Perturbation-based Regret Analysis of MPC in LTV systems with General Well-Conditioned Costs

*Yiheng Lin, **Yang Hu**, Haoyuan Sun,*
Guannan Qu, Guanya Shi, Adam Wierman

IIIS, Tsinghua University

March 28, 2022

Optimal Control: Formulation

Consider a *Linear Time-Varying* (*LTV*) system

$$x_{t+1} = A_t x_t + B_t u_t + w_t.$$

We want to design a controller $u_t = f(x_{0:t}, \cdots)$ that minimizes

$$\mathcal{J}(u_{0:T-1}) = \sum_{t=0}^{T} \big(f_t(x_t) + c_t(u_{t-1})\big).$$

where f_t is *hitting cost*, and c_t is *transition cost*.

Optimal Control: Formulation

Consider a *Linear Time-Varying* (*LTV*) system

$$x_{t+1} = A_t x_t + B_t u_t + w_t.$$

We want to design a controller $u_t = f(x_{0:t}, \cdots)$ that minimizes

$$\mathcal{J}(u_{0:T-1}) = \sum_{t=0}^{T} \big(f_t(x_t) + c_t(u_{t-1})\big).$$

where f_t is *hitting cost*, and c_t is *transition cost*.

- The problem adopts a finite *horizon* of T steps.
- We allow a large family of online controllers, not restricted to linear feedback controllers.

Optimal Control: History and Today

- In classical optimal control theory ...
  - Linear feedback controllers ($u_t = -Kx_t$) in LTI systems.
  - Dynamical matrices (A, B) are known to the controller.
  - Model uncertainty formulated as additive noises.
  - Explicit controller design (Riccati) and optimality (HJB).

Optimal Control: History and Today

- In classical optimal control theory ...
  - Linear feedback controllers ($u_t = -Kx_t$) in LTI systems.
  - Dynamical matrices (A, B) are known to the controller.
  - Model uncertainty formulated as additive noises.
  - Explicit controller design (Riccati) and optimality (HJB).
- In online learning setting ("control meets learning") ...
  - Controller may be optimized over an arbitrary family ($K \in \mathcal{K}$).
  - Deal with general LTV systems and unknown dynamics.
  - Adopt performance metrics in online learning context (regret, competitive ratio, sample complexity, etc.)

Optimal Control: Taking a Learning Lens

- Existing literature covers an interesting range of topics:
  - Regret in predictive setting: [1], [2], [3], etc.
  - Regret in adversarial setting: [4], [5], etc.
  - Competitive ratio: [6], [7], [8] etc.
  - Sample complexity: [9], [10], etc.

Optimal Control: Taking a Learning Lens

- Existing literature covers an interesting range of topics:
  - Regret in predictive setting: [1], [2], [3], etc.
  - Regret in adversarial setting: [4], [5], etc.
  - Competitive ratio: [6], [7], [8] etc.
  - Sample complexity: [9], [10], etc.
- However, there are also certain limitations in these results:
  - Most still focus on LTI systems and/or quadratic costs.
  - Use modified controller to achieve better performance.

Optimal Control: Taking a Learning Lens

- Existing literature covers an interesting range of topics:
  - Regret in predictive setting: [1], [2], [3], etc.
  - Regret in adversarial setting: [4], [5], etc.
  - Competitive ratio: [6], [7], [8] etc.
  - Sample complexity: [9], [10], etc.
- However, there are also certain limitations in these results:
  - Most still focus on LTI systems and/or quadratic costs.
  - Use modified controller to achieve better performance.
- How about standard controllers in general settings?
  - Go beyond *LQR* to *LTV systems* + *well-conditioned costs*?
  - Analyze standard *Model Predictive Control* (*MPC*)?

Model Predictive Control

At time step t, the MPC controller:

- has access to (exact) prediction of dynamics of k future steps
$$\{\vartheta_\tau = (A_\tau, B_\tau, w_\tau, f_\tau, c_\tau)\}_{\tau=t}^{t+k};$$
- observes current state x_t (and past trajectory, if necessary);
- is required to output control input u_t for *one* step.

Model Predictive Control

At time step t, the MPC controller:

- has access to (exact) prediction of dynamics of k future steps

$$\{\vartheta_\tau = (A_\tau, B_\tau, w_\tau, f_\tau, c_\tau)\}_{\tau=t}^{t+k};$$

- observes current state x_t (and past trajectory, if necessary);
- is required to output control input u_t for *one* step.

Why do we allow exact predictions?

- ~~That's all what we can do before NeurIPS deadline.~~
- Approximate some realistic cases (e.g., power plant control).
- Help understand the idea, but not to get stuck in details.

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

ϑ_t	ϑ_{t+1}	ϑ_{t+2}	ϑ_{t+3}	ϑ_{t+4}	ϑ_{t+5}	ϑ_{t+6}	ϑ_{t+7}

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

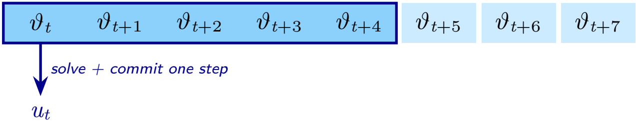

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

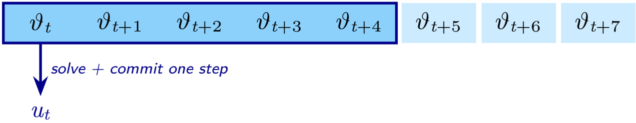

In each step, it solves the optimization problem

$$\tilde{\psi}_t^k(x, \zeta; F) := \underset{y_{0:k}, v_{0:k-1}}{\arg\min} \quad \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}\left(v_{\tau-1}\right) + F(y_k)$$

$$\text{subject to} \quad y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1},$$

$$y_0 = x.$$

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

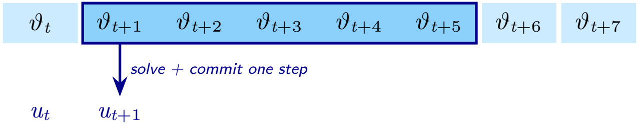

In each step, it solves the optimization problem

$$\tilde{\psi}_t^k(x, \zeta; F) := \underset{y_{0:k}, v_{0:k-1}}{\arg\min} \quad \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k)$$

$$\text{subject to} \quad y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1},$$

$$y_0 = x.$$

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

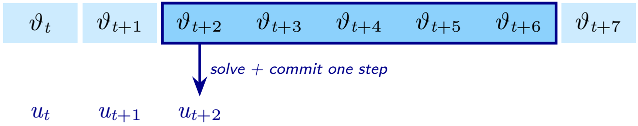

In each step, it solves the optimization problem

$$\begin{aligned} \tilde{\psi}_t^k(x, \zeta; F) := \underset{y_{0:k}, v_{0:k-1}}{\arg\min} \quad & \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k) \\ \text{subject to} \quad & y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1}, \\ & y_0 = x. \end{aligned}$$

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

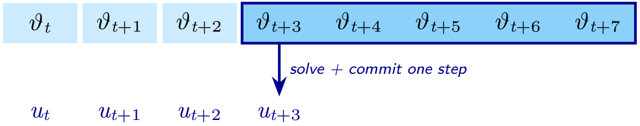

In each step, it solves the optimization problem

$$\tilde{\psi}_t^k(x, \zeta; F) := \underset{y_{0:k}, v_{0:k-1}}{\arg\min} \quad \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k)$$

$$\text{subject to} \quad y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1},$$

$$y_0 = x.$$

Model Predictive Control

Basically, MPC = online control with a *receding horizon*.

ϑ_t	ϑ_{t+1}	ϑ_{t+2}	ϑ_{t+3}	ϑ_{t+4}	ϑ_{t+5}	ϑ_{t+6}	ϑ_{t+7}

u_t u_{t+1} u_{t+2} u_{t+3} $\cdots$

In each step, it solves the optimization problem

$$\tilde{\psi}_t^k(x, \zeta; F) := \arg\min_{y_{0:k}, v_{0:k-1}} \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k)$$

$$\text{subject to} \quad y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1},$$

$$y_0 = x.$$

Preliminaries: Performance Metrics

Existing literature adopts standard online learning metrics:

- Static regret (against linear controller K^*)[5, 11]

$$r_{\mathrm{s}} := \sup_{x_0, w_{0:T-1}} \left\{ \mathcal{J}(\mathrm{ALG}) - \inf_{K} \mathcal{J}(K) \right\}.$$

- Dynamic regret (against optimal controller u^*)[1, 3]

$$r_{\mathrm{d}} := \sup_{x_0, w_{0:T-1}} \left\{ \mathcal{J}(\mathrm{ALG}) - \inf_{u_{0:T-1}} \mathcal{J}(u_{0:T-1}) \right\}.$$

- Competitive ratio (against optimal controller u^*)[6, 8]

$$c := \sup_{x_0, w_{0:T-1}} \left\{ \frac{\mathcal{J}(\mathrm{ALG})}{\inf_{u_{0:T-1}} \mathcal{J}(u_{0:T-1})} \right\}.$$

Preliminaries: Assumptions on System Model

Assumption 1. Assume f_t and c_t are well-conditioned, i.e.,

- strongly convex and strongly smooth;
- twice continuously differentiable;
- non-negative, and $f_t(0) = c_t(0) = 0$.

Assumption 2. Assume the dynamical matrices are bounded as

$$\|A_t\| \leq a, \ \|B_t\| \leq b, \ \|B_t^{\dagger}\| \leq b',$$

where $B_t^{\dagger}$ denotes the Moore-Penrose pseudo-inverse of B_t.

Preliminaries: Assumption on Controllability

We make an assumption slightly stronger than controllability:

Assumption 3. The system is (d, σ)-*uniform controllable*, i.e., there exists a positive constant σ such that

$$\sigma_{\min}(M(t, d)) \geq \sigma, \ \forall t = 0, \cdots, T - d.$$

Here the *controllability matrix* is defined as

$$M(t, p) := \left[\Phi(t+p, t+1)B_t, \Phi(t+p, t+2)B_{t+1}, \cdots, \Phi(t+p, t+p)B_{t+p}\right],$$

where $\Phi(t_2, t_1) := A_{t_2-1}A_{t_2-2}\cdots A_{t_1}$, and the *controllability index* d is the smallest p such that $M(t, p)$ is of full row rank.

Preliminaries: Optimization Problems

- Single-ended optimization problem (solved by MPC$_k$):

$$\tilde{\psi}_t^k(x, \zeta; F) := \arg\min_{y_{0:k}, v_{0:k-1}} \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1}) + F(y_k)$$
$$\text{subject to} \quad y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1},$$
$$y_0 = x.$$

- Double-ended optimization problem (auxiliary):

$$\psi_t^k(x, \zeta, z) := \arg\min_{y_{0:k}, v_{0:k-1}} \sum_{\tau=1}^{k} f_{t+\tau}(y_\tau) + \sum_{\tau=1}^{k} c_{t+\tau}(v_{\tau-1})$$
$$\text{subject to} \quad y_\tau = A_{t+\tau-1} y_{\tau-1} + B_{t+\tau-1} v_{\tau-1} + \zeta_{\tau-1},$$
$$y_0 = x, \; y_k = z.$$

Structure of the Presentation

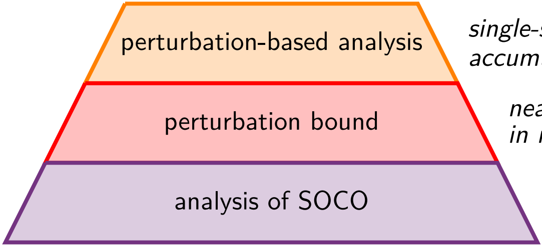

single-step errors do not accumulate over time

nearby initial states result in nearby trajectories

banded matrix inverse implies exp-decay

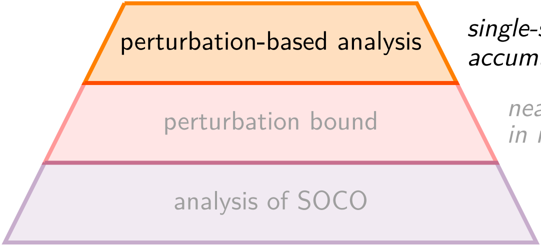

single-step errors do not accumulate over time

nearby initial states result in nearby trajectories

banded matrix inverse implies exp-decay

Assuming Exponential-Decaying Perturbation Bound ...

- Intuitively, nearby initial states result in nearby trajectories.
  - We want the trajectories that only differ in initial states converge exponentially fast to each other, i.e.,

$$\|\tilde{\psi}_t^p(x, \zeta; F)_{y_h} - \tilde{\psi}_t^p(x', \zeta; F)_{y_h}\| \leq C\lambda^h \|x - x'\|,$$
$$\|\psi_t^p(x, \zeta, z)_{y_h} - \psi_t^p(x', \zeta, z)_{y_h}\| \leq C\lambda^h \|x - x'\|.$$

  - Further, the perturbation bound can be extended to ζ and z, so that the right-hand side is in the form

$$C\left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{|h-\tau|} \|\zeta_\tau - \zeta'_\tau\| + \lambda^{p-h} \|z - z'\|\right).$$

What Does the Perturbation Bound Imply?

Lemma 1 (stability of $\tilde{\psi}$)

For any F, $\|\tilde{\psi}_t^p(x, \zeta; F)_{y_h}\| \le C\lambda^h \|x\| + \frac{2C}{1-\lambda} \sup_\tau \|\zeta_\tau\|$.

What Does the Perturbation Bound Imply?

Lemma 1 (stability of $\tilde{\psi}$)

For any F, $\|\tilde{\psi}_t^p(x,\zeta;F)_{y_h}\| \le C\lambda^h\|x\| + \frac{2C}{1-\lambda}\sup_\tau\|\zeta_\tau\|$.

Proof Sketch. Compare with a noiseless trajectory starting from 0, where the optimal trajectory is to stay at 0. Formally,

$$\begin{aligned}\|\tilde{\psi}_t^p(x,\zeta;F)_{y_h}\| &= \left\|\tilde{\psi}_t^p(x,\zeta;F)_{y_h} - \tilde{\psi}_t^p(0,0;F)_{y_h}\right\| \\ &\le C\left(\lambda^h\|x\| + \sum_{\tau=0}^{p-1}\lambda^{|h-\tau|}\|\zeta_\tau\|\right) \\ &\le C\lambda^h\|x\| + \frac{2C}{1-\lambda}\sup_\tau\|\zeta_\tau\|.\end{aligned}$$

What Does the Perturbation Bound Imply?

Lemma 1 (stability of $\tilde{\psi}$)

For any F, $\|\tilde{\psi}_t^p(x, \zeta; F)_{y_h}\| \leq C\lambda^h \|x\| + \frac{2C}{1-\lambda} \sup_\tau \|\zeta_\tau\|$.

- Basically, an exponential-decaying perturbation bound implies *stability* of all $\tilde{\psi}$-trajectories.
  - Intuitively, it approaches $0, \cdots, 0$ (recall 0 is an equilibrium), another $\tilde{\psi}$-trajectory, exponentially fast.

What Does the Perturbation Bound Imply?

Lemma 1 (stability of $\tilde{\psi}$)

For any F, $\|\tilde{\psi}_t^p(x, \zeta; F)_{y_h}\| \leq C\lambda^h \|x\| + \frac{2C}{1-\lambda} \sup_\tau \|\zeta_\tau\|$.

- Basically, an exponential-decaying perturbation bound implies *stability* of all $\tilde{\psi}$-trajectories.
  - Intuitively, it approaches $0, \cdots, 0$ (recall 0 is an equilibrium), another $\tilde{\psi}$-trajectory, exponentially fast.
- A helpful corollary is that OPT $\approx$ MPC$_\infty$.
  - ∞ refers to a sufficiently long prediction window.
  - $\approx$ comes from the fact that OPT has no terminal cost.
  - Both are stable $\tilde{\psi}$-trajectories, so their distance is at most twice the constant in Lemma 1.

What Does the Perturbation Bound Imply?

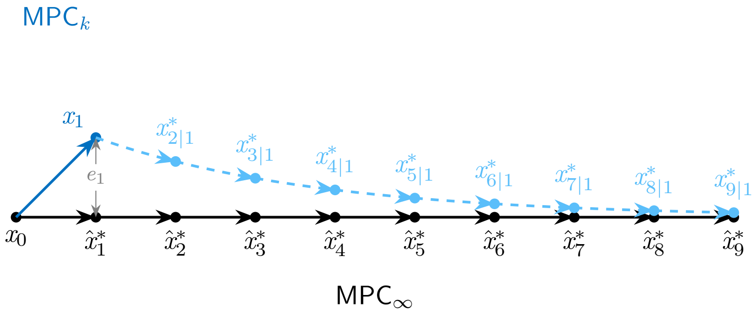

What Does the Perturbation Bound Imply?

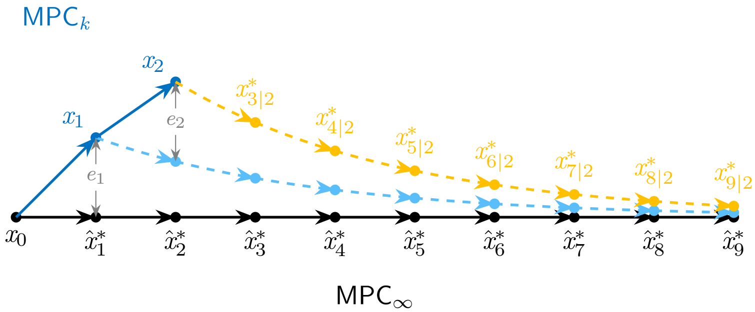

What Does the Perturbation Bound Imply?

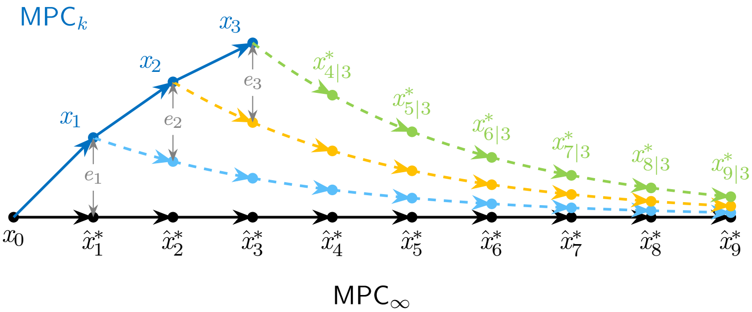

What Does the Perturbation Bound Imply?

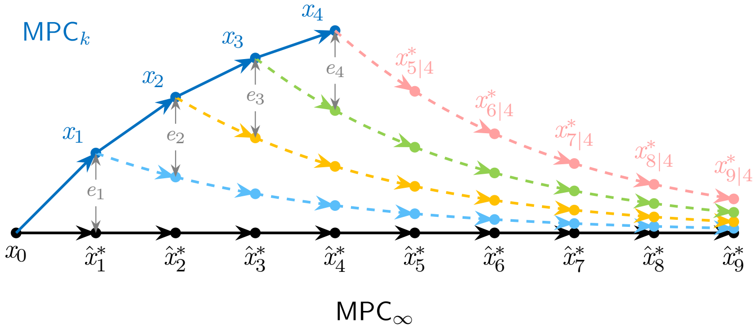

What Does the Perturbation Bound Imply?

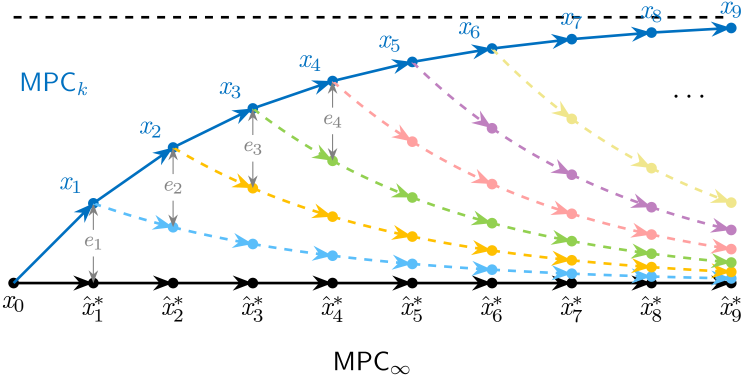

What Does the Perturbation Bound Imply?

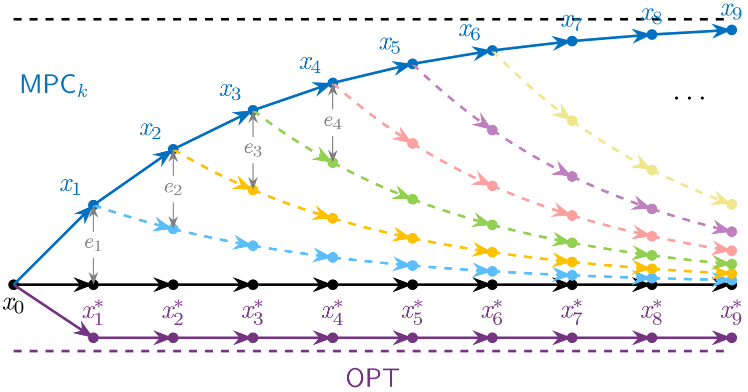

What Does the Perturbation Bound Imply?

Lemma 2

If the one-future-step error satisfies $e_t := \|x_{t+1} - \hat{x}^*_{t+1|t}\| = O(\lambda^k)$, *then we also have* $\|x_t - x_t^*\| = O(\lambda^k)$.

What Does the Perturbation Bound Imply?

Lemma 2

If the one-future-step error satisfies $e_t := \|x_{t+1} - \hat{x}^*_{t+1|t}\| = O(\lambda^k)$, *then we also have* $\|x_t - x^*_t\| = O(\lambda^k)$.

Proof sketch. First, bound the difference $\|x_t - \hat{x}^*_t\|$:

$$\begin{aligned}\|x_t - \hat{x}^*_t\| &\le \|x_t - x^*_{t|t-1}\| + \sum_{i=1}^{t-1} \|x^*_{t|t-i} - x^*_{t|t-(i+1)}\| \\ &\le \|x_t - x^*_{t|t-1}\| + \sum_{i=1}^{t-1} C\lambda^i \|x_{t-i} - x^*_{t-i|t-(i+1)}\| \\ &= O(\lambda^k).\end{aligned}$$

What Does the Perturbation Bound Imply?

Lemma 2

If the one-future-step error satisfies $e_t := \|x_{t+1} - \hat{x}^*_{t+1|t}\| = O(\lambda^k)$*, then we also have* $\|x_t - x^*_t\| = O(\lambda^k)$*.*

Proof sketch. Next, to bound $\|\hat{x}^*_t - x^*_t\|$, note that by Lemma 1,

$$\|\hat{x}^*_T - x^*_T\| \le \|\hat{x}^*_T\| + \|x^*_T\| \le 2C\lambda^T \|x\| + \frac{4C}{1-\lambda} \sup_\tau \|\zeta_\tau\| .$$

Therefore, by perturbation bound on terminal state z,

$$\|\hat{x}^*_t - x^*_t\| \le \lambda^{T-t} \|\hat{x}^*_T - x^*_T\| \le \lambda^k \left(2C\lambda^T \|x\| + \frac{4C}{1-\lambda} \sup_\tau \|\zeta_\tau\| \right)$$

for any $t \le T - k$. For the remaining part, we simply revise the algorithm to enforce $\hat{x}^*_t = x^*_t$ for $t > T - k$.

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT $\approx$ MPC$_\infty$.

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT ≈ MPC_∞.

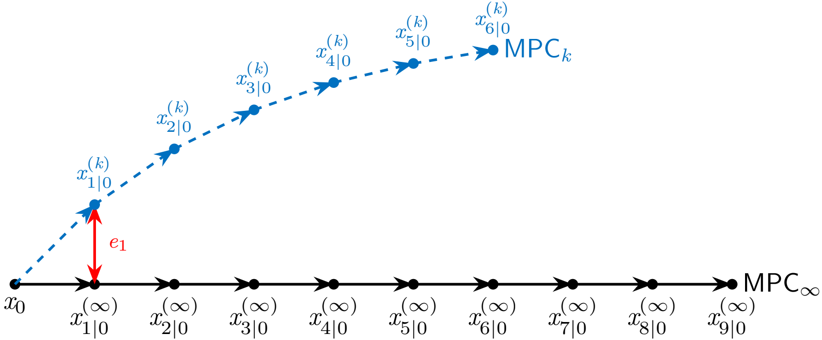

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT $\approx$ MPC$_\infty$.

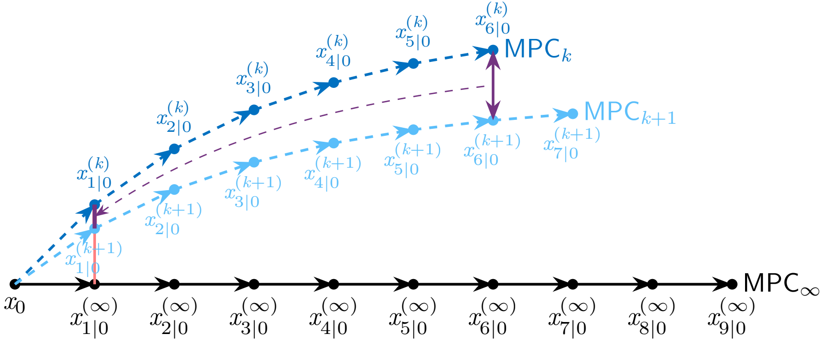

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT $\approx$ MPC$_\infty$.

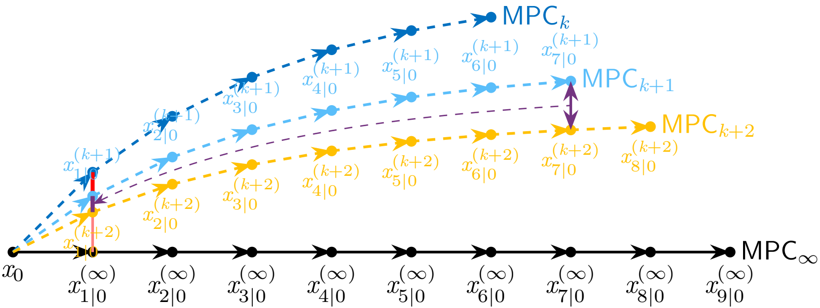

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT $\approx$ MPC$_\infty$.

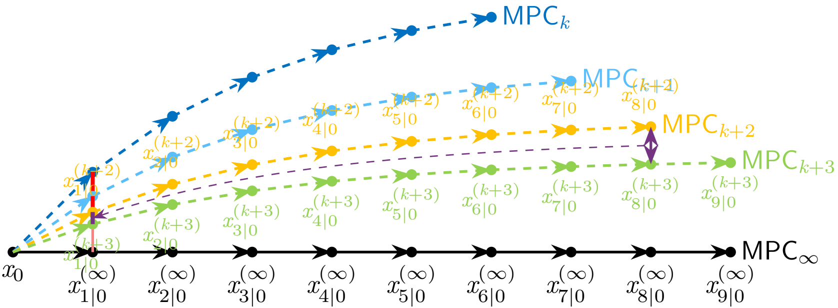

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT $\approx$ MPC$_\infty$.

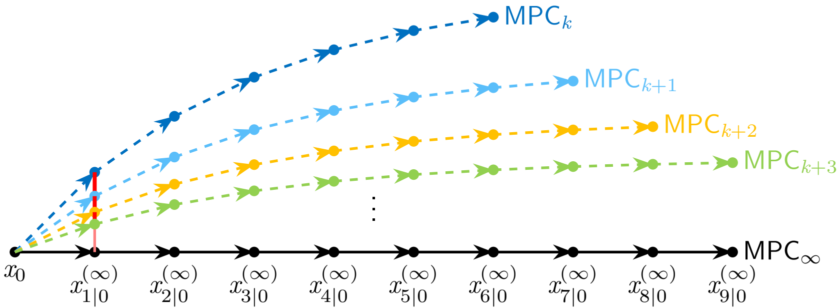

How to Bound the One-Future-Step Error?

- To show $e_t = O(\lambda^k)$, *just telescope over k!*
  - Intuitively, incrementing k by 1 does not change the trajectory much (in fact, by only $O(\lambda^k)$), and OPT $\approx$ MPC$_\infty$.

Lemma 3 (one-future-step error)

With large enough k, for any $p \geq h \geq 1$ and $t < T - p$,

$$\left\| \tilde{\psi}_t^p(x_t, \cdots)_{y_h} - \tilde{\psi}_t^{p+1}(x_t, \cdots)_{y_h} \right\| \leq 2C\lambda^{p-h} \left(C\lambda^p \|x_t\| + \frac{2C}{1-\lambda} \sup_{0 \leq \tau \leq T-1} \|w_\tau\| \right).$$

Thus $\|\tilde{\psi}_t^k(x_t, \cdots)_{y_h} - \tilde{\psi}_t^\infty(x_t, \cdots)_{y_h}\| = O(\lambda^k)$ if $\|w_t\| \leq D$.

How to Bound the One-Future-Step Error?

Proof sketch. Let $D := \sup_\tau \|w_\tau\|$. It is obvious that

$$\left\|\tilde{\psi}_t^p(x_t,\zeta;F)_{y_h} - \tilde{\psi}_t^{p+1}(x_t,\zeta;F)_{y_h}\right\| = \|\psi_t^p(x_t,\zeta,z)_{y_h} - \psi_t^p(x_t,\zeta,z')_{y_h}\| \le C\lambda^{p-h}\|z - z'\| \le 2C\lambda^{p-h}\left(C\lambda^p\|x_t\| + \frac{2C}{1-\lambda}D\right),$$

where $z := \tilde{\psi}_t^p(x_t,\zeta;F)_{y_p}$ and $z' := \tilde{\psi}_t^{p+1}(x_t,\zeta;F)_{y_p}$, and the last inequality is due to the stability of $\tilde{\psi}$ (*see Lemma 1*).

How to Bound the One-Future-Step Error?

Proof sketch. Therefore, for $t \leq T - k$,

$$
\begin{aligned}
&\left\| \tilde{\psi}_t^k(x_t, \zeta; F)_{y_1} - \tilde{\psi}_t^{T-t}(x_t, \zeta; F)_{y_1} \right\| \\
\leq & \sum_{p=k}^{T-t} \left\| \tilde{\psi}_t^p(x_t, \zeta; F)_{y_1} - \tilde{\psi}_t^{p+1}(x_t, \zeta; F)_{y_1} \right\| \\
\leq & \sum_{p=k}^{\infty} 2C\lambda^{p-1} \left(C\lambda^p \|x_t\| + \frac{2C}{1-\lambda} D \right) \\
= & \frac{2C^2}{\lambda(1-\lambda^2)} \cdot \lambda^{2k} \|x_t\| + \frac{4C^2}{\lambda(1-\lambda)^2} \cdot \lambda^k D \\
= & O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta} \right) \lambda^k \right).
\end{aligned}
$$

Combining the Building Blocks ...

Theorem 4 (Main Theorem, informal)

For MPC$_k$, MPC with large enough prediction window k,

1. *Its closed-loop dynamics of is input-to-state stable;*
2. *It achieves an $O(\lambda^k T)$ dynamic regret if $\|w_t\| \leq D$;*
3. *It achieves a $1 + O(\lambda^k)$ competitive ratio if $F(x) = \mathbf{1}_{\{x \neq 0\}} \cdot \infty$.*

Combining the Building Blocks ...

Proof sketch. To show the regret bound, note that for any $\eta > 0$,

$$
\begin{aligned}
&\mathrm{cost}(\mathsf{MPC}_k) - (1+\eta)\mathrm{cost}(\mathsf{OPT}) \\
&= \left(\sum_{t=0}^{T-k-1} \iota_t^1(x_t, x_{t+1}) + \iota_{T-k}^k(x_{T-k}, x_T)\right) \\
&\qquad - (1+\eta)\left(\sum_{t=0}^{T-k-1} \iota_t^1(x_t^*, x_{t+1}^*) + \iota_{T-k}^k(x_{T-k}^*, x_T^*)\right) \\
&= \sum_{t=0}^{T-k-1} \left(\iota_t^1(x_t, x_{t+1}) - (1+\eta)\iota_t^1(x_t^*, x_{t+1}^*)\right) \\
&\qquad + \left(\iota_{T-k}^k(x_{T-k}, x_T) - (1+\eta)\iota_{T-k}^k(x_{T-k}^*, x_T^*)\right)
\end{aligned}
$$

Combining the Building Blocks ...

Proof sketch. To show the regret bound, note that for any $\eta > 0$,

$$\begin{aligned}
&\mathrm{cost}(\mathsf{MPC}_k) - (1+\eta)\mathrm{cost}(\mathsf{OPT}) \\
= &\sum_{t=0}^{T-k-1} \left(\iota_t^1(x_t, x_{t+1}) - (1+\eta)\iota_t^1(x_t^*, x_{t+1}^*)\right) \\
&\quad + \left(\iota_{T-k}^k(x_{T-k}, x_T) - (1+\eta)\iota_{T-k}^k(x_{T-k}^*, x_T^*)\right) \\
\le &\sum_{t=0}^{T-k-1} \left(\iota_t^1(x_t, x_{t+1}) - (1+\eta)\iota_t^1(x_t^*, x_{t+1}^*)\right) \\
&\quad + \left(\iota_{T-k}^k(x_{T-k}, x_T^*) - (1+\eta)\iota_{T-k}^k(x_{T-k}^*, x_T^*)\right)
\end{aligned}$$

Combining the Building Blocks ...

Proof sketch. To show the regret bound, note that for any $\eta > 0$,

$$
\begin{aligned}
&\mathrm{cost}(\mathsf{MPC}_k) - (1+\eta)\mathrm{cost}(\mathsf{OPT}) \\
\leq\ & \sum_{t=0}^{T-k-1} \left(\iota_t^1(x_t, x_{t+1}) - (1+\eta)\iota_t^1(x_t^*, x_{t+1}^*)\right) \\
&\quad + \left(\iota_{T-k}^k(x_{T-k}, x_T^*) - (1+\eta)\iota_{T-k}^k(x_{T-k}^*, x_T^*)\right) \\
\leq\ & \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_4}{2} \sum_{t=0}^{T-k-1} \left(\left\|x_t - x_t^*\right\|^2 + \left\|x_{t+1} - x_{t+1}^*\right\|^2\right) \\
&\quad + \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_0 + \ell_f}{2} \left\|x_{T-k} - x_{T-k}^*\right\|^2
\end{aligned}
$$

* Note that ι inherits Lipschitzness from f and c.

Combining the Building Blocks ...

Proof sketch. To show the regret bound, note that for any $\eta > 0$,

$$
\begin{aligned}
&\mathrm{cost}(\mathsf{MPC}_k) - (1+\eta)\mathrm{cost}(\mathsf{OPT}) \\
&\le \left(1 + \frac{1}{\eta}\right) \cdot L_4 \sum_{t=0}^{T-k-1} \|x_t - x_t^*\|^2 + \left(1 + \frac{1}{\eta}\right) \cdot \frac{L_0 + \ell_f}{2} \left\|x_{T-k} - x_{T-k}^*\right\|^2 \\
&\le \left(1 + \frac{1}{\eta}\right) O\left(\left(D + \frac{\lambda^k(\|x_0\| + D)}{\delta}\right)^2 \lambda^{2k} T\right),
\end{aligned}
$$

To bound $\mathrm{cost}(\mathsf{OPT})$, we consider a suboptimal controller inspired by the d-hop transformation (see next section), which forces the state to 0 every d steps. In this way we can show that $\mathrm{cost}(\mathsf{OPT}) \le O(D^2 T + \|x_0\|^2)$. Set $\eta = O(\lambda^k)$, and we obtain the regret bound.

Combining the Building Blocks ...

Proof sketch. The competitive ratio result can be shown in a similar (but more involved) way. The major concern here is that we need to somehow "lower bound" $\mathrm{cost}(\mathsf{OPT})$, which turns out to require a different one-future-step bound that involves x_t^*, namely

$$\left\|\tilde{\psi}_t^p(x_t, w_{t:t+p-1}; F)_{y_h} - \tilde{\psi}_t^{p+1}(x_t, w_{t:t+p}; F)_{y_h}\right\| \leq C\lambda^{p-h}\left(\left\|x_{t+p}^*\right\| + C\lambda^p \left\|x_t - x_t^*\right\| + C\left\|x_{t+p+1}^*\right\|\right).$$

On the other hand, the ISS result simply follows from

$$\|x_t\| \leq \|x_t - \hat{x}_t^*\| + \|\hat{x}_t^*\|.$$

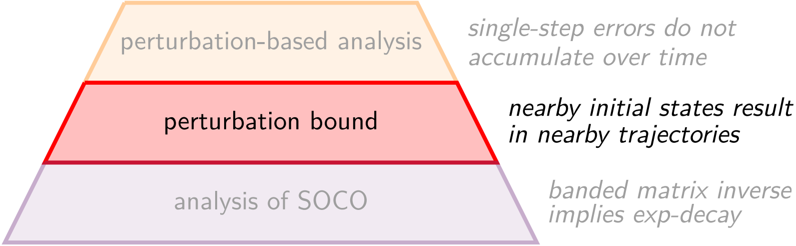
perturbation-based analysis
perturbation bound
analysis of SOCO
single-step errors do not accumulate over time
nearby initial states result in nearby trajectories
banded matrix inverse implies exp-decay

Step 1: Reduce MPC to SOCO

We basically want to remove the dynamics constraints to make the reduced optimization problem more approachable.

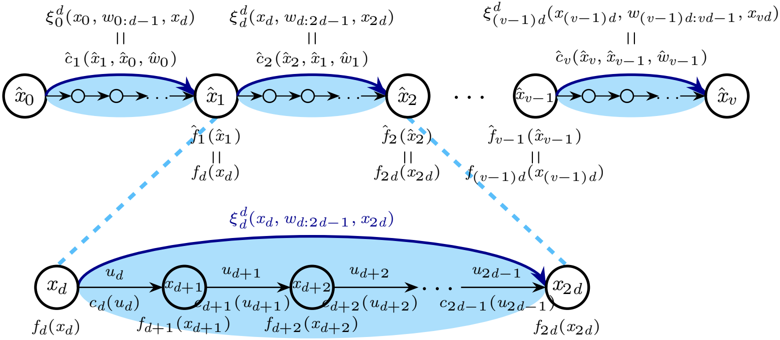

Step 1: Reduce MPC to SOCO

Lemma 5

For $p \geq d$, where d is the ***controllability index****, the above reduction satisfies the following:*

1. *The total cost of interval $(t, t+p)$, $\xi_t^p(x, \zeta, z)$, is convex and $L_2(p)$-strongly smooth in (x, ζ, z).*
2. *The trajectory of interval $(t, t+p)$ can be recovered by solving a double-ended constrained optimization problem, where the solution $\psi_t^p(x, \zeta, z)$ is $L_1(p)$-Lipschitz in (x, ζ, z);*

Step 1: Reduce MPC to SOCO

Proof sketch. The proof is based on the following fact.

Lemma 6

Consider an optimization problem over objective $f(x, y)$ parametrized by x, which is assume to be convex, L-strongly smooth in (x, y), μ-strongly convex in y, and continuously differentiable. Then

- *the optimal trajectory $y^*(x) := \arg\min_y f(x, y)$ is $\frac{L}{\mu}$-Lipschitz;*
- *the optimal cost $f^*(x) := \min_y f(x, y)$ is $\left(L + \frac{L^2}{\mu}\right)$-strongly smooth.*

It only suffices to note that the switching cost ξ_t^p is a solution to an unconstrained optimization problem, which can be done by "merging" states, control inputs, and noises between consecutive decision points.

Step 2: Establish Perturbation Bound for SOCO

Lemma 7 (perturbation bound for SOCO)

Given well-conditioned $\hat{f}_\tau$, $\hat{c}_\tau$, the optimal solution of SOCO

$$\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p) := \arg\min_{\hat{x}_{1:p-1}} \sum_{\tau=1}^{p-1} \hat{f}_\tau(\hat{x}_\tau) + \sum_{\tau=1}^{p} \hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1}),$$

satisfies

$$\left\| \hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p)_h - \hat{\psi}(\hat{x}_0', \hat{w}', \hat{x}_p')_h \right\| \leq$$

$$C_0 \left(\lambda_0^{h-1} \left\| \hat{x}_0 - \hat{x}_0' \right\| + \sum_{\tau=0}^{p-1} \lambda_0^{|h-\tau|-1} \left\| \hat{w}_\tau - \hat{w}_\tau' \right\| + \lambda_0^{p-h-1} \left\| \hat{x}_p - \hat{x}_p' \right\| \right).$$

Proof sketch. See the next section.

Step 3: Deduce Perturbation Bound for MPC

Theorem 8 (perturbation bound for MPC)

Given any (x, ζ, z) *and* (x', ζ', z'), *for all time steps* t,

$$\|\tilde{\psi}_t^p(x, \zeta; F)_{y_h} - \tilde{\psi}_t^p(x', \zeta'; F)_{y_h}\| \leq C\left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{|h-\tau|} \|\zeta_\tau - \zeta'_\tau\|\right)$$

$$\|\psi_t^p(x, \zeta, z)_{y_h} - \psi_t^p(x', \zeta', z')_{y_h}\| \leq C\left(\lambda^h \|x - x'\| + \sum_{\tau=0}^{p-1} \lambda^{|h-\tau|} \|\zeta_\tau - \zeta'_\tau\| + \lambda^{p-h} \|z - z'\|\right).$$

Step 3: Deduce Perturbation Bound for MPC

Proof sketch. Suppose $ud \le h < (u+1)d$ and $p = vd + r$, and we shall select the decision points as

$$y_0, y_d, \cdots, y_{(u-1)d}, y_h, y_{(u+2)d}, \cdots, y_{(v-1)d}, y_p,$$

or equivalently $y_{i_0}, \cdots, y_{i_{v-1}}$. Let $\hat{w}_{\tau-1} := \zeta_{i_{\tau-1}:i_\tau - 1}$ be the noise for SOCO, and $\hat{\psi}(x_t, \zeta, x_{t+p})$ be the optimal solution of SOCO. We know $\xi_t^{i_\tau - i_{\tau-1}}(x_{i_{\tau-1}}, \hat{w}_{\tau-1}, x_{i_\tau})$ is convex and L_0-strongly smooth, so we have

$$\|\psi_t^p(x,\zeta,z)_{y_h} - \psi_t^p(x',\zeta',z')_{y_h}\| = \left\|\hat{\psi}(x,\hat{w},z)_u - \hat{\psi}(x',\hat{w}',z')_u\right\|$$

$$\le C_0 \left(\lambda_0^{u-1}\|x - x'\|_2 + \sum_{\tau=0}^{v-2} \lambda_0^{|u-\tau|-1}\|\hat{w}_\tau - \hat{w}'_\tau\|_2 + \lambda_0^{(v-1)-u-1}\|z - z'\|_2\right)$$

Step 3: Deduce Perturbation Bound for MPC

Proof sketch. Now we only have to "expand" within each interval:

$$\begin{aligned}
&\left\|\psi_t^p(x,\zeta,z)_{y_h} - \psi_t^p(x',\zeta',z')_{y_h}\right\| \\
&\le C_0\left(\lambda_0^{u-1}\left\|x-x'\right\|_2 + \sum_{\tau=0}^{v-2}\lambda_0^{|u-\tau|-1}\left\|\hat{w}_\tau - \hat{w}'_\tau\right\|_2 + \lambda_0^{(v-1)-u-1}\left\|z-z'\right\|_2\right) \\
&= C_0\left(\lambda_0^{u-1}\left\|x-x'\right\|_2 + \sum_{\tau=0}^{v-2}\lambda_0^{|u-\tau|-1}\sum_{j=i_\tau}^{i_{\tau+1}-1}\left\|\zeta_j - \zeta'_j\right\|_2 + \lambda_0^{(v-1)-u-1}\left\|z-z'\right\|_2\right) \\
&\le \frac{C_0}{\lambda_0}\left(\lambda^{i_u-i_0}\left\|x-x'\right\|_2 + \sum_{\tau=0}^{v-2}\sum_{j=i_\tau}^{i_{\tau+1}-1}\lambda^{|j-i_u|}\left\|\zeta_j - \zeta'_j\right\|_2 + \lambda^{i_{v-1}-i_u}\left\|z-z'\right\|_2\right) \\
&= C\left(\lambda^h\left\|x-x'\right\| + \sum_{\tau=0}^{p-1}\lambda^{|h-\tau|}\left\|\zeta_\tau - \zeta'_\tau\right\| + \lambda^{p-h}\left\|z-z'\right\|\right).
\end{aligned}$$

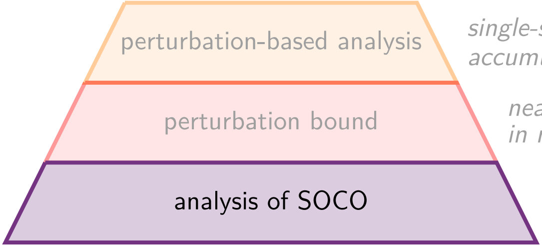

single-step errors do not accumulate over time

nearby initial states result in nearby trajectories

banded matrix inverse implies exp-decay

(Block-)Banded Matrix Inverse Implies Exponential Decay

Lemma 9 (block-banded matrix inverse)

Suppose $A \in \mathbb{S}^{\omega n}$ *is a positive definite matrix formed by* $\omega \times \omega$ *blocks* $A_{i,j} \in \mathbb{R}^{n \times n}$*, where its singular spectrum* $\sigma(A) \subseteq [a_0 . b_0]$*, and* A *is* q*-banded (i.e.,* $A_{i,j} = 0,\ \forall |i-j| > q/2$*). Suppose* $D = \operatorname{diag}(D_1, \ldots, D_\omega)$*, where* $D_i \in \mathbb{S}^n$ *is positive semi-definite, and* $M = \left((A+D)^{-1}\right)_{S_R, S_C}$*, where* $S_R, S_C \subseteq \{1, \ldots, \omega\}$*. Then we have* $\|M\| \leq C\gamma^{\hat{d}}$*, where*

$$C = \frac{2}{a_0},\ \gamma = \left(\frac{\sqrt{\operatorname{cond}(A)} - 1}{\sqrt{\operatorname{cond}(A)} + 1}\right)^{2/q},\ \hat{d} = \min_{i \in S_R, j \in S_c} |i-j|.$$

(Block-)Banded Matrix Inverse Implies Exponential Decay

Proof sketch. We first consider the case $D = 0$, as inspired by [12].

(1) $\hat{d} \neq 0$. Write $\hat{d} = \nu q/2 + \kappa$, where $\nu \geq 0, 1 \leq \kappa \leq q/2$. Then there exists a polynomial p_ν of degree ν, where

$$\left\|A^{-1} - p_\nu(A)\right\| \leq \frac{1}{a_0} \cdot \frac{\left(1 + \sqrt{\operatorname{cond}(A)}\right)^2}{2\operatorname{cond}(A)} \gamma^{\hat{d}} \leq C\gamma^{\hat{d}}.$$

Since p_v has degree $v < \frac{2\hat{d}}{q}$ and A is q-banded, the matrix $p_\nu(A)$ satisfies $(p_\nu(A))_{i,j} = 0$ for any $i \in S_R$ and $j \in S_C$. We then obtain

$$\|P\| = \left\| \left(A^{-1} - p_\nu(A)\right)_{S_R, S_C} \right\| \leq \left\|A^{-1} - p_\nu(A)\right\| \leq C\gamma^{\hat{d}}.$$

(2) $\hat{d} = 0$. Clearly $\|P\| = \left\| \left(A^{-1}\right)_{S_R, S_C} \right\| \leq \left\|A^{-1}\right\| = \frac{1}{a_0} \leq C$.

(Block-)Banded Matrix Inverse Implies Exponential Decay

Proof sketch. Then we reduce the general case to $D = 0$.

Let $N := (a_0 I + D) \in \mathbb{S}^{n\omega}$, and $H := N^{-\frac{1}{2}}(A + D)N^{-\frac{1}{2}} \in \mathbb{S}^{n\omega}$. We can show that $I \preceq H \preceq \frac{b_0}{a_0} I$ by bounding $x^\top Hx$. Since H is also q-banded and $\mathrm{cond}(H) \leq \frac{b_0}{a_0} = \mathrm{cond}(A)$, we know from the special case that

$$\left\|(H^{-1})_{S_R,S_C}\right\| \leq 2\gamma_H^{\hat{d}} \leq 2\gamma^{\hat{d}},$$

where $\gamma_H = \left(\frac{\sqrt{\mathrm{cond}(H)}-1}{\sqrt{\mathrm{cond}(H)}+1}\right)^{2/q} \leq \gamma$. Consequently,

$$\begin{aligned}
\|P\| &= \left\|\left(N^{-\frac{1}{2}} H^{-1} N^{-\frac{1}{2}}\right)_{S_R,S_C}\right\| \\
&\leq \left\|(a_0 I + D_{S_R})^{-\frac{1}{2}}\right\| \cdot \left\|(H^{-1})_{S_R,S_C}\right\| \cdot \left\|(a_0 I + D_{S_C})^{-\frac{1}{2}}\right\| \\
&\leq \frac{1}{a_0}\left\|(H^{-1})_{S_R,S_C}\right\| \leq C\gamma^{\hat{d}}.
\end{aligned}$$

Exponential Decay Property of SOCO

Proof sketch of Lemma 7. Recall that the objective of SOCO is

$$\hat{\psi}(\hat{x}_0, \hat{w}, \hat{x}_p) := \arg\min_{\hat{x}_{1:p-1}} \sum_{\tau=1}^{p-1} \hat{f}_\tau(\hat{x}_\tau) + \sum_{\tau=1}^{p} \hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1}).$$

Let $\hat{\zeta} := (\hat{x}_0, \hat{w}, \hat{x}_p)$ be the system parameters, $\hat{h}(\hat{x}_{1:p-1}, \hat{\zeta})$ be the objective, and $e := (e_0, \delta_{0:p-1}, e_p)$ be a direction of perturbation.

By the first-order criterion, for any $\theta \in \mathbb{R}$,

$$\frac{\partial}{\partial \hat{x}_{1:p-1}} \hat{h}(\hat{\psi}(\hat{\zeta} + \theta e), \hat{\zeta} + \theta e) = 0.$$

Exponential Decay Property of SOCO

Proof sketch of Lemma 7. Take derivative with respect to θ to get

$$\begin{aligned}
&\frac{\partial^2}{(\partial \hat{x}_{1:p-1})^2}\hat{h}(\hat{\psi}(\hat{\zeta}+\theta e),\hat{\zeta}+\theta e)\cdot\frac{\mathrm{d}}{\mathrm{d}\theta}\hat{\psi}(\hat{\zeta}+\theta e)\\
=&-\frac{\partial}{\partial \hat{x}_0}\frac{\partial}{\partial \hat{x}_{1:p-1}}\hat{h}(\hat{\psi}(\hat{\zeta}+\theta e),\hat{\zeta}+\theta e)\,e_0\\
&-\frac{\partial}{\partial \hat{x}_p}\frac{\partial}{\partial \hat{x}_{1:p-1}}\hat{h}(\hat{\psi}(\hat{\zeta}+\theta e),\hat{\zeta}+\theta e)\,e_p\\
&-\sum_{\tau=0}^{p-1}\frac{\partial}{\partial \hat{w}_\tau}\frac{\partial}{\partial \hat{x}_{1:p-1}}\hat{h}(\hat{\psi}(\hat{\zeta}+\theta e),\hat{\zeta}+\theta e)\delta_\tau\\
=:&\;R^{(0)}e_0+R^{(p)}e_p+\sum_{\tau=0}^{p-1}K^{(\tau)}\delta_\tau.
\end{aligned}$$

Exponential Decay Property of SOCO

Proof sketch of Lemma 7. Therefore, the derivative is given by

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\hat{\psi}(\hat{\zeta}+\theta e) = H^{-1}\left(R^{(0)}e_0 + R^{(p)}e_p + \sum_{\tau=0}^{p-1} K^{(\tau)}\delta_\tau\right).$$

where $H := \frac{\partial^2}{\partial \hat{x}_{1:p-1}^2}\hat{h}(\hat{\psi}(\hat{\zeta}+\theta e), \hat{\zeta}+\theta e)$ is the Hessian of $\hat{h}(\cdot, \hat{\zeta}+\theta e)$. Since $\hat{h}(\cdot, \hat{\zeta}+\theta e)$ only involves correlation of adjacent variables, the Hessian H is block-tridiagonal. Meanwhile, $\hat{h}$ can be decomposed as

$$\hat{h} = \underbrace{\sum_{\tau=1}^{p-1}\frac{\mu}{2}\|\hat{x}_\tau\|^2 + \sum_{\tau=1}^{p}\hat{c}_\tau(\hat{x}_\tau, \hat{x}_{\tau-1}, \hat{w}_{\tau-1})}_{\hat{h}_1(\hat{x}_{1:p-1}, \hat{\zeta})} + \underbrace{\sum_{\tau=1}^{p-1}\left(\hat{f}_\tau(\hat{x}_\tau) - \frac{\mu}{2}\|\hat{x}_\tau\|^2\right)}_{\hat{h}_2(\hat{x}_{1:p-1}, \hat{\zeta})}.$$

Exponential Decay Property of SOCO

Proof sketch of Lemma 7. Accordingly, we have $H = H_1 + H_2$, where $\mu I \preceq H_1 \preceq (\mu + 2\ell)I$ and $H_2 \succeq 0$. Hence we may apply Lemma 9 to get

$$\begin{aligned}\left\| \frac{\mathrm{d}}{\mathrm{d}\theta}\hat{\psi}(\hat{\zeta}+\theta e)_h \right\| &\leq \ell \left\|(H^{-1})_{h,1}\right\| \|e_0\| + \ell \left\|(H^{-1})_{h,p-1}\right\| \|e_p\| \\ &\quad + \ell \left\|(H^{-1})_{h,1}\right\| \|\delta_0\| + \ell \left\|(H^{-1})_{h,p-1}\right\| \|\delta_{p-1}\| \\ &\quad + \sum_{\tau=1}^{p-2} \ell \left\|(H^{-1})_{h,\tau:\tau+1}\right\| \|\delta_\tau\| \\ &\leq \cdots (\textit{plug in upper bounds}) \\ &\leq C_0 \left(\lambda_0^{h-1}\|e_0\| + \sum_{\tau=0}^{p-1} \lambda_0^{|h-\tau|-1}\|\delta_\tau\| + \lambda_0^{p-h-1}\|e_p\| \right).\end{aligned}$$

Exponential Decay Property of SOCO

Proof sketch of Lemma 7. Finally, we have

$$\begin{aligned}\left\|\hat{\psi}(\hat{\zeta})_h - \hat{\psi}(\hat{\zeta}+e)_h\right\| &= \left\|\int_0^1 \frac{\mathrm{d}}{\mathrm{d}\theta}\hat{\psi}(\hat{\zeta}+\theta e)_h \mathrm{d}\theta\right\| \\ &\le \int_0^1 \left\|\frac{\mathrm{d}}{\mathrm{d}\theta}\hat{\psi}(\hat{\zeta}+\theta e)_h\right\| \mathrm{d}\theta \\ &\le C_0 \left(\lambda_0^{h-1}\|e_0\| + \sum_{\tau=0}^{p-1} \lambda_0^{|h-\tau|-1}\|\delta_\tau\| + \lambda_0^{p-h-1}\|e_p\|\right).\end{aligned}$$

This finishes the proof.

Review of the Proof Structure

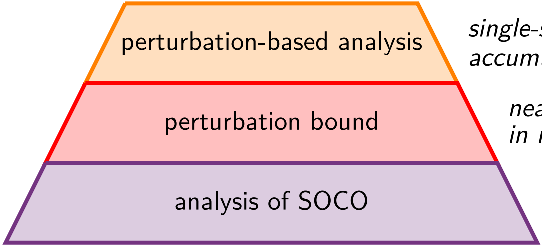

single-step errors do not accumulate over time

nearby initial states result in nearby trajectories

banded matrix inverse implies exp-decay

What to Be Done Next?

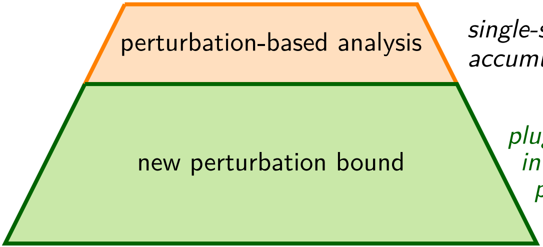

single-step errors do not accumulate over time

plug in perturbation bound in new settings to obtain performance guarantees

What to Be Done Next?

- Prediction error.
  - What if the oracle gives inexact predictions on dynamics?
- Constraint set.
  - What if there are constraints on states and control inputs?
- Other standard controllers.
  - Does the perturbation-based analysis framework help to provide guarantees for other controllers?

References I

[1] Yingying Li, Xin Chen, and Na Li. Online optimal control with linear dynamics and predictions: Algorithms and regret analysis. *Advances in Neural Information Processing Systems*, 32, 2019.

[2] Runyu Zhang, Yingying Li, and Na Li. On the regret analysis of online lqr control with predictions. *arXiv preprint arXiv: 2102.01309*, 2021.

[3] Chenkai Yu, Guanya Shi, Soon-Jo Chung, Yisong Yue, and Adam Wierman. The power of predictions in online control. *Advances in Neural Information Processing Systems*, 33, 2020.

[4] Naman Agarwal, Brian Bullins, Elad Hazan, Sham M. Kakade, and Karan Singh. Online control with adversarial disturbances. *arXiv preprint arXiv:1902.08721v1*, 2019.

[5] Naman Agarwal, Elad Hazan, and Karan Singh. Logarithmic regret for online control. *Advances in Neural Information Processing Systems*, 32, 2019.

References II

[6] Guanya Shi, Yiheng Lin, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Online optimization with memory and competitive control. In *Advances in Neural Information Processing Systems*, pages 20636–20647, 2020.

[7] Gautam Goel and Babak Hassibi. Competitive control. *arXiv preprint arXiv:2107.13657*, 2021.

[8] Chenkai Yu, Guanya Shi, Soon-Jo Chung, Yisong Yue, and Adam Wierman. Competitive control with delayed imperfect information. *arXiv preprint arXiv:2010.11637*, 2020.

[9] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*, 2017.

[10] Maryam Fazel, Rong Ge, Sham M Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. *arXiv preprint arXiv:1801.05039*, 2018.

References III

[11] Max Simchowitz and Dylan Foster. Naive exploration is optimal for online LQR. In *International Conference on Machine Learning*, pages 8937–8948. PMLR, 2020.

[12] Stephen Demko, William F Moss, and Philip W Smith. Decay rates for inverses of band matrices. *Mathematics of computation*, 43(168):491–499, 1984.

Thanks for your attention!