analysis of Perturbed LQR via Performance Estimation Problem

Anthony Couthures, Gianluca Bianchin, and Julien M. Hendrickx

***Abstract*— We propose an exact formulation to quantify the finite-horizon robustness of a Linear Quadratic Regulator (LQR) subject to norm-bounded Linear Time-Invariant (LTI) parametric uncertainties, using the Performance Estimation Problem (PEP) framework. Originally developed for optimization, this method allows us to reformulate the worst-case performance search as a rank-constrained Semidefinite Program (SDP). By applying operator interpolation, this formulation retains the sequence variables, allowing the worst-case to be explicitly reconstructed when a rank-one solution is found. This approach also opens new possibilities for the finite-horizon certification of broader system classes, such as data-driven models or nonlinear components. Using this framework, we observe that finite-horizon LTI perturbations do not simply follow asymptotic strategies. Instead, they align slowly decaying invariant manifolds with specific initial conditions to drive transient growth, causing a bifurcation in the worst-case behavior as the horizon length changes. Finally, we show numerically that our performance bounds sit between the upper bounds of Integral Quadratic Constraint (IQC) relaxations and a lower bound derived from a pseudospectral heuristic.**

I. Introduction

In the current control literature, there is a gap between the theoretical design of optimal controllers, such as a Linear Quadratic Regulator (LQR), and their real-world applications. These control laws are typically designed assuming a perfectly known, nominal model over an infinite horizon, whereas practical applications often involve deploying them for finite durations on systems subject to parametric uncertainties. Such parameter variations, which create a mismatch between the design model and the true physical system, can arise, for example, from changing operating conditions or unmodeled nonlinearities like a varying resistance in an electronic circuit.

To handle these parametric uncertainties, it is natural to turn to robust control to provide guarantees on the desired behaviors [1]–[3]. Robust control approaches include, for example, μ-analysis [4] (which evaluates worst-case stability and performance against structured uncertainties), robust Lyapunov functions and Integral Quadratic Constraints (IQC) [5] (a framework that bounds the input-output behavior of complex uncertainties and nonlinearities using quadratic inequalities). These methods provide theoretical certificates under the presence of such perturbations.

However, while powerful, most robust control certificates evaluate only asymptotic behavior or specific perturbations, limiting their use in finite-horizon applications. For instance, μ-analysis computes exact robust stability margins for Linear Time-Invariant (LTI) systems [4], but it evaluates asymptotic stability (see Example 1). Similarly, while recent structured pseudospectra algorithms provide robust transient bounds [6], [7], they rely on continuous-time resolvent envelopes rather than finite-horizon trajectories. Time-domain IQC methods handle finite horizons using dissipation inequalities [8], but rely on dual bounding techniques, which create two limitations. First, standard time-domain IQC formulations relax the static LTI parameter variations into a larger class of Linear Time-Varying (LTV) uncertainties, introducing conservatism. Second, the dual approach acts primarily as a bounding oracle. Because the optimization variables are dual multipliers rather than system trajectories, extracting the exact physical worst-case perturbation matrices or initial conditions remains difficult.

By contrast, in this paper we evaluate the finite-horizon performance of an LQR controller under norm-bounded LTI uncertainties using the Performance Estimation Problem (PEP) framework [9], [10] and exhibit worst-case behaviors. Originally developed to analyze optimization algorithms, PEP treats worst-case analysis as an optimization problem over discrete trajectories. Using necessary and sufficient operator interpolation conditions, we formulate the search for worst-case LTI trajectories as Quadratically Constrained Quadratic Program (QCQP) and then as a rank-constrained Semidefinite Program (SDP). This primal approach avoids LTV conservatism. By retaining the sequence variables, we can extract the exact worst-case initial conditions and perturbation matrices. Instead of just providing scalar bounds, PEP acts as a diagnostic tool to uncover geometric failure modes when these controllers are deployed over a finite horizon. Finally, we compare our rank-relaxed SDP and numerically extracted lower bounds against standard IQC results to quantify the conservatism gap introduced by the LTV relaxation.

Notation. We denote by $\mathbb{R}$ the set of real numbers. Vectors of dimension n will be denoted as $x \in \mathbb{R}^n$ in lowercase and a matrix $A \in \mathbb{R}^{n\times m}$ in uppercase letters.

A. Couthures, G. Bianchin and J. M. Hendrickx are with the ICTEAM institute, UCLouvain, Louvain-la-Neuve, Belgium. A. Couthures, G. Bianchin and J. M. Hendrickx are supported by the "SIDDARTA" Concerted Research Action (ARC) granted by the UCLouvain. J. M. Hendrickx is also supported by "InterpoControl" F.R.S.-FNRS research project. anthony.couthures@uclouvain.be

The concatenation of a sequence of vectors $\{x_k\}_{k=0}^{N-1}$ is denoted by $\boldsymbol{x} = [x_0, \dots, x_{N-1}] \in \mathbb{R}^{n\times N}$. For two matrices $A, B \in \mathbb{R}^{n\times n}$, we denote by $A \succeq B$ the Loewner partial order, i.e., $A - B$ is positive semidefinite.

II. Problem Formulation

Consider a discrete-time linear system $x_{k+1} = Ax_k + Bu_k$ with state trajectory $\{x_k\}_{k=0}^{N} \subset \mathbb{R}^n$, control sequence $\{u_k\}_{k=0}^{N-1} \subset \mathbb{R}^m$, for the nominal matrices $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$. Let $u_k = Kx_k$ be a nominal state-feedback control law, where $K = -(R+B^\top PB)^{-1}B^\top PA$ is the standard infinite-horizon LQR gain derived from state and input penalty matrices $Q \succeq 0$ and $R \succ 0$, and $P \succ 0$ is the solution to the discrete algebraic Riccati equation: $P = A^\top PA - A^\top PB(R+B^\top PB)^{-1}B^\top PA + Q$.

As discussed in Section I, deploying this nominal controller in practice subjects the system to parametric mismatches. We therefore evaluate the performance of the fixed controller K under unknown but bounded LTI perturbations on the system matrices:

$$x_{k+1} = (A + \Delta_A)x_k + (B + \Delta_B)u_k \tag{1}$$

where $\Delta_A \in \mathbb{R}^{n\times n}$ and $\Delta_B \in \mathbb{R}^{n\times m}$ are static perturbations bounded in spectral norm by $\varepsilon_A, \varepsilon_B \geq 0$:

$$\Delta_A \in \mathcal{D}_A := \{M \in \mathbb{R}^{n\times n} \mid \|M\|_2 \leq \varepsilon_A\},$$
$$\Delta_B \in \mathcal{D}_B := \{M \in \mathbb{R}^{n\times m} \mid \|M\|_2 \leq \varepsilon_B\}.$$

Under the nominal control law $u_k = Kx_k$, the closed-loop perturbed dynamics becomes:

$$x_{k+1} = (\Phi + \Delta_A + \Delta_B K)x_k, \tag{2}$$

where $\Phi = A + BK$ is the nominal closed-loop state matrix.

To quantify the transient performance degradation caused by these uncertainties, we consider a finite-horizon evaluation cost over N steps, with a terminal penalty $S \succeq 0$:

$$J_N(x_0, \Delta) = \sum_{k=0}^{N-1} \left(x_k^\top Q x_k + u_k^\top R u_k\right) + x_N^\top S x_N,$$

where the state trajectory depends on the perturbations $\Delta = (\Delta_A, \Delta_B)$ through the closed-loop dynamics defined in (2). Substituting $u_k = Kx_k$, this cost simplifies to a quadratic function of the state trajectory:

$$J_N(x_0, \Delta) = \sum_{k=0}^{N-1} x_k^\top Q_K x_k + x_N^\top S x_N, \tag{3}$$

where $Q_K = Q + K^\top RK \succeq 0$ is the equivalent closed-loop stage cost matrix, and the state sequence $\{x_k\}_{k=0}^{N}$ is generated by the perturbed dynamics (2). As an infinite-horizon gain, K remains suboptimal for this finite-horizon cost even if $\Delta = 0$.

Given an initial state bounded by the unit ball, the worst-case performance is:

$$\begin{aligned} J_N^*(\varepsilon_A, \varepsilon_B, Q_K, S) := \max_{x_0, \Delta_A, \Delta_B} \quad & J_N(x_0, \Delta) \\ \text{s.t.} \quad & \|x_0\|_2^2 \leq 1 \\ & x_{k+1} = (\Phi + \Delta_A + \Delta_B K)x_k, \quad \forall k \in \mathcal{K} \\ & \Delta_A \in \mathcal{D}_A, \quad \Delta_B \in \mathcal{D}_B \end{aligned} \tag{4}$$

where $\mathcal{K} = \{0, \dots, N-1\}$. We restrict x_0 to the unit ball without loss of generality. Because the dynamics is linear and the cost is quadratic, the worst-case cost under an arbitrary bound $\|x_0\|_2^2 \leq \gamma$ equals γJ_N^*.

Example 1 (Worst-case with nominal pole): Consider

$$A = \begin{bmatrix} 0 & \frac{5}{4} \\ 0 & \frac{1}{4} \end{bmatrix}, \quad B = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}, \quad Q = \mathrm{diag}(1, \frac{3}{4}), \quad R = 1.$$

The Riccati solution is $P = \mathrm{diag}(1, 2)$, so that $K = [\,0 \;\; -1/2\,]$, $\Phi = \left[\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right]$ and $Q_K = \mathrm{Id}$. Take $N = 2$, $S = 0$, $\varepsilon_A = 3/8$ and $\varepsilon_B = 1/4$.

Writing $\Phi_\Delta = \Phi + \Delta_A + \Delta_B K$, the cost reduces to $J_2(x_0, \Delta) = \|x_0\|_2^2 + \|\Phi_\Delta x_0\|_2^2$, hence $J_2^* = 1 + \max_\Delta \|\Phi_\Delta\|_2^2$. The triangle inequality gives $\|\Phi_\Delta\|_2 \leq \|\Phi\|_2 + \varepsilon_A + \varepsilon_B\|K\|_2 = 3/2$, and this bound is attained by

$$\Delta_A^* = \begin{bmatrix} 0 & \frac{3}{8} \\ 0 & 0 \end{bmatrix}, \quad \Delta_B^* = \begin{bmatrix} -\frac{1}{4} \\ 0 \end{bmatrix}, \quad \Phi_{\Delta^*} = \begin{bmatrix} 0 & \frac{3}{2} \\ 0 & 0 \end{bmatrix}.$$

From $x_0 = [\,0 \;\; 1\,]^\top$, this gives $J_2^* = 1 + 9/4 = 3.25$, against 2 for the nominal loop. Yet Φ and Φ_{Δ^*} have the same eigenvalues $\{0, 0\}$: the worst-case perturbation leaves the poles in place and strengthens the coupling from the second state to the first. An analysis based on eigenvalue location, such as μ-analysis or pseudospectra, sees no degradation here. In our computations, the perturbations that maximize the spectral radius ($\rho \approx 0.79$) yield at most $J_2 \approx 3.05$. The IQC bound is tight in this case, since with $N = 2$ only one transition is penalized and a time-varying perturbation has no extra freedom.

IQC is exact in this short-horizon example, but not in general. Over longer horizons, a time-varying perturbation can steer the state in a different direction at each step, which a single matrix of the same norm cannot always do. The LTV relaxation then overestimates the LTI worst-case, as Section V shows numerically.

III. The Performance Estimation Problem (PEP) Reformulation

Our main result, Theorem 1, uses the PEP framework to recast the exact worst-case optimization problem (4) into a rank-constrained SDP. To derive this formulation, we abstract the problem over discrete sequences, enforce the LTI behavior of the perturbations through operator interpolation, and lift the sequence variables into a positive semidefinite Gram matrix.

A. The Abstract Optimization Problem

To bypass the coupling between the perturbation matrices Δ_A, Δ_B and the state x_k in (2), we treat the perturbation outputs as independent sequence variables. Over the horizon N, we define:

$$\begin{aligned}\boldsymbol{x} &= [x_0, \dots, x_{N-1}] \in \mathbb{R}^{n\times N},\\ \boldsymbol{v}_A &= [v_{A,0}, \dots, v_{A,N-1}] \in \mathbb{R}^{n\times N},\\ \boldsymbol{v}_B &= [v_{B,0}, \dots, v_{B,N-1}] \in \mathbb{R}^{n\times N},\end{aligned}$$

where at each time step $k \in \mathcal{K}$, the pointwise perturbations are given by:

$$v_{A,k} = \Delta_A x_k, \qquad v_{B,k} = \Delta_B K x_k.$$

By substituting these vectors into the system dynamics, the state update becomes a linear equation independent of the explicit matrices Δ_A and Δ_B:

$$x_{k+1} = \Phi x_k + v_{A,k} + v_{B,k}. \tag{5}$$

Instead of optimizing over the matrices Δ_A, Δ_B, our decision variables are now the sequences of vectors: the states $\{x_k\}_{k=0}^{N}$ and the perturbations $\{v_{A,k}\}_{k=0}^{N-1}$ and $\{v_{B,k}\}_{k=0}^{N-1}$.

The optimization problem (4) can thus be rewritten as the abstract optimization problem over sequences:

$$\begin{aligned}&J_N^*(\varepsilon_A, \varepsilon_B, Q_K, S) = \max_{x_k, v_{A,k}, v_{B,k}} \quad J_N(x_0, \boldsymbol{v}_A, \boldsymbol{v}_B)\\ \text{s.t.}\quad &\|x_0\|_2^2 \le 1 \\ &x_{k+1} = \Phi x_k + v_{A,k} + v_{B,k}, \quad \forall k \in \mathcal{K}\\ &\exists \Delta_A \in \mathcal{D}_A \text{ s.t. } v_{A,k} = \Delta_A x_k, \quad \forall k \in \mathcal{K}\\ &\exists \Delta_B \in \mathcal{D}_B \text{ s.t. } v_{B,k} = \Delta_B K x_k, \quad \forall k \in \mathcal{K}\end{aligned} \tag{6}$$

By shifting to this sequence-based formulation, the objective and the dynamics are now quadratic and linear, respectively. However, the final two constraints ensuring the perturbation sequences can be generated by single norm-bounded LTI operators cannot be directly encoded in a solver. The PEP framework works by translating these abstract existence constraints into algebraic inequalities.

B. Operator Interpolation: Enforcing LTI Behavior

Let $\mathcal{L}_{[0,L]}$ denote the set of matrices $M \in \mathbb{R}^{d_2\times d_1}$ with an induced 2-norm bounded by $L \ge 0$. We use operator interpolation to enforce LTI behavior on the perturbation variables. Let $\boldsymbol{p} = [p_0, \dots, p_{N-1}] \in \mathbb{R}^{d_1\times N}$ and $\boldsymbol{q} = [q_0, \dots, q_{N-1}] \in \mathbb{R}^{d_2\times N}$ be matrices containing vector sequences. We say the sequences are $\mathcal{L}_{[0,L]}$-interpolable if there exists $M \in \mathcal{L}_{[0,L]}$ such that $q_k = Mp_k$ for all k. For $\mathcal{L}_{[0,L]}$ the interpolation conditions are given by algebraic conditions over the data matrices from [11].

Lemma 1 ($\mathcal{L}_{[0,L]}$-Interpolation Condition): Let $\boldsymbol{p} \in \mathbb{R}^{d_1\times N}$ and $\boldsymbol{q} \in \mathbb{R}^{d_2\times N}$ be data matrices. The sequences are $\mathcal{L}_{[0,L]}$-interpolable if and only if the following LMI holds:

$$L^2 \boldsymbol{p}^\top \boldsymbol{p} - \boldsymbol{q}^\top \boldsymbol{q} \succeq 0.$$

Lemma 1 allows us to map the operator existence condition in (6) to quadratic constraints on the data matrix, which turn the problem into a QCQP.

Proposition 1 (QCQP formulation): The worst-case finite-horizon performance J_N^* defined in (6) equals the optimal value of the following QCQP over the finite vector sequences $\boldsymbol{x}, \boldsymbol{v}_A$, and $\boldsymbol{v}_B$:

$$\begin{aligned}&J_N^*(\varepsilon_A, \varepsilon_B, Q_K, S) = \max_{\boldsymbol{x}, \boldsymbol{v}_A, \boldsymbol{v}_B} \quad J_N(x_0, \boldsymbol{v}_A, \boldsymbol{v}_B)\\ \text{s.t.}\quad &\|x_0\|_2^2 \le 1 \\ &x_{k+1} = \Phi x_k + v_{A,k} + v_{B,k}, \quad \forall k \in \mathcal{K}\\ &\varepsilon_A^2 \boldsymbol{x}^\top \boldsymbol{x} - \boldsymbol{v}_A^\top \boldsymbol{v}_A \succeq 0\\ &\varepsilon_B^2 \boldsymbol{x}^\top K^\top K \boldsymbol{x} - \boldsymbol{v}_B^\top \boldsymbol{v}_B \succeq 0\end{aligned} \tag{7}$$

Proof: Using Lemma 1, we map the existence constraints $\exists \Delta_A \in \mathcal{D}_A$ in the optimization problem (6) to the LMI $\varepsilon_A^2 \boldsymbol{x}^\top \boldsymbol{x} - \boldsymbol{v}_A^\top \boldsymbol{v}_A \succeq 0$. Similarly, $\exists \Delta_B \in \mathcal{D}_B$ is replaced by $\varepsilon_B^2 \boldsymbol{x}^\top K^\top K \boldsymbol{x} - \boldsymbol{v}_B^\top \boldsymbol{v}_B \succeq 0$, yielding (7). ∎

Proposition 1 provides an equivalent formulation of (6) and thus to (4). However, the resulting QCQP remains non-convex due to the quadratic matrix inequalities.

C. The Gram Matrix SDP Formulation

To bypass the intractability of the QCQP, we use a Gram lifting approach [12]. Because the objective function and constraints in (7) depend solely on quadratic inner products of the sequence variables, we can linearize them by shifting our decision variables to a positive semidefinite Gram matrix.

We concatenate the initial state and the independent perturbation sequences into a single lifted vector $\boldsymbol{z}$:

$$\boldsymbol{z} = \begin{bmatrix} x_0^\top & v_{A,0}^\top \dots v_{A,N-1}^\top & v_{B,0}^\top \dots v_{B,N-1}^\top \end{bmatrix}^\top.$$

We define block-selector matrices $E_0, E_{A,k}, E_{B,k} \in \mathbb{R}^{n\times n(1+2N)}$ to extract the corresponding blocks from $\boldsymbol{z}$:

$$x_0 = E_0 \boldsymbol{z}, \quad v_{A,k} = E_{A,k} \boldsymbol{z}, \quad v_{B,k} = E_{B,k} \boldsymbol{z}.$$

Because the linear dynamics determines the state trajectory, every state x_k is a linear combination of the components of $\boldsymbol{z}$. Using the dynamics from Proposition 1, we define the state extraction matrices $H_{x,k}$ recursively:

$$H_{x,0} = E_0, \quad H_{x,k+1} = \Phi H_{x,k} + E_{A,k} + E_{B,k}, \quad \forall k \in \mathcal{K}.$$

Thus, $x_k = H_{x,k} \boldsymbol{z}$ for all $k \in \mathcal{K}$. Any scalar inner product in the system can be expressed as a linear trace operation on the $(1+2N)n \times (1+2N)n$ positive semidefinite Gram matrix $Z = \boldsymbol{z}\boldsymbol{z}^\top$, which has rank at most 1. Indeed, for any symmetric $M \in \mathbb{R}^{n\times n}$, one has:

$$x_i^\top M x_j = \boldsymbol{z}^\top H_{x,i}^\top M H_{x,j} \boldsymbol{z} = \operatorname{Tr}\left(H_{x,j}^\top M H_{x,i} Z\right).$$

By the trace property, the (i,j)-th entry of this matrix is $\operatorname{Tr}(H_{x,j}^\top H_{x,i} Z)$. Similarly, the cost J_N can be expressed as a linear operator acting on Z. Explicitly, the symmetric cost matrix $M_{J_N} \in \mathbb{R}^{(1+2N)n\times(1+2N)n}$ is obtained

from the state extraction matrices $H_{x,k}$ as:

$$M_{J_N} = \sum_{k=0}^{N-1} H_{x,k}^\top Q_K H_{x,k} + H_{x,N}^\top S H_{x,N},$$

which yields $J_N(x_0, \boldsymbol{v}_A, \boldsymbol{v}_B) = \mathrm{Tr}(M_{J_N} Z)$.

Let $\mathcal{G}_{\boldsymbol{x}}(Z) \in \mathbb{S}^{N\times N}$ denote the linear operator mapping Z to the $N\times N$ sequence inner-product matrix $\boldsymbol{x}^\top \boldsymbol{x}$. We define analogous linear operators $\mathcal{G}_{\boldsymbol{v}_A}(Z), \mathcal{G}_{\boldsymbol{v}_B}(Z)$, and $\mathcal{G}_{K\boldsymbol{x}}(Z)$ for the other sequences. This allows us to express the interpolation conditions as LMIs over Z.

Theorem 1 (Rank-constrained SDP formulation): The worst-case finite-horizon performance J_N^* equals the optimal value of the following rank-constrained SDP over the Gram matrix Z:

$$\begin{aligned} & J_N^*(\varepsilon_A, \varepsilon_B, Q_K, S) = \max_{Z\succeq 0} \quad \mathrm{Tr}(M_{J_N} Z) \\ \text{s.t.} \quad & \mathrm{Tr}\left(E_0^\top E_0 Z\right) \le 1 \\ & \varepsilon_A^2 \mathcal{G}_{\boldsymbol{x}}(Z) - \mathcal{G}_{\boldsymbol{v}_A}(Z) \succeq 0 \\ & \varepsilon_B^2 \mathcal{G}_{K\boldsymbol{x}}(Z) - \mathcal{G}_{\boldsymbol{v}_B}(Z) \succeq 0 \\ & \mathrm{rank}(Z) \le 1 \end{aligned} \tag{8}$$

This reformulates the problem into an SDP with a rank constraint. While still NP-hard, isolating the non-convexity allows us to drop the rank constraint to form a tractable convex relaxation. Moreover, the rank then serves as an admissibility check: if the relaxed solution has rank one, the bound is tight and produces an exact, physically realizable worst-case trajectory. However, the convex relaxation suffers from conservatism. Indeed, the optimal relaxed Gram matrix Z_{SDP} often has a rank strictly greater than one [12]. In this case, the relaxation only provides an upper bound on J_N^* that is not necessarily physical, as it may correspond to a superposition of trajectories.

IV. Extracting Worst-Case via PEP

As noted above, the solution to the relaxed problem (8), denoted Z_{SDP}, is often of rank strictly greater than one, already for relatively short planning horizons ($N \ge 5$). To find a valid LTI trajectory, we keep the SDP result as an upper bound and search for an admissible worst case directly over the perturbations. For fixed (Δ_A, Δ_B), the worst initial condition is the dominant eigenvector of the finite-horizon Gramian, so the search reduces to maximizing J_N over the perturbations alone. We do so by projected gradient ascent on the norm balls $\mathcal{D}_A \times \mathcal{D}_B$, following the norm-constrained gradient flows of [6], [7]. The ascent is started from the principal eigenvector of Z_{SDP}, from the best Monte Carlo (MC) sample and from random perturbations. The result is an admissible LTI perturbation, so its cost is a lower bound on J_N^*, and it satisfies first-order stationarity conditions. When $\mathrm{rank}(Z_{\mathrm{SDP}}) = 1$, it matches the SDP bound.

We provide an interactive interface [13] for changing the parameters (matrix A, B, Q, R, S and hyperparameter $N, \varepsilon_A, \varepsilon_B$) in real time using the solver MOSEK.

A. Pseudospectral Targeting and Trajectories for Large N

To illustrate the PEP approach, we apply it to three 2D systems (Fig. 1). We chose the following nominal matrices A because they each represent one of the three distinct worst-case behaviors we observed, making the differences visually clear:

$$A_1 = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix},\ A_2 = \begin{bmatrix} -0.9 & 1 \\ 0 & 1 \end{bmatrix} \text{ and } A_3 = \begin{bmatrix} 1 & 0.9 \\ -0.8 & 1 \end{bmatrix},$$

with $B = \begin{bmatrix} 0 & 1 \end{bmatrix}^\top$ and $Q = R = S = \mathrm{Id}$ for $N = 20$.

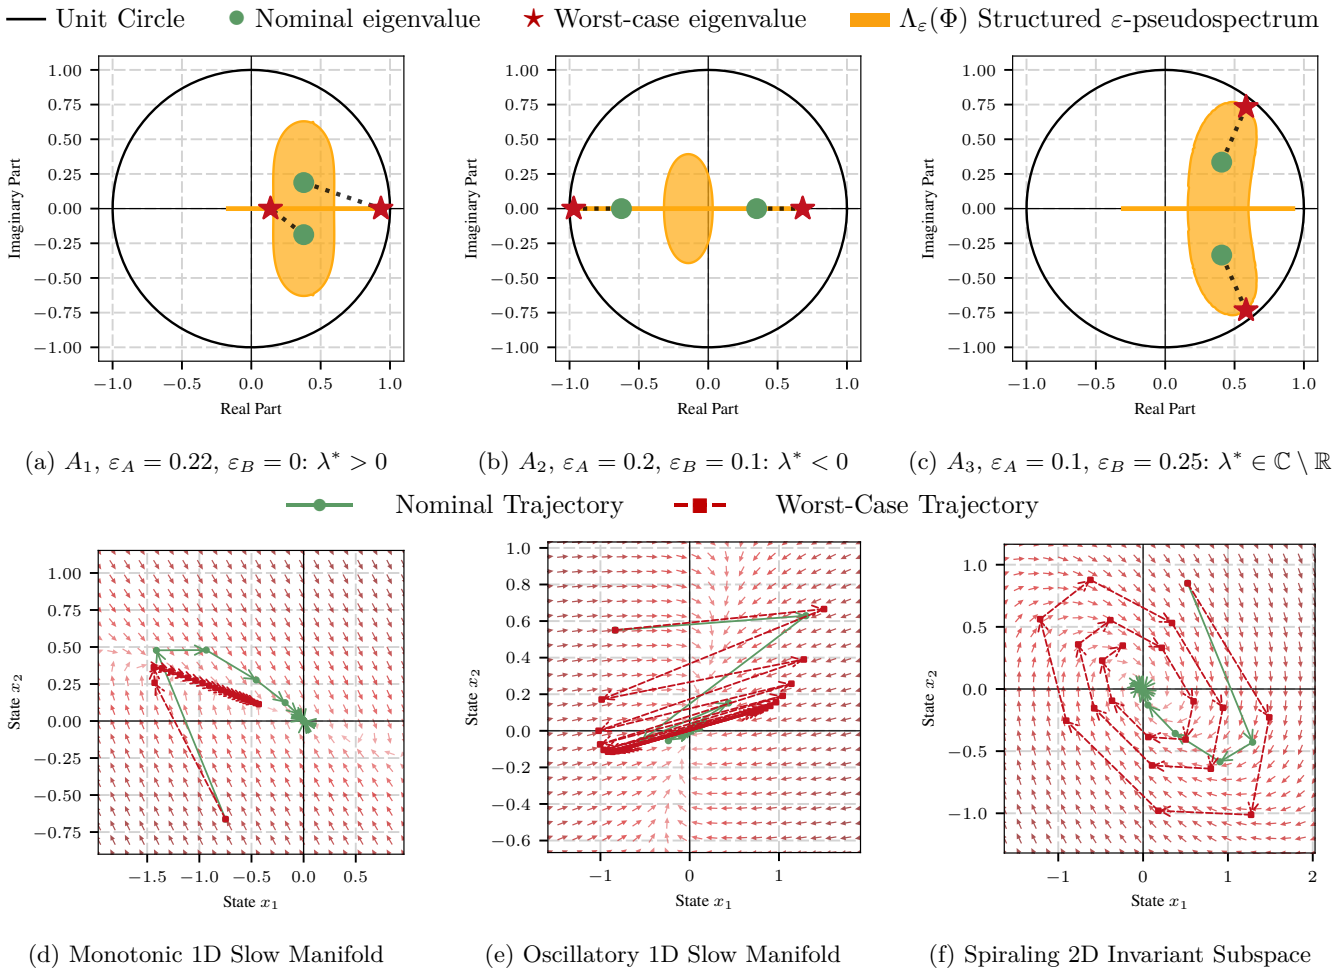
Varying the nominal system matrices reveals three classes of worst-case behavior, which we interpret through their spectral geometry (Fig. 1, top row). This spectral view is used as a diagnostic of the PEP solution: in general, the finite-horizon objective is not determined by the perturbed eigenvalues alone (Example 1). For given uncertainty bounds ε_A and ε_B, the structured ε-pseudospectrum of the closed-loop system [7] is the set of all possible eigenvalues for the perturbed matrix $\Phi_\Delta = \Phi + \Delta_A + \Delta_B K$ under valid LTI perturbations, computed exactly via the real stability radius formula [14]:

$$\begin{aligned} \Lambda_\varepsilon(\Phi) := \{\lambda \in \mathbb{C} \mid & \det(\lambda I - (\Phi + \Delta_A + \Delta_B K)) = 0, \\ & \text{for some } \|\Delta_A\|_2 \le \varepsilon_A,\ \|\Delta_B\|_2 \le \varepsilon_B\}. \end{aligned}$$

As the red stars show, the optimal perturbation often pushes a dominant eigenvalue toward the boundary of $\Lambda_\varepsilon(\Phi)$. For large horizons, this behavior approaches the spectral-radius-maximizing perturbations predicted by classical μ-analysis and pseudospectral theory [4], [15]. However, these asymptotic tools do not fully characterize the finite-horizon optimum: PEP captures transient dynamics and invariant-direction effects, so the worst-case cost is not determined by the eigenvalues alone. In our numerical evaluations, the pseudospectral heuristic that maximizes the spectral radius of Φ_Δ over the admissible perturbations [6] therefore yields a lower cost than PEP (e.g., on system A_1, $J_{20}^\rho \approx 20.95$ versus $J_{20}^{\mathrm{PEP}} \approx 21.06$). Since this maximizer is not unique, J_N^ρ denotes the largest finite-horizon cost among the sampled maximizers, the most favorable choice for the heuristic.

This occurs because the finite-horizon cost depends on transient amplification, whereas asymptotic tools mainly focus on pole locations. PEP does not impose or target any particular spectral configuration: it optimizes directly over admissible finite-horizon trajectories. A posteriori, we observe that the resulting perturbations may remain slightly inside $\Lambda_\varepsilon(\Phi)$ depending on the cost matrices while modifying the associated invariant directions in a way that produces a larger transient cost.

The maximum spectral radius on the boundary can occur at a real positive, real negative, or complex eigenvalue, depending on the nominal eigenvalues. This shift changes the state-space trajectories (bottom row of Fig. 1). The perturbation modifies the dominant



(a) A_1, $\varepsilon_A = 0.22$, $\varepsilon_B = 0$: $\lambda^* > 0$ (b) A_2, $\varepsilon_A = 0.2$, $\varepsilon_B = 0.1$: $\lambda^* < 0$ (c) A_3, $\varepsilon_A = 0.1$, $\varepsilon_B = 0.25$: $\lambda^* \in \mathbb{C} \setminus \mathbb{R}$

(d) Monotonic 1D Slow Manifold (e) Oscillatory 1D Slow Manifold (f) Spiraling 2D Invariant Subspace

Fig. 1: Worst-case perturbation: spectral representation (top row) and trajectories (bottom row) extracted via PEP and projected gradient ascent.

eigenspace of Φ_{Δ^*} to confine the state energy to a slowly decaying invariant manifold. In our computations, we observe three cases:

- *Monotonic 1D Slow Manifold ($\lambda^* > 0$, Fig. 1d):* When the dominant worst-case eigenvalue is real and positive, the optimal perturbation applies a rank-one shear opposing the nominal damping. The trajectory collapses onto a 1D invariant line (a "slow manifold") and moves monotonically toward the origin.
- *Oscillatory 1D Slow Manifold ($\lambda^* < 0$, Fig. 1e):* If the perturbed eigenvalue on the boundary is real and negative, the trajectory again collapses onto a 1D line. The negative eigenvalue causes the state to alternate signs, jumping back and forth across the origin at each step while slowly decaying.
- *Spiraling 2D Invariant Subspace ($\lambda^* \in \mathbb{C}\setminus\mathbb{R}$, Fig. 1f):* When the worst-case eigenvalue is a complex conjugate pair, Φ_{Δ^*} acts on a 2D invariant plane, and the state trajectory spirals. The perturbation changes the rotational frequency and minimizes radial contraction, creating a slowly decaying spiral.

In classical matrix perturbation theory, asymptotic variations that maximize the spectral radius correspond to rank-one (for real eigenvalues) or rank-two (for complex conjugate pairs) perturbations [16]. The PEP solutions recover this geometry. The primal formulation collapses finite-horizon trajectories onto the slow manifolds dictated by worst-case μ-analysis.

B. Pole-Placement bifurcation with respect to N

The optimal adversarial LTI perturbation depends directly on the length of the evaluation horizon. We illustrate this using the system A_1 with the same B, Q and R but terminal cost $S = \mathrm{diag}(0, 1)$, which penalizes only the second state component. We set the uncertainty bounds $\varepsilon_A = 0.15$ and $\varepsilon_B = 0.05$. The different worst-case pole placements under those settings are shown in Fig. 2.

As the horizon N increases, the initial transient growth contributes less to the total cumulative cost J_N. The slow asymptotic decay of the state begins to drive the cost instead. As N approaches 15, the perturbation stops manipulating transient eigenvectors. The poles migrate toward the rightmost edge of the structured ε-

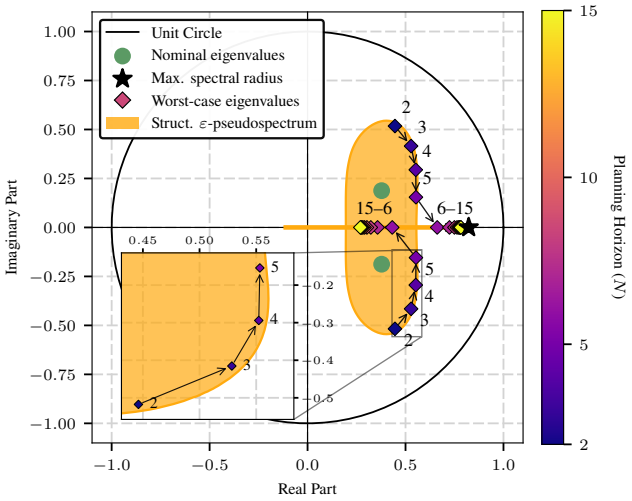


Fig. 2: Spectral migration of the worst-case closed-loop poles as a function of the finite planning horizon N. The colored diamonds represent the worst-case poles extracted from the PEP solution.

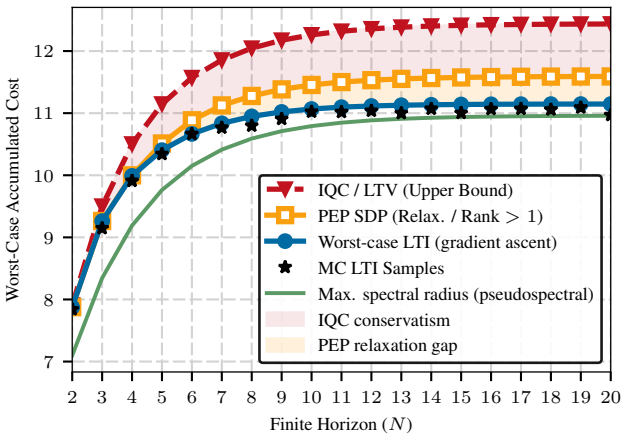


Fig. 3: Finite-horizon accumulated cost bounds for system A_1 with $\varepsilon_A = 0.15$ and $\varepsilon_B = 0.05$ for the different approaches.

pseudospectrum to maximize the asymptotic time constant. Along the way, the poles undergo a geometric bifurcation. For horizons $N \leq 5$, the optimal perturbation maintains a complex conjugate pole pair. For $N \geq 6$, the poles meet on the real axis and split. This implies that the worst-case trajectories transition from 2D spiraling behavior (Fig. 1f) to a 1D slow manifold one (Fig. 1d).

Evaluating controllers over finite horizons therefore requires more than identifying the worst-case asymptotic poles on the pseudospectrum boundary.

V. Performance: Classical vs. PEP

We now compare the worst-case performance obtained from the different approaches discussed in this paper: standard IQC [5], the relaxed PEP SDP, the worst-case LTI perturbation extracted by projected gradient ascent, and the pseudospectral heuristic defined in Section IV-A. We evaluate these methods on system A_1 to solve problem (4), using the same matrix B as in the previous section. The cost matrices are set to $Q = R = S = \mathrm{Id}$, with perturbation bounds $\varepsilon_A = 0.15$ and $\varepsilon_B = 0.05$. Fig. 3 shows how the accumulated cost changes with the horizon N under each approach.

The IQC framework (red dashed line) bounds the cost over all allowable parameter variations, including LTV sequences [5]. Because the physical perturbation in (4) is LTI, the relaxed PEP formulation (orange squares) tightens this bound by enforcing time invariance. The pink shaded area shows the cost difference between the LTV and LTI bounds.

As discussed at the end of Section III-C, dropping the rank constraint in the PEP SDP introduces conservatism, as it admits infeasible trajectories; this gap is shown by the orange shaded region. This gap initially grows as N increases then plateaus similarly to the IQC vs Relaxed SDP one. However, the LTI worst case extracted by projected gradient ascent (blue circles) is admissible, so its cost is a lower bound on J_N^*. It is first-order stationary but not certified globally optimal. Its cost is higher than what random Monte Carlo sampling (black stars) typically finds, as random sampling struggles to hit the narrow geometric manifold where the worst-case parameters lie.

The pseudospectral heuristic (green solid line) yields the lowest cost of the group. As discussed in Section IV-A, maximizing the spectral radius places the dominant eigenvalue on the boundary of the pseudospectrum and thus minimizes the long-term decay rate. However, it misses the short-term transient growth. The heuristic does not align the eigenvectors with a worst-case initial condition x_0, leaving out a large portion of the finite-horizon cost for short horizons.

VI. Conclusions and Perspectives

Extracting worst-case trajectories with the PEP framework shows how finite-horizon robustness differs from infinite-horizon stability. We found that LTI perturbations align with specific initial conditions to cause short-term transient growth. For example, we observed that, as the planning horizon N changes, the failure mode of a system can transition from 2D spiraling invariant subspaces to 1D monotonic slow manifolds. Because of this transient behavior, analyzing the boundary eigenvalues of the structured pseudospectrum does not capture the full finite-time cost.

Comparing these PEP solutions to the standard IQC framework highlights an open question regarding LTV relaxations. While these relaxations usually introduce conservatism when bounding LTI uncertainties, we observed numerical instances where the IQC upper bounds exactly or closely match the LTI cost from our PEP

approach. Finding the algebraic conditions that explain when a worst-case LTV sequence inherently collapses to an LTI structure will help determine when the IQC (LTV) relaxation is tight.

Finally, sequence-based optimization offers a direct path for extending this analysis. Operator interpolation relies on algebraic inequalities rather than specific physical models. This makes it possible to replace the norm-bounded LTI constraints with data-driven uncertainty sets or sector-bounded nonlinearities. Future work could focus on integrating these finite-horizon certificates into the controller design process to synthesize control laws optimized against worst-case transient behavior.

References

[1] F. Amato, M. Ariola, and P. Dorato, "Finite-Time Control of Linear Systems Subject to Parametric Uncertainties and Disturbances," *Automatica*, vol. 37, no. 9, pp. 1459–1463, 2001.

[2] D. Q. Mayne, M. M. Seron, and S. V. Raković, "Robust Model Predictive Control of Constrained Linear Systems with Bounded Disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

[3] I. Abel, D. Steeves *et al.*, "Prescribed-Time Safety Design for Strict-Feedback Nonlinear Systems," *IEEE Transactions on Automatic Control*, vol. 69, no. 3, pp. 1464–1479, 2024.

[4] A. Packard and J. Doyle, "The Complex Structured Singular Value," *Automatica*, vol. 29, no. 1, pp. 71–109, 1993.

[5] A. Megretski and A. Rantzer, "System Analysis via Integral Quadratic Constraints," *IEEE Transactions on Automatic Control*, vol. 42, no. 6, pp. 819–830, 1997.

[6] N. Guglielmi, C. Lubich, and S. Sicilia, "Rank-1 Matrix Differential Equations for Structured Eigenvalue Optimization," *SIAM Journal on Numerical Analysis*, vol. 61, no. 4, pp. 1737–1762, 2023.

[7] N. Guglielmi and C. Lubich, "Transient Dynamics Under Structured Perturbations: Bridging Unstructured and Structured Pseudospectra," *SIAM Journal on Numerical Analysis*, vol. 63, no. 2, pp. 908–930, 2025.

[8] P. Seiler, R. M. Moore *et al.*, "Finite Horizon Robustness Analysis of LTV Systems Using Integral Quadratic Constraints," *Automatica*, vol. 100, pp. 135–143, 2019.

[9] Y. Drori and M. Teboulle, "Performance of First-Order Methods for Smooth Convex Minimization: A Novel Approach," *Mathematical Programming*, vol. 145, no. 1, pp. 451–482, 2014.

[10] A. B. Taylor, J. M. Hendrickx, and F. Glineur, "Smooth Strongly Convex Interpolation and Exact Worst-Case Performance of First-Order Methods," *Mathematical Programming*, vol. 161, no. 1, pp. 307–345, 2017.

[11] N. Bousselmi, J. M. Hendrickx, and F. Glineur, "Interpolation Conditions for Linear Operators and Applications to Performance Estimation Problems," *arXiv*, 2024.

[12] Z.-Q. Luo, W.-K. Ma *et al.*, "Semidefinite Relaxation of Quadratic Optimization Problems," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 20–34, 2010.

[13] A. Couthures, "Worst-case LTI perturbations of finite-horizon LQR: A performance estimation problem approach," 10.5281/zenodo.23161489, 2026. https://github.com/AnthonyCouthures/PEP-LQR

[14] L. Qiu, B. Bernhardsson *et al.*, "A Formula for Computation of the Real Stability Radius," *Automatica*, vol. 31, no. 6, pp. 879–890, 1995.

[15] D. Hinrichsen and A. J. Pritchard, "Real and Complex Stability Radii: A Survey," in *Control of Uncertain Systems: Proceedings of an International Workshop Bremen, West Germany, June 1989*, D. Hinrichsen and B. Mårtensson, Eds. Boston, MA: Birkhäuser, 1990, pp. 119–162.

[16] L. N. Trefethen and M. Embree, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press, 2005.