Bounded-Regret MPC via Perturbation Analysis: Prediction Error, Constraints, and Nonlinearity

Yiheng Lin[1], Yang Hu[2], Guannan Qu[3], Tongxin Li[4], Adam Wierman[1]

[1] Caltech CMS [2] Harvard CS [3] CMU ECE [3] CUHK Shenzhen, SDS

Model Predictive Control (MPC)

We consider an optimal online control problem in finite horizon T:

$$\min_{x_{0:T}, u_{0:T-1}} \sum_{t=0}^{T-1} f_t(x_t, u_t; \xi_t^*) + F_T(x_T; \xi_T^*) \qquad \text{(total costs)}$$
$$\text{s.t. } x_{t+1} = g_t(x_t, u_t; \xi_t^*), \quad \forall 0 \le t < T, \text{(dynamics)}$$
$$s_t(x_t, u_t; \xi_t^*) \le 0, \quad \forall 0 \le t < T, \text{(constraints)}$$
$$x_0 = x(0). \quad \text{(initial state)}$$

$x_t \in \mathbb{R}^n$: state; $u_t \in \mathbb{R}^m$: action; $\xi_t^* \in \Xi_t$: unknown ground-truth uncertainty parameter. At time t, the controller observes

$$\underbrace{x_t, \{(f_\tau, g_\tau, s_\tau)\}_{\tau=t,\dots,t+k}}_{\text{Exact predictions}}, \underbrace{\{\xi_{\tau|t}\}_{\tau=t,\dots,t+k}}_{\text{Inexact predictions}}.$$

Def. *Power of τ-step-away predictions:* $P(\tau) = \sum_{t=0}^{T-\tau} \left\| \xi_{t+\tau|t} - \xi_{t+\tau}^* \right\|^2$.

In this work, we consider the following MPC controller at time t:

Alg. *If $t < T - k$, commit $u_t \leftarrow \psi_t^{t+k}(x_t, \xi_{t:t+k-1|t}, \zeta_{t+k|t}; \mathbb{I})_{v_t}$, where $\zeta_{t+k|t} = \arg\min_x \min_u f_{t+k}(x, u; \xi_{t+k|t})$ and $\mathbb{I}$ is the indicator function. Else, commit $u_t \leftarrow \psi_t^T(x_t, \xi_{t:T|t}; F_T)$.*

Here, $\psi_{t_1}^{t_2}$ denotes the solution to the optimal control problem:

$$\psi_{t_1}^{t_2}(z, \xi_{t_1:t_2-1}, \zeta_{t_2}; F_{t_2}) = \arg\min_{y_{t_1:t_2}, v_{t_1:t_2-1}} \sum_{t=t_1}^{t_2-1} f_t(y_t, v_t; \xi_t) + F_{t_2}(y_{t_2}; \xi_{t_2})$$
$$\text{s.t. } y_{t+1} = g_t(y_t, v_t; \xi_t), \forall t_1 \le t < t_2,$$
$$s_t(y_t, v_t; \xi_t) \le 0, \forall t_1 \le t < t_2,$$
$$y_{t_1} = z.$$

Objective: Bound the dynamic regret $\mathrm{cost}(\textsc{mpc}) - \mathrm{cost}(\textsc{opt})$, where $\textsc{opt}$ denotes the offline optimal trajectory $x_{0:T}^*, u_{0:T-1}^*$.

Per-step Error & Decaying Perturbation

Per-step error is the error injected by MPC at every time step.

Def. *The per-step error e_t incurred by MPC at time t is the distance between its actual action u_t and the clairvoyant optimal action, i.e.,*

$$e_t = \left\| u_t - \psi_t^T(x_t, \xi_{t:T}^*; F_T)_{v_t} \right\|, \text{ where } u_t = \psi_t^{t+k}(x_t, \xi_{t:t+k|t}, F_{t+k})_{v_t}.$$

We need to verify two kinds of decaying perturbation bounds.

1) Perturb the uncertainty parameters given a fixed initial state, i.e.,

$$\left\| \psi_{t_1}^{t_2}(z, \xi_{t_1:t_2-1}, \zeta_{t_2}; F_{t_2})_{v_{t_1}} - \psi_{t_1}^{t_2}(z, \xi'_{t_1:t_2-1}, \zeta'_{t_2}; F_{t_2})_{v_{t_1}} \right\|$$
$$\le \sum_{t=t_1}^{t_2} q_1(t - t_1)\delta_t \cdot \|z\| + \sum_{t=t_1}^{t_2} q_2(t - t_1)\delta_t, \tag{1}$$

where $\delta_t = \|\xi_t - \xi'_t\|$ for $t = t_1, \dots, t_2 - 1$ and $\delta_{t_2} = \|\zeta_{t_2} - \zeta'_{t_2}\|$.

2) Perturb the initial state given fixed uncertainty parameters, i.e.,

$$\left\| \psi_{t_1}^{t_2}(z, \xi_{t_1:t_2-1}, \zeta_{t_2}; F_{t_2})_{y_t/v_t} - \psi_{t_1}^{t_2}(z', \xi_{t_1:t_2-1}, \zeta_{t_2}; F_{t_2})_{y_t/v_t} \right\|$$
$$\le q_3(t - t_1) \|z - z'\|, \text{ for } t \in [t_1, t_2]. \tag{2}$$

Prop 1. *For $R > 0$ and $\sum_{t=0}^{\infty} q_i(t) \le C_i$ for $i = 1, 2, 3$, $C_3 \ge 1$,* (1) *holds with $t_1 = t$ and $t_2 = t + k$ for $t < T - k$, $F_{t+k} = \mathbb{I}$*

$$z \in \mathcal{B}(x_t^*, R);\ \xi'_{t:t+k-1} = \xi^*_{t:t+k-1};\ \xi_{t+k}, \xi'_{t+k} \in \mathcal{B}(x^*_{t+k}, R) \subseteq \mathbb{R}^n;$$

with $t_1 = t$ and $t_2 = T$ for $t \ge T - k$, (1) *holds for $z \in \mathcal{B}(x_t^*, R)$; $\xi_{t:T} = \xi^*_{t:T}$; $F = F_T$. Further,* (2) *holds for any $z, z' \in \mathcal{B}(x_t^*, R)$ and $\xi_{t_1:t_2} = \xi^*_{t_1:t_2}$.*

We propose a general pipeline to reduce the derivation of **dynamic regret** for Model Predictive Control in nonlinear time-varying systems to deriving **decaying perturbation bounds** of optimal trajectories.

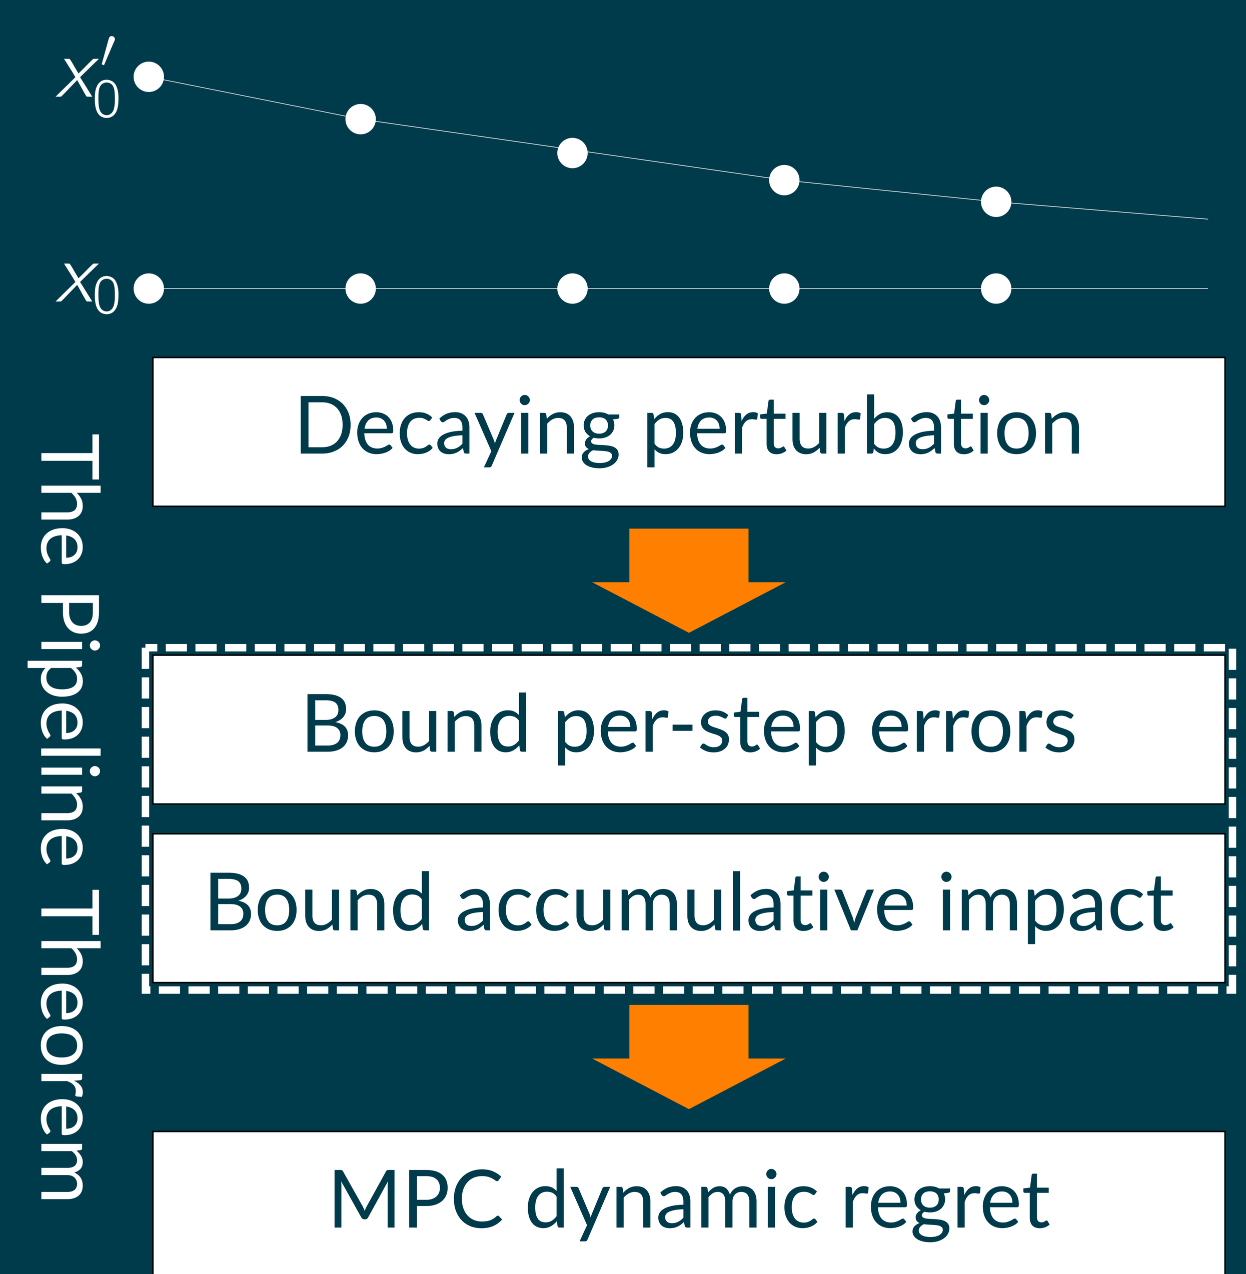

Take a picture to download the full paper

The Pipeline Theorem

We need the following assumptions for the pipeline theorem to hold.

- *Stability of* $\textsc{opt}$: $\exists D_{x^*} > 0$, s.t. $\|x_t^*\| \le D_{x^*}$.
- *Lipschitz dynamics:* $\|g_t(x_t, u_t; \xi_t^*) - g_t(x_t, u'_t; \xi_t^*)\| \le L_g \|u_t - u'_t\|$.
- *Well-conditioned costs:* every stage cost $f_t(\cdot, \cdot; \xi_t^*)$ and the terminal cost $F_T(\cdot; \xi_T^*)$ are nonnegative, convex, and ℓ-smooth.

Bound the per-step errors: Given Property 1 holds, we have

Lemma 1. *Let Property 1 hold. Suppose $x_t \in \mathcal{B}(x_t^*, R/C_3)$ and $\zeta_{t+k|t} \in \mathcal{B}(x^*_{t+k}, R)$ for $t < T - k$. Then, the per-step error of MPC is bounded by*

$$e_t \le \sum_{\tau=0}^{k} ((R/C_3 + D_{x^*}) \cdot q_1(\tau) + q_2(\tau)) \rho_{t,\tau}$$
$$+ 2R ((R/C_3 + D_{x^*}) \cdot q_1(k) + q_2(k)). \tag{3}$$

Bound the accumulative impact: Once Property 1 holds and the per-step errors are sufficiently small, the trajectory of MPC "remains close" to OPT and the dynamic regret can be bounded.

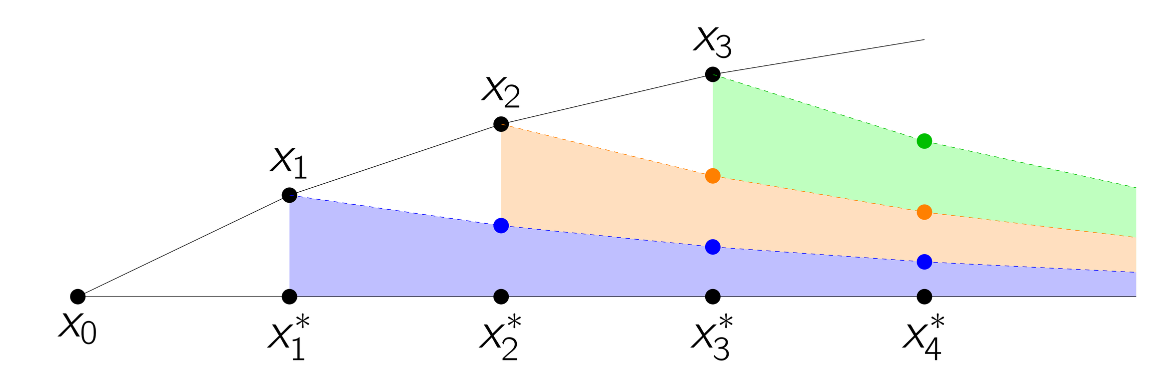

Lemma 2. *Let Property 1 hold. If $e_\tau \le R/(C_3^2 L_g)$ for all $\tau < t$, then $x_t \in \mathcal{B}(x_t^*, R/C_3)$ for MPC and its dynamic regret is upper bounded by*

$$\mathrm{cost}(\textsc{mpc}) - \mathrm{cost}(\textsc{opt}) = O\left(\sqrt{\mathrm{cost}(\textsc{opt}) \cdot \sum_{t=0}^{T-1} e_t^2} + \sum_{t=0}^{T-1} e_t^2 \right).$$

Combining Lemma 1 and 2 gives **the Pipeline Theorem**:

Thm 1. *Let Property 1 hold. Suppose $x_t \in \mathcal{B}(x_t^*, R/C_3)$ and $\zeta_{t+k|t} \in \mathcal{B}(x^*_{t+k}, R)$ for $t < T - k$. If k and $\rho_{t,\tau}$ satisfy that (3) $\le R/(C_3^2 L_g)$, then*

$$\mathrm{cost}(\textsc{mpc}) - \mathrm{cost}(\textsc{opt}) = O\left(\sqrt{\mathrm{cost}(\textsc{opt}) \cdot (E_1 + E_2)} + (E_1 + E_2) \right),$$

where $E_1 = \sum_{\tau=0}^{k-1} (q_1(\tau) + q_2(\tau)) P(\tau)$ and $E_2 = (q_1(k)^2 + q_2(k)^2)\, T$.

The dynamic regret depends on both $\{P(\tau)\}_{\tau \ge 0}$ and k. We see that

1. No need to predict far future accurately;
2. As k increases, $E_1 \nearrow$, $E_2 \searrow$. The trade-off motivates the problem of MPC horizon selection in our more recent work.

Apply to General Systems

The pipeline theorem reduces the MPC dynamic regret problem to verifying Prop 1, which is challenging in general and requires:

- All cost/dynamics/constraint functions are in $\mathcal{C}^2$;
- Strong second order sufficient condition (SSOSC) holds;
- Linear Independence Constraint Quantification (LICQ) holds;
- Uniform singular spectrum bound for reduced Hessian.

We provide several positive and negative examples for Prop 1:

Example. ***Positive examples:***

- *General costs, $g_t(x_t, u_t; \xi_t) = A_t x_t + B_t u_t + w_t(\xi_t)$, and unconstrained.*
- *$f_t(x_t, u_t; \xi_t) = (x_t - \bar{x}_t(\xi_t))^\top Q_t(\xi_t)(x_t - \bar{x}_t(\xi_t)) + u_t^\top R_t(\xi_t) u_t$, $g_t(x_t, u_t; \xi_t) = A_t(\xi_t) \cdot x_t + B_t(\xi_t) \cdot u_t + w_t(\xi_t)$, and unconstrained.*
- *$n = m = 1$, general costs, $x_{t+1} = x_t + u_t$, $x_t \in [-1, 1]$, $u_t \ge -0.8$.*

Negative example:

- *$n = m = 1$, general costs, $x_{t+1} = x_t + u_t$, $x_t \in [-1, 1]$, $u_t \in [-0.8, 0.8]$.*