Data-Enabled Predictive Control of Autonomous Energy Systems

Florian Dörfler

Automatic Control Laboratory, ETH Zürich

Acknowledgements

Jeremy Coulson

Linbin Huang

Paul Beuchat

John Lygeros

Ivan Markovsky

Ezzat Elokda

Perspectives on model-based control

Single system level:

- modeling & system ID are very expensive
- models not always useful for control
- need for end-to-end automation solutions

From experiment design to closed-loop control[☆]

Håkan Hjalmarsson*

1. Introduction

Ever increasing productivity demands and environmental standards necessitate more and more advanced control methods to be employed in industry. However, such methods usually require a model of the process and modeling and system identification are expensive. Quoting (Ogunnaike, 1996):

"*It is also widely recognized, however, that obtaining the process model is the single most time consuming task in the application of model-based control.*"

In Hussain (1999) it is reported that three quarters of the total costs associated with advanced control projects can be attributed to modeling. It is estimated that models exist for far less than one percent of all processes in regulatory control. According to Desborough and Miller (2001), one of the few instances when the cost of dynamic modeling can be justified is for the commissioning of model predictive controllers.

It has also been recognized that models for control pose special considerations. Again quoting (Ogunnaike, 1996):

"*There is abundant evidence in industrial practice that when modeling for control is not based on criteria related to the actual end use, the results can sometimes be quite disappointing.*"

Hence, efficient modeling and system identification techniques suited for industrial use and tailored for control design applications have become important enablers for industrial advances. The Panel for Future Directions in Control, (Murray, Åström, Boyd, Brockett, & Stein, 2003), has identified *automatic synthesis of control algorithms, with integrated validation and verification* as one of the major future challenges in control. Quoting (Murray et al., 2003):

"*Researchers need to develop much more powerful design tools that automate the entire control design process from*

Critical infrastructure level: (especially in energy)

- subsystem (device) models & controls are proprietary
- infrastructure (network) owned by many entities/countries
- operating points/modes are in flux & constantly changing

nobody has any dynamic models . . .

Control in a data-rich world

- ever-growing trend in CS & applications: ***data-driven control*** by-passing models
- canonical problem: ***black/gray-box system control*** based on I/O samples

Q: Why give up physical modeling and reliable model-based algorithms ?

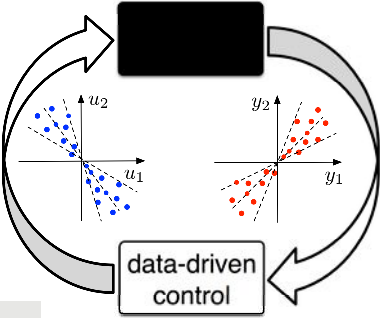

Data-driven control is ***viable alternative*** when

- models are too complex to be useful (e.g., fluids, wind farms, & building automation)
- first-principle models are not conceivable (e.g., human-in-the-loop, biology, & perception)
- modeling & system ID is too cumbersome (e.g., robotics & electronics applications)

Central promise: *It is often easier to learn control policies directly from data, rather than learning a model.*

Example: PID

Snippets from the literature

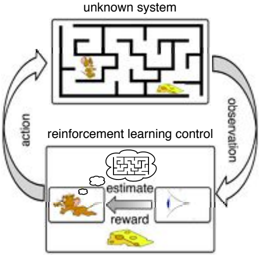

1. ***reinforcement learning*** / stochastic adaptive control / dual control / approximate dynamic programming

ø not suitable for physical, real-time, & safety-critical

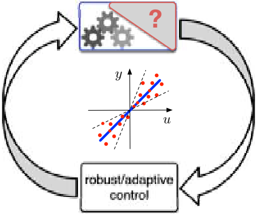

2. gray-box ***safe learning & control*** (adaptive)

ø limited applicability: need a-priori safety

3. sequential ***system ID + UQ + control***

→ recent finite-sample & end-to-end ID + UQ + control pipelines out-performing RL

ø ID seeks best but not most useful model

→ "easier to learn policies than models"

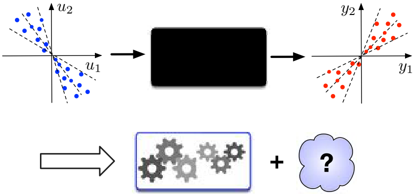

Colorful idea

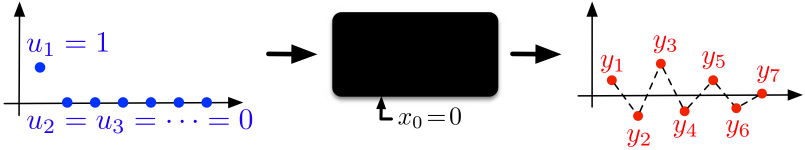

If you had the ***impulse response*** of a LTI system, then …

- can ***identify model*** (e.g., transfer function or Kalman-Ho realization)
- …but can also build ***predictive model directly from raw data*** :

$$y_{\text{future}}(t) = \begin{bmatrix} y_1 & y_2 & y_3 & \cdots \end{bmatrix} \cdot \begin{bmatrix} u_{\text{future}}(t) \\ u_{\text{future}}(t-1) \\ u_{\text{future}}(t-2) \\ \vdots \end{bmatrix}$$

- ***model predictive control*** from data: dynamic matrix control (DMC)
- ***today:*** can we do so with arbitrary, finite, and corrupted I/O samples ?

Contents

I. Data-Enabled Predictive Control (DeePC): Basic Idea

J. Coulson, J. Lygeros, and F. Dörfler. *Data-Enabled Predictive Control: In the Shallows of the DeePC*. arxiv.org/abs/1811.05890.

II. From Heuristics & Numerical Promises to Theorems

J. Coulson, J. Lygeros, and F. Dörfler. *Regularized and Distributionally Robust Data-Enabled Predictive Control*. arxiv.org/abs/1903.06804.

III. Application: End-to-End Automation in Energy Systems

L. Huang, J. Coulson, J. Lygeros, and F. Dörfler. *Data-Enabled Predictive Control for Grid-Connected Power Converters*. arxiv.org/abs/1903.07339.

Preview

complex 4-area power ***system***:

large ($n = 208$), few sensors (8), nonlinear, noisy, stiff, input constraints, & decentralized control specifications

control objective: damping of inter-area oscillations via HVDC link but $\boxed{\text{without model}}$

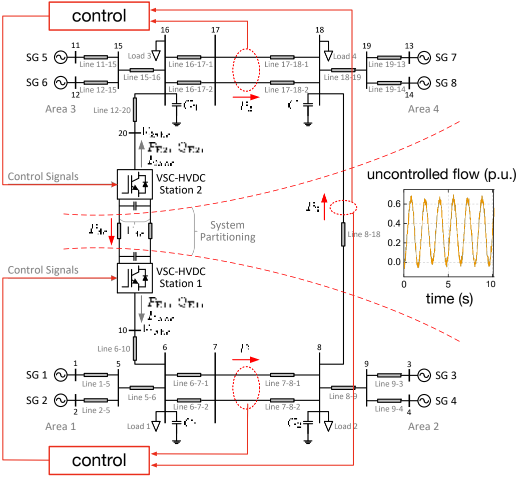

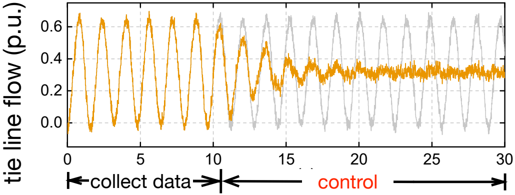

seek a method that ***works reliably***, can be ***efficiently*** implemented, & ***certifiable***

$\rightarrow$ automating ourselves

Behavioral view on LTI systems

Definition: A discrete-time ***dynamical system*** is a 3-tuple $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathscr{B})$ where

(i) $\mathbb{Z}_{\geq 0}$ is the discrete-time axis,

(ii) $\mathbb{W}$ is a signal space, and

(iii) $\mathscr{B} \subseteq \mathbb{W}^{\mathbb{Z}_{\geq 0}}$ is the behavior.

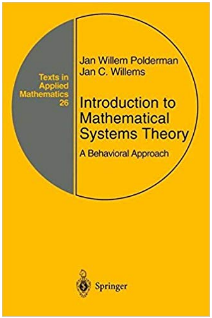

Definition: The dynamical system $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathscr{B})$ is

(i) ***linear*** if $\mathbb{W}$ is a vector space & $\mathscr{B}$ is a subspace of $\mathbb{W}^{\mathbb{Z}_{\geq 0}}$

(ii) and ***time-invariant*** if $\mathscr{B} \subseteq \sigma\mathscr{B}$, where $\sigma w_t = w_{t+1}$.

$\mathscr{B} =$ ***set of trajectories*** & $\mathscr{B}_T$ is ***restriction*** to $t \in [0, T]$

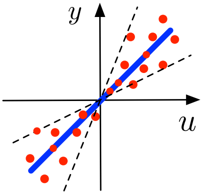

LTI systems and matrix time series

foundation of state-space subspace system ID & signal recovery algorithms

$\big(u(t), y(t)\big)$ satisfy recursive ***difference equation***

$$b_0 u_t + b_1 u_{t+1} + \ldots + b_n u_{t+n} + a_0 y_t + a_1 y_{t+1} + \ldots + a_n y_{t+n} = 0$$

(ARMA / kernel representation)

$$\Longrightarrow$$

$\Longleftarrow$

under assumptions

$\begin{bmatrix} b_0 & a_0 & b_1 & a_1 & \ldots & b_n & a_n \end{bmatrix}$ spans left nullspace of ***Hankel matrix*** (collected from data)

$$\mathscr{H}_L \begin{pmatrix} u \\ y \end{pmatrix} = \begin{bmatrix} \binom{u_1}{y_1} & \binom{u_2}{y_2} & \binom{u_3}{y_3} & \cdots & \binom{u_{T-L+1}}{y_{T-L+1}} \\ \binom{u_2}{y_2} & \binom{u_3}{y_3} & \binom{u_4}{y_4} & \cdots & \vdots \\ \binom{u_3}{y_3} & \binom{u_4}{y_4} & \binom{u_5}{y_5} & \cdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \binom{u_L}{y_L} & \cdots & \cdots & \cdots & \binom{u_T}{y_T} \end{bmatrix}$$

The Fundamental Lemma

Definition : The signal $u = \text{col}(u_1, \ldots, u_T) \in \mathbb{R}^{mT}$ is ***persistently exciting of order*** L ***if*** $\mathscr{H}_L(u) = \begin{bmatrix} u_1 & \cdots & u_{T-L+1} \\ \vdots & \ddots & \vdots \\ u_L & \cdots & u_T \end{bmatrix}$ is of full row rank,

i.e., if the signal is ***sufficiently rich*** *and* ***long*** $(T - L + 1 \geq mL)$.

Fundamental Lemma [Willems et al, '05] : Let $T, t \in \mathbb{Z}_{>0}$, Consider

- a *controllable* LTI system $(\mathbb{Z}_{\geq 0}, \mathbb{R}^{m+p}, \mathscr{B})$, and
- a T-sample long *trajectory* col$(u, y) \in \mathscr{B}_T$, where
- u is *persistently exciting* of order $t + n$ (prediction span + # states).

Then $\boxed{\text{colspan}\left(\mathscr{H}_t\left(\begin{smallmatrix} u \\ y \end{smallmatrix}\right)\right) = \mathscr{B}_t}$.

Cartoon of Fundamental Lemma

persistently exciting

controllable LTI

sufficiently many samples

$$
\underbrace{\begin{aligned} x_{k+1} &= A x_k + B u_k \\ y_k &= C x_k + D u_k \end{aligned}}_{\text{parametric state-space model}}
\quad \Longleftrightarrow \quad
\underbrace{\text{colspan} \begin{bmatrix} \begin{pmatrix} u_1 \\ y_1 \end{pmatrix} & \begin{pmatrix} u_2 \\ y_2 \end{pmatrix} & \begin{pmatrix} u_3 \\ y_3 \end{pmatrix} & \cdots \\ \begin{pmatrix} u_2 \\ y_2 \end{pmatrix} & \begin{pmatrix} u_3 \\ y_3 \end{pmatrix} & \begin{pmatrix} u_4 \\ y_4 \end{pmatrix} & \cdots \\ \begin{pmatrix} u_3 \\ y_3 \end{pmatrix} & \begin{pmatrix} u_4 \\ y_4 \end{pmatrix} & \begin{pmatrix} u_5 \\ y_5 \end{pmatrix} & \cdots \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix}}_{\text{non-parametric model from raw data}}
$$

all trajectories constructible from finitely many previous trajectories

Data-driven simulation [Markovsky & Rapisarda '08]

Problem: predict future output $y \in \mathbb{R}^{p \cdot T_{\text{future}}}$ based on

- input signal $u \in \mathbb{R}^{m \cdot T_{\text{future}}}$ $\rightarrow$ to predict forward
- past data $\text{col}(u^{\text{d}}, y^{\text{d}}) \in \mathscr{B}_{T_{\text{data}}}$ $\rightarrow$ to form Hankel matrix

Assume: $\mathscr{B}$ controllable & u^{d} persistently exciting of order $T_{\text{future}} + n$

Solution: given $(u_1, \ldots, u_{T_{\text{future}}}) \rightarrow$ compute $\boldsymbol{g}$ & $(y_1, \ldots, y_{T_{\text{future}}})$ from

$$\left[\begin{array}{cccc} u^{\text{d}}_1 & u^{\text{d}}_2 & \cdots & u^{\text{d}}_{T-N+1} \\ \vdots & \vdots & \ddots & \vdots \\ u^{\text{d}}_{T_{\text{future}}} & u^{\text{d}}_{T_{\text{future}}+1} & \cdots & u^{\text{d}}_{T} \\ \hline y^{\text{d}}_1 & y^{\text{d}}_2 & \cdots & y^{\text{d}}_{T-N+1} \\ \vdots & \vdots & \ddots & \vdots \\ y^{\text{d}}_{T_{\text{future}}} & y^{\text{d}}_{T_{\text{future}}+1} & \cdots & y^{\text{d}}_{T} \end{array}\right] \boldsymbol{g} = \left[\begin{array}{c} u_1 \\ \vdots \\ u_{T_{\text{future}}} \\ \hline y_1 \\ \vdots \\ y_{T_{\text{future}}} \end{array}\right]$$

Issue: predicted output is not unique $\rightarrow$ need to set initial conditions!

Refined problem: predict future output $y \in \mathbb{R}^{p \cdot T_{\text{future}}}$ based on

- initial trajectory $\text{col}(u_{\text{ini}}, y_{\text{ini}}) \in \mathbb{R}^{(m+p)T_{\text{ini}}}$ $\rightarrow$ to estimate initial x_{ini}
- input signal $u \in \mathbb{R}^{m \cdot T_{\text{future}}}$ $\rightarrow$ to predict forward
- past data $\text{col}(u^{\text{d}}, y^{\text{d}}) \in \mathscr{B}_{T_{\text{data}}}$ $\rightarrow$ to form Hankel matrix

Assume: $\mathscr{B}$ controllable & u^{d} persist. exciting of order $T_{\text{ini}} + T_{\text{future}} + n$

Solution: given $(u_1, \dots, u_{T_{\text{future}}})$ & $\text{col}(u_{\text{ini}}, y_{\text{ini}})$

$\rightarrow$ compute $\boldsymbol{g}$ & $(y_1, \dots, y_{T_{\text{future}}})$ from

$$\begin{bmatrix} U_{\text{p}} \\ Y_{\text{p}} \\ U_{\text{f}} \\ Y_{\text{f}} \end{bmatrix} \boldsymbol{g} = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix}$$

$\Rightarrow$ if $T_{\text{ini}} \geq$ lag of system, then y is unique

$$\begin{bmatrix} U_{\text{p}} \\ U_{\text{f}} \end{bmatrix} \triangleq \begin{bmatrix} u_1^{\text{d}} & \cdots & u_{T-T_{\text{future}}-T_{\text{ini}}+1}^{\text{d}} \\ \vdots & \ddots & \vdots \\ u_{T_{\text{ini}}}^{\text{d}} & \cdots & u_{T-T_{\text{future}}}^{\text{d}} \\ u_{T_{\text{ini}}+1}^{\text{d}} & \cdots & u_{T-T_{\text{future}}+1}^{\text{d}} \\ \vdots & \ddots & \vdots \\ u_{T_{\text{ini}}+T_{\text{future}}}^{\text{d}} & \cdots & u_T^{\text{d}} \end{bmatrix} \qquad \begin{bmatrix} Y_{\text{p}} \\ Y_{\text{f}} \end{bmatrix} \triangleq \begin{bmatrix} y_1^{\text{d}} & \cdots & y_{T-T_{\text{future}}-T_{\text{ini}}+1}^{\text{d}} \\ \vdots & \ddots & \vdots \\ y_{T_{\text{ini}}}^{\text{d}} & \cdots & y_{T-T_{\text{future}}}^{\text{d}} \\ y_{T_{\text{ini}}+1}^{\text{d}} & \cdots & y_{T-T_{\text{future}}+1}^{\text{d}} \\ \vdots & \ddots & \vdots \\ y_{T_{\text{ini}}+T_{\text{future}}}^{\text{d}} & \cdots & y_T^{\text{d}} \end{bmatrix}$$

Control from Hankel matrix data

A note on persistency of excitation

Jan C. Willems[a], Paolo Rapisarda[b], Ivan Markovsky[a,*], Bart L.M. De Moor[a]

[a]ESAT, SCD/SISTA, K.U. Leuven, Kasteelpark Arenberg 10, B 3001 Leuven, Heverlee, Belgium
[b]Department of Mathematics, University of Maastricht, 6200 MD Maastricht, The Netherlands

Received 3 June 2004; accepted 7 September 2004
Available online 30 November 2004

We are all writing merely the dramatic corollaries . . .

implicit (computational)

$\rightarrow$ Ivan Markovsky & ourselves

explicit (control policy)

$\rightarrow$ **Claudio de Persis** & Pietro Tesi

recently gaining lots of momentum with contributions by
C. Scherer, F. Allgöwer, K. Camlibel, H. Trentelman, . . .

Output **M**odel **P**redictive **C**ontrol

The canonical receding-horizon ***MPC optimization problem***:

$$
\begin{aligned}
\underset{u,x,y}{\text{minimize}} \quad & \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 \\
\text{subject to} \quad & x_{k+1} = Ax_k + Bu_k, \; \forall k \in \{0, \dots, T_{\text{future}} - 1\}, \\
& y_k = Cx_k + Du_k, \quad \forall k \in \{0, \dots, T_{\text{future}} - 1\}, \\
& x_{k+1} = Ax_k + Bu_k, \; \forall k \in \{-T_{\text{ini}} - 1, \dots, -1\}, \\
& y_k = Cx_k + Du_k, \quad \forall k \in \{-T_{\text{ini}} - 1, \dots, -1\}, \\
& u_k \in \mathcal{U}, \quad \forall k \in \{0, \dots, T_{\text{future}} - 1\}, \\
& y_k \in \mathcal{Y}, \quad \forall k \in \{0, \dots, T_{\text{future}} - 1\}
\end{aligned}
$$

quadratic cost with $R \succ 0, Q \succeq 0$ & ref. r

model for **prediction** over $k \in [0, T_{\text{future}} - 1]$

model for **estimation** (many variations)

hard operational or safety **constraints**

For a deterministic LTI plant and an exact model of the plant,
MPC is the ***gold standard of control***: safe, optimal, tracking, ...

Data-Enabled Predictive Control

DeePC uses non-parametric and data-based Hankel matrix time series as prediction/estimation model inside MPC optimization problem:

$$
\begin{aligned}
\underset{g,u,y}{\text{minimize}} \quad & \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 \\
\text{subject to} \quad & \begin{bmatrix} U_{\mathrm{p}} \\ Y_{\mathrm{p}} \\ U_{\mathrm{f}} \\ Y_{\mathrm{f}} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix}, \\
& u_k \in \mathcal{U}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}, \\
& y_k \in \mathcal{Y}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}
\end{aligned}
$$

quadratic cost with $R \succ 0, Q \succeq 0$ & ref. r

non-parametric model for **prediction** and **estimation**

hard operational or safety **constraints**

- Hankel matrix with $T_{\text{ini}} + T_{\text{future}}$ rows from past data $\begin{bmatrix} U_{\mathrm{p}} \\ U_{\mathrm{f}} \end{bmatrix} = \mathscr{H}_{T_{\text{ini}}+T_{\text{future}}}(u^{\mathrm{d}})$ and $\begin{bmatrix} Y_{\mathrm{p}} \\ Y_{\mathrm{f}} \end{bmatrix} = \mathscr{H}_{T_{\text{ini}}+T_{\text{future}}}(y^{\mathrm{d}})$

collected **offline** (could be adapted online)

- past $T_{\text{ini}} \geq$ lag samples $(u_{\text{ini}}, y_{\text{ini}})$ for x_{ini} estimation

updated **online**

Correctness for LTI Systems

Theorem: Consider a ***controllable LTI system*** and the DeePC & MPC optimization problems with ***persistently exciting*** data of order $T_{\text{ini}}+T_{\text{future}}+n$. Then the ***feasible sets of DeePC & MPC coincide***.

Corollary: If $\mathcal{U}, \mathcal{Y}$ are ***convex***, then also the ***trajectories coincide***.

Aerial robotics case study :

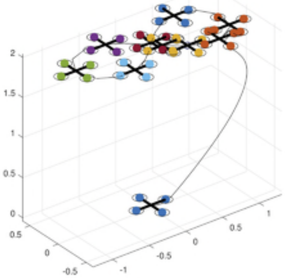

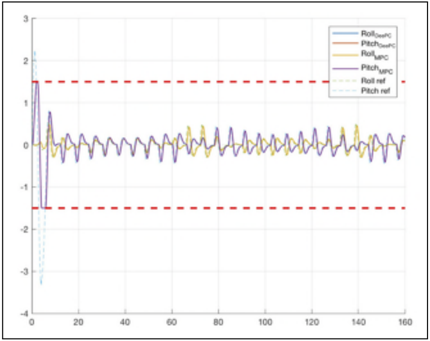

Thus, ***MPC carries over to DeePC***

…at least in the ***nominal case***.

(see e.g. [Berberich, Köhler, Müller, & Allgöwer '19] for stability proofs)

Beyond LTI, what about measurement noise, corrupted past data, and nonlinearities ?

Noisy real-time measurements

$$\underset{g,u,y}{\text{minimize}} \quad \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 + \lambda_y \|\sigma_y\|_1$$

$$\text{subject to} \quad \begin{bmatrix} U_{\text{p}} \\ Y_{\text{p}} \\ U_{\text{f}} \\ Y_{\text{f}} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_y \\ 0 \\ 0 \end{bmatrix},$$

$$u_k \in \mathcal{U}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\},$$

$$y_k \in \mathcal{Y}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}$$

Solution : add **slack** to ensure feasibility with ℓ_1**-penalty**

$\Rightarrow$ for λ_y sufficiently large $\sigma_y \neq 0$ only if constraint infeasible

c.f. ***sensitivity analysis*** over randomized sims

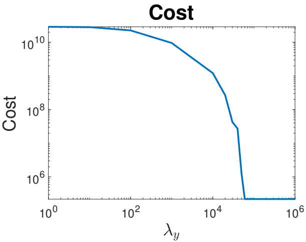

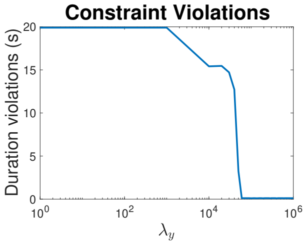

Hankel matrix corrupted by noise

$$
\begin{aligned}
\underset{g,u,y}{\text{minimize}} \quad & \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 + \lambda_g \|g\|_1 \\
\text{subject to} \quad & \begin{bmatrix} U_{\mathrm{p}} \\ Y_{\mathrm{p}} \\ U_{\mathrm{f}} \\ Y_{\mathrm{f}} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix}, \\
& u_k \in \mathcal{U}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}, \\
& y_k \in \mathcal{Y}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}
\end{aligned}
$$

Solution : add a ℓ_1**-penalty** on g

intuition: ℓ_1 sparsely selects {Hankel matrix columns} = {past trajectories} = {motion primitives}

c.f. ***sensitivity analysis*** over randomized sims

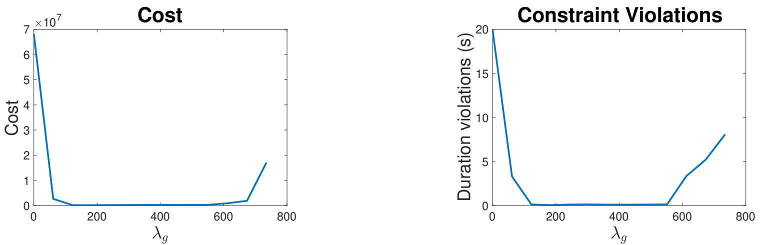

Towards nonlinear systems ...

Idea : lift nonlinear system to large/∞-dimensional bi-/linear system
$\rightarrow$ Carleman, Volterra, Fliess, Koopman, Sturm-Liouville methods
$\rightarrow$ nonlinear dynamics can be approximated LTI on finite horizons

$\rightarrow$ ***exploit size rather than nonlinearity*** and find features in data
$\rightarrow$ ***regularization*** singles out relevant features / basis functions

case study : regularization for g and σ_y

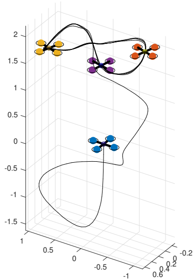

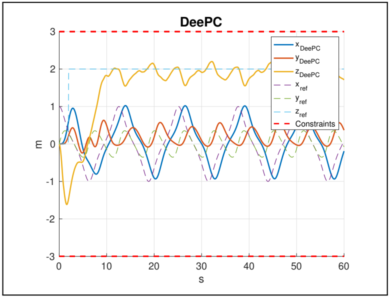

Experimental snippet

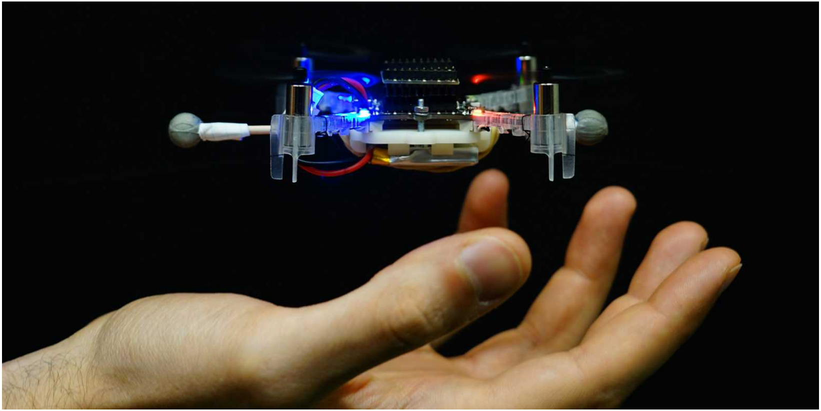

recall the ***central promise***:

it is easier to learn control policies directly from data, rather than learning a model

Comparison to system ID + MPC

Setup: nonlinear stochastic quadcopter model with full state info

DeePC + ℓ_1-regularization for g and σ_y

MPC: system ID via prediction error method + nominal MPC

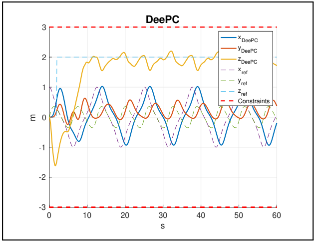

single fig-8 run

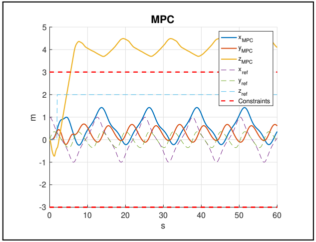

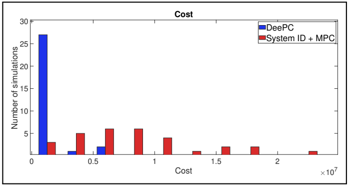

random sims

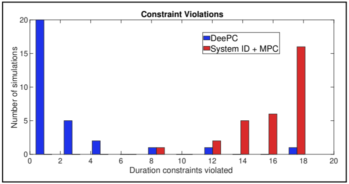

from heuristics &
numerical promises
to ***theorems***

Robust problem formulation

1. the ***nominal problem*** (without g-regularization)

$$\begin{aligned} \underset{g,u,y}{\text{minimize}} \quad & \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 + \lambda_y \|\sigma_y\|_1 \\ \text{subject to} \quad & \begin{bmatrix} U_{\mathrm{p}} \\ \widehat{Y_{\mathrm{p}}} \\ U_{\mathrm{f}} \\ \widehat{Y_{\mathrm{f}}} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ \widehat{y_{\text{ini}}} \\ u \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_y \\ 0 \\ 0 \end{bmatrix}, \\ & u_k \in \mathcal{U}, \quad \forall k \in \{0, \dots, T_{\text{future}} - 1\} \end{aligned}$$

where $\widehat{\cdot}$ denotes *measured* & thus possibly corrupted data

2. ***abstraction*** of the problem after eliminating (u, y, σ_y): $\underset{g \in G}{\text{minimize}} \;\; c\left(\widehat{\xi}, g\right)$

with samples $\widehat{\xi} = \left(\widehat{Y_{\mathrm{p}}}, \widehat{Y_{\mathrm{f}}}, \widehat{y_{\text{ini}}}\right)$ & $G = \{g : U_{\mathrm{p}} g = u_{\text{ini}} \;\&\; U_{\mathrm{f}} g \in \mathcal{U}\}$

3. a ***further abstraction***

$$\underset{g \in G}{\text{minimize}} \; c\left(\widehat{\xi}, g\right) = \underset{g \in G}{\text{minimize}} \; \mathbb{E}_{\widehat{\mathbb{P}}}\left[c\left(\xi, g\right)\right]$$

where $\widehat{\mathbb{P}} = \delta_{\widehat{\xi}}$ denotes the *empirical distribution* from which we obtained $\widehat{\xi}$

$\Rightarrow$ ***poor out-of-sample performance*** of above sample-average solution $g^\star$

for ***real problem***: $\mathbb{E}_{\mathbb{P}}\left[c\left(\xi, g^\star\right)\right]$ where $\mathbb{P}$ is the *unknown* distribution of ξ

4. ***distributionally robust*** formulation:

$$\inf_{g \in G} \; \sup_{Q \in \mathbb{B}_\epsilon(\widehat{P})} \; \mathbb{E}_Q\left[c\left(\xi, g\right)\right]$$

where the *ambiguity set* $\mathbb{B}_\epsilon(\widehat{P})$ is an ***ϵ-Wasserstein ball centered at $\widehat{P}$***:

$$\mathbb{B}_\epsilon(\widehat{P}) = \left\{ P \,:\, \inf_{\Pi} \int \left\|\xi - \hat{\xi}\right\|_W d\Pi \leq \epsilon \right\}$$

where Π has marginals $\hat{P}$ and P

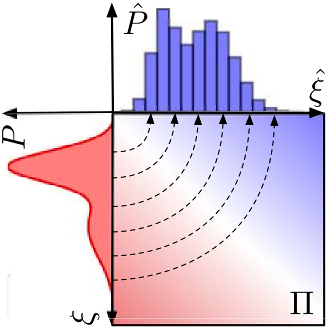

note: Wasserstein ball does not only include LTI systems with additive Gaussian noise but “everything” (integrable)

4. ***distributionally robust*** formulation $\inf_{g \in G} \sup_{Q \in \mathbb{B}_\epsilon(\widehat{P})} \mathbb{E}_Q \left[c\left(\xi, g\right) \right]$

where the *ambiguity set* $\mathbb{B}_\epsilon(\widehat{P})$ is an ***ϵ-Wasserstein ball centered at $\widehat{P}$***:

$$\mathbb{B}_\epsilon(\widehat{P}) = \left\{ P \,:\, \inf_{\Pi} \int \left\| \xi - \hat{\xi} \right\|_W d\Pi \,\leq\, \epsilon \right\}$$ where Π has marginals $\hat{P}$ and P

Theorem: Under minor technical conditions:

$$\inf_{g \in G} \sup_{Q \in \mathbb{B}_\epsilon(\widehat{P})} \mathbb{E}_Q \left[c\left(\xi, g\right) \right] \;\equiv\; \min_{g \in G} c\left(\widehat{\xi}, g\right) + \epsilon\, \mathrm{Lip}(c) \cdot \|g\|_W^\star$$

Cor: ℓ_∞-robustness in trajectory space $\Leftrightarrow$ ℓ_1-regularization of DeePC

Proof uses methods by Kuhn & Esfahani:
semi-infinite problem becomes finite after marginalization & for discrete worst case

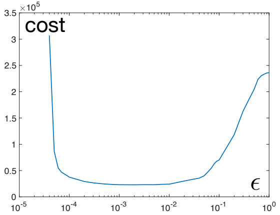

Explicit relation to system ID & MPC

1. ***regularized DeePC*** problem

$$\begin{array}{ll} \underset{g,\, u \in \mathcal{U},\, y \in \mathcal{Y}}{\text{minimize}} & f(u,y) + \lambda_g \|g\|_2^2 \\ \text{subject to} & \begin{bmatrix} U_\mathrm{p} \\ Y_\mathrm{p} \\ U_\mathrm{f} \\ Y_\mathrm{f} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix} \end{array}$$

2. standard model-based ***MPC*** (ARMA parameterization)

$$\begin{array}{ll} \underset{u \in \mathcal{U},\, y \in \mathcal{Y}}{\text{minimize}} & f(u,y) \\ \text{subject to} & y = K \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix} \end{array}$$

3. ***subspace ID*** $\boxed{y = Y_\mathrm{f}\, g^\star}$

where $g^\star = g^\star(u_{\text{ini}}, y_{\text{ini}}, u)$ solves

$$\begin{array}{ll} \underset{g}{\arg\min} & \|g\|_2^2 \\ \text{subject to} & \begin{bmatrix} U_\mathrm{p} \\ Y_\mathrm{p} \\ U_\mathrm{f} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix} \end{array}$$

4. equivalent ***prediction error ID***

$$\underset{K}{\text{minimize}} \quad \sum_j \left\| y_j^{\text{d}} - K \begin{bmatrix} u_{\text{ini}}{}_j^{\text{d}} \\ y_{\text{ini}}{}_j^{\text{d}} \\ u_j^{\text{d}} \end{bmatrix} \right\|^2$$

$$\rightarrow \quad \boxed{y = K \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix} = Y_\mathrm{f}\, g^\star}$$

subsequent ***ID & MPC***

$$\begin{aligned} \underset{u \in \mathcal{U}, y \in \mathcal{Y}}{\text{minimize}} \quad & f(u, y) \\ \text{subject to} \quad & y = K \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix} \end{aligned}$$

where K solves

$$\underset{K}{\arg\min} \quad \sum_j \left\| y_j - K \begin{bmatrix} u_{\text{ini}\,j} \\ y_{\text{ini}\,j} \\ u_j \end{bmatrix} \right\|^2$$

$$\equiv$$

$$\begin{aligned} \underset{u \in \mathcal{U}, y \in \mathcal{Y}}{\text{minimize}} \quad & f(u, y) \\ \text{subject to} \quad & \begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} Y_{\text{f}} \\ U_{\text{f}} \end{bmatrix} g \end{aligned}$$

where g solves

$$\begin{aligned} \underset{g}{\arg\min} \quad & \|g\|_2^2 \\ \text{subject to} \quad & \begin{bmatrix} U_{\text{p}} \\ Y_{\text{p}} \\ U_{\text{f}} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix} \end{aligned}$$

regularized DeePC

$$\begin{aligned} \underset{g, u \in \mathcal{U}, y \in \mathcal{Y}}{\text{minimize}} \quad & f(u, y) + \lambda_g \|g\|_2^2 \\ \text{subject to} \quad & \begin{bmatrix} U_{\text{p}} \\ Y_{\text{p}} \\ U_{\text{f}} \\ Y_{\text{f}} \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix} \end{aligned}$$

$\Rightarrow$ feasible set of ID & MPC

$\subseteq$ feasible set for DeePC

$\Rightarrow$ DeePC $\leq$ MPC + $\lambda_g \cdot$ ID

"easier to learn control policies from data rather than models"

DeePC vs. System ID & MPC

"It is easier to learn control policies from data rather than models."

1) Optimality certificate for subspace & prediction error ID methods

$$\underbrace{\text{control cost} + \lambda_g \cdot \text{regularizer}}_{\text{cost of DeePC}} \leq \underbrace{\text{control cost} + \lambda_g \cdot \text{ID loss function}}_{\text{cost of model-based approach}}$$

Proof sketch: both problems have the same feasible set, but finding the best control subject to a model minimizing fit criterion is a bi-level problem

2) Data informativity [Camlibel, Trentelman et al. '19]

data-driven (DeePC) control is feasible even data is not rich enough for ID

3) DeePC = ID for control: model-fit criterion biased by control objective

Example: objective is to track $\sin(\omega\, t) \Rightarrow$ identify best model near ω

DeePC vs. System ID & MPC

4) Observations across many case studies from robotics & energy:

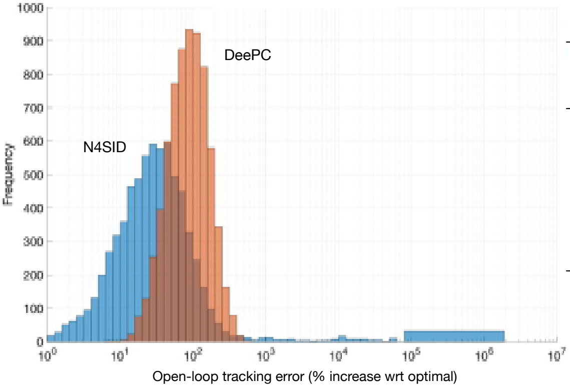

→ often ***similar*** performance

→ direct (DeePC) approach appears ***more robust*** to outliers than indirect (ID + MPC) approaches

→ direct ***often outperforms*** indirect — almost always in nonlinear closed loop

to be further explored . . .

application: ***end-to-end automation*** in energy systems

Power system case study

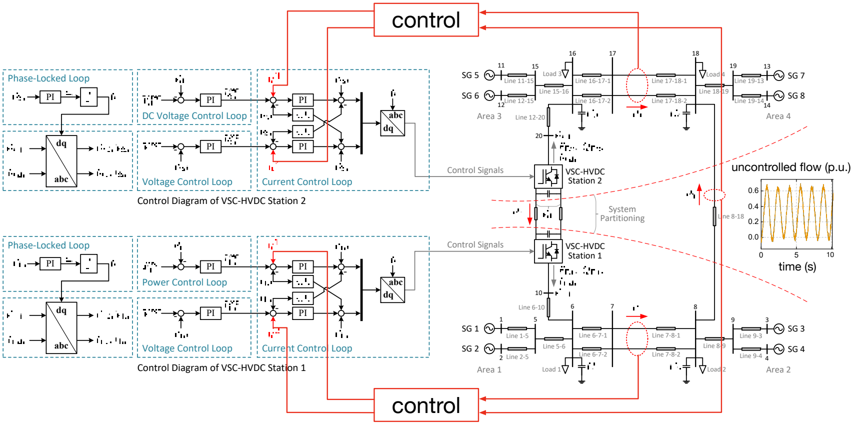

- ***complex*** 4-area power ***system***: large ($n = 208$), few measurements (8), nonlinear, noisy, stiff, input constraints, & decentralized control
- ***control objective:*** damping of inter-area oscillations via HVDC link
- ***real-time*** MPC & DeePC prohibitive $\rightarrow$ choose T, T_{ini}, & T_{future} wisely

Centralized control

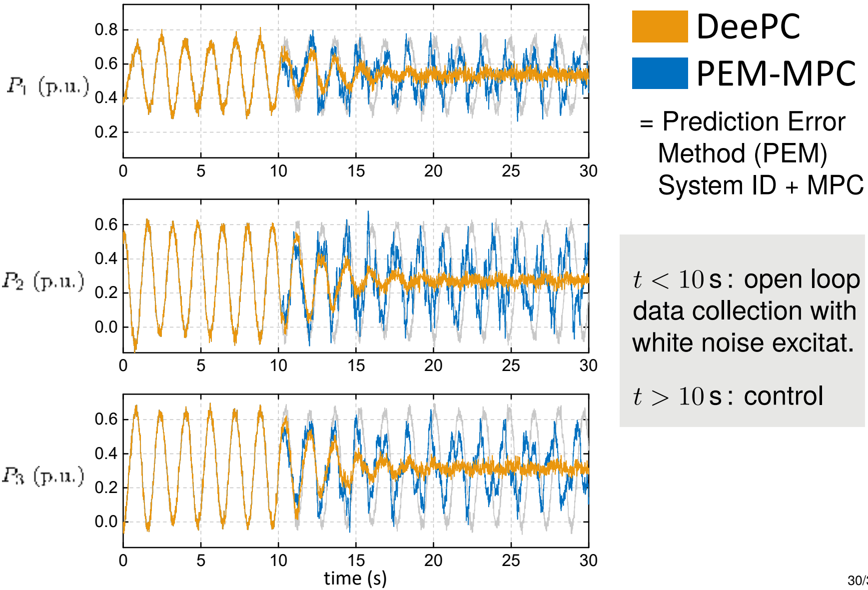

= Prediction Error Method (PEM) System ID + MPC

$t < 10\,$s: open loop data collection with white noise excitat.

$t > 10\,$s: control

Performance: DeePC wins (clearly!)

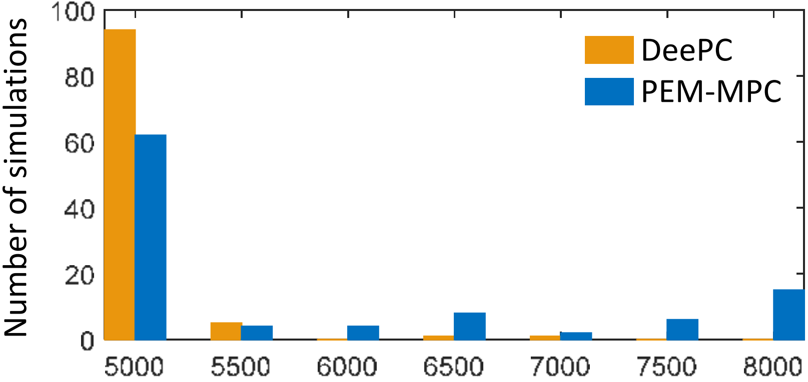

Measured closed-loop cost $= \sum_k \|y_k - r_k\|_Q^2 + \|u_k\|_R^2$

DeePC hyper-parameter tuning

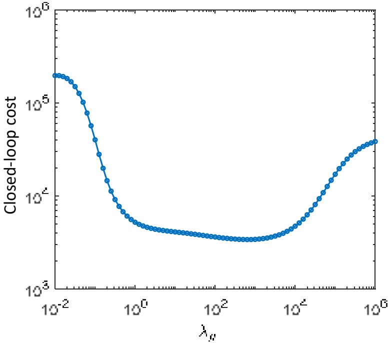

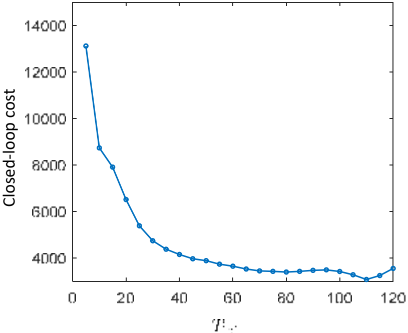

regularizer λ_g

- for distributional robustness $\approx$ radius of Wasserstein ball
- wide range of sweet spots

$\rightarrow$ choose $\lambda_g = 20$

estimation horizon T_{ini}

- for model complexity $\approx n$
- $T_{\text{ini}} \geq 50$ is sufficient & low computational complexity

$\rightarrow$ choose $T_{\text{ini}} = 60$

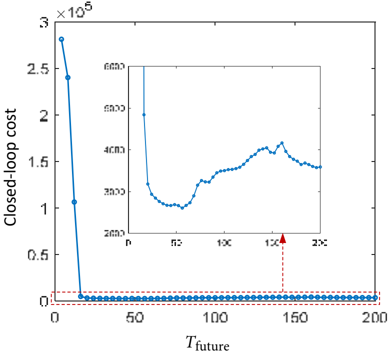

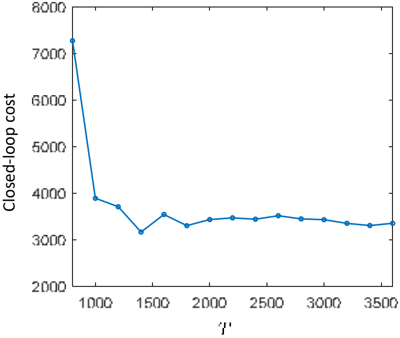

prediction horizon T_{future}

- long enough for stability

  $\rightarrow$ choose $T_{\text{future}} = 120$ and apply first 60 input steps

data length T

- long enough for persistent excitation but accordingly card$(g) = T - T_{\text{ini}} - T_{\text{future}} + 1$

  $\rightarrow$ choose $T = 1500$
  (Hankel matrix $\approx$ square)

Computational cost

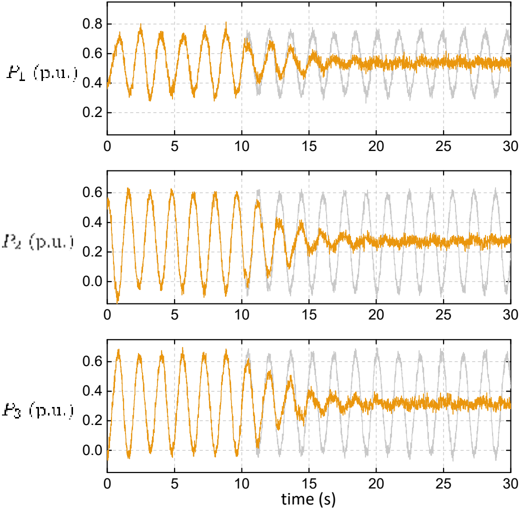

- $T = 1500$
- $\lambda_g = 20$
- $T_{\text{ini}} = 60$
- $T_{\text{future}} = 120$ and apply first 60 input steps
- sampling time = 0.02 s
- solver (OSQP) time = 1 s (on Intel Core i5 7200U)

$\Rightarrow$ **implementable**

Decentralized implementation

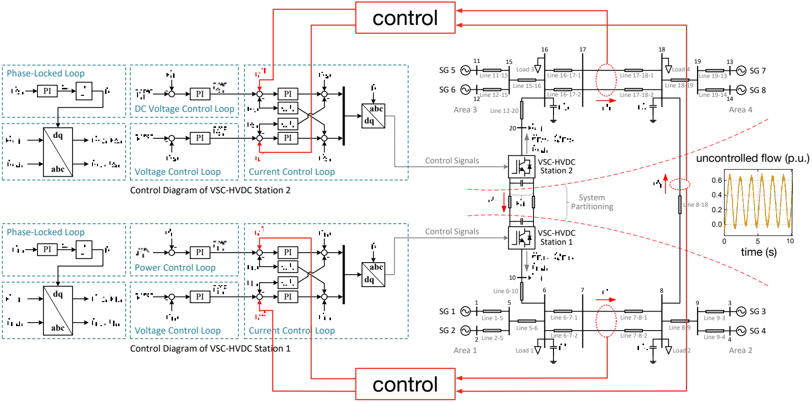

- ***plug'n'play MPC:*** treat interconnection P_3 as disturbance variable w with past disturbance w_{ini} measurable & future $w_{\text{future}} \in \mathcal{W}$ uncertain
- for each controller ***augment Hankel matrix*** with data W_p and W_f
- decentralized ***robust min-max DeePC:*** $\min_{g,u,y} \max_{w \in \mathcal{W}}$

Decentralized control performance

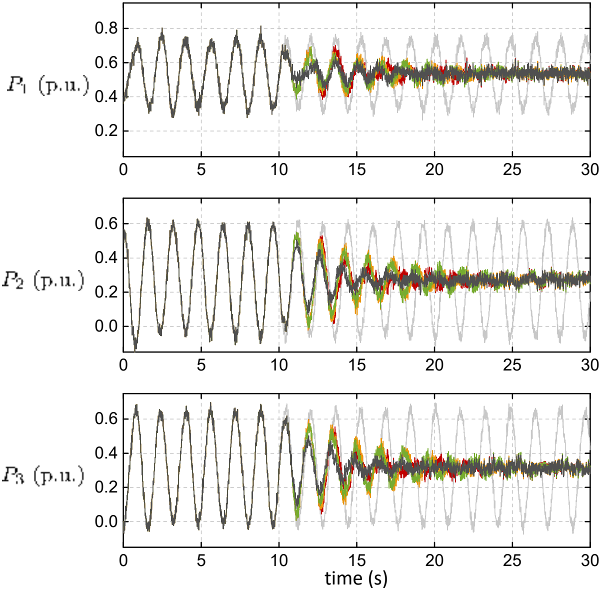

- colors correspond to different hyper-parameter settings (not discernible)
- ambiguity set $\mathcal{W}$ is ∞-ball (box)
- for computational efficiency $\mathcal{W}$ is downsampled (piece-wise linear)
- solver time ≈ 2.6 s

$\Rightarrow$ **implementable**

Summary & conclusions

- fundamental lemma from behavioral systems
- matrix time series serves as predictive model
- data-enabled predictive control (DeePC)

✓ certificates for deterministic LTI systems

✓ distributional robustness via regularizations

✓ outperforms ID + MPC in optimization metric

→ certificates for nonlinear & stochastic setup

→ adaptive extensions, explicit policies, . . .

→ applications to building automation, bio, etc.

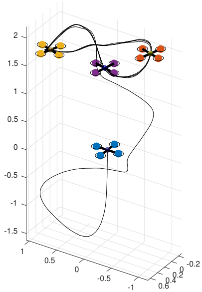

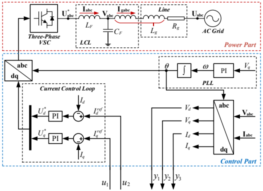

Why have these powerful ideas not been mixed long before ?

Willems '07: *"[MPC] has perhaps too little system theory and too much brute force computation in it."*

The other side often proclaims *"behavioral systems theory is beautiful but did not prove utterly useful"*