Regularized & Distributionally Robust Data-Enabled Predictive Control

Florian Dörfler

ETH Zürich

CST Seminar @ Technion

Acknowledgements

Jeremy Coulson

Linbin Huang

Paul Beuchat

John Lygeros

Ivan Markovsky

Ezzat Elokda

Perspectives on model-based control

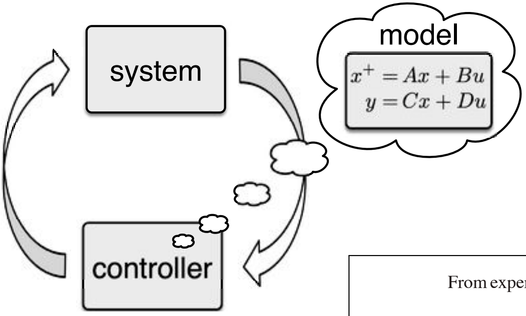

→ ***models*** useful for system analysis, design, estimation, . . . ***control***

→ ***modeling*** from first principles & ***system ID***

recurring themes

- 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*

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

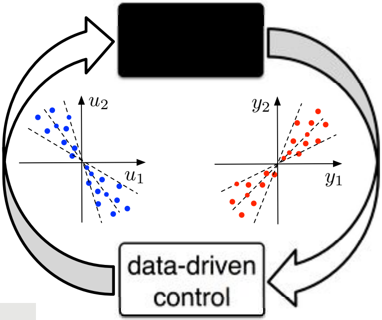

Q: Why give up physical modeling and reliable model-based algorithms ?

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 [Åström, '73]

Snippets from the literature

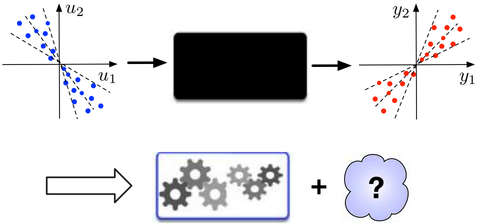

indirect data-driven control:

sequential system ID + uncertainty quantification + robust control

→ recent end-to-end design pipelines with finite-sample guarantees

ø ID seeks best but not most useful model: "easier to learn policies . . . "

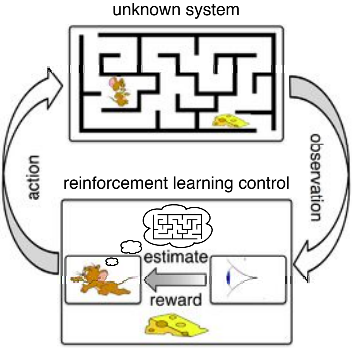

direct data-driven control:

reinforcement learning / stochastic adaptive control / approximate dynamic programming

→ spectacular theoretic & practical advances

→ more brute force storage/computation/data

ø not suitable for physical systems: real-time, safety-critical, continuous

Abstraction reveals pros & cons

indirect data-driven control

minimize	control cost (x, u)	} outer optimization
subject to	(x, u) satisfy state-space model	
where	x estimated from (u, y) & model	} middle opt.
where	model identified from (u^d, y^d) data	} inner opt.

outer + middle opt.: separation & certainty equivalence ($\rightarrow$ LQG case)

inner opt.: **no** separation ($\rightarrow$ ID-4-control)

$\rightarrow$ nested multi-level optimization problem

direct data-driven control

minimize	control cost (u, y)
subject to	(u, y) consistent with (u^d, y^d) data

$\rightarrow$ ***trade-offs***

modular vs. end-2-end
suboptimal (?) vs. optimal
convex vs. non-convex (?)

Additionally: all above should be min-max or $\mathbb{E}(\cdot)$ accounting for ***uncertainty*** ...

Colorful idea

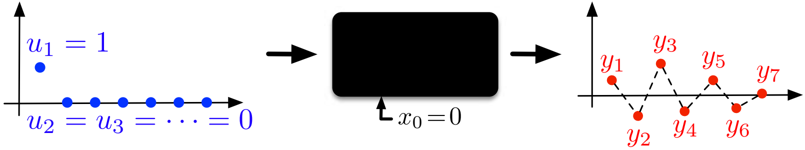

If you had the ***impulse response*** of a LTI system, then ...

- can ***identify model*** (e.g., transfer function or Kalman-Ho realization)
- or ***predictive control directly from raw data*** (dynamic matrix control)

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

- ***insight*** : single trajectory generates all others — at least conceptually
- ***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. *Distributionally Robust Chance Constrained Data-enabled Predictive Control*. https://arxiv.org/abs/2006.01702.

III. Application: End-to-End Automation in Energy & Robotics

L. Huang, J. Coulson, J. Lygeros, and F. Dörfler. *Data-Driven Wide-Area Control*. arxiv.org/abs/1911.12151.

E. Elokda, J. Coulson, P. Beuchat, J. Lygeros, and F. Dörfler. *Data-Enabled Predictive Control for Quadcopters*. https://www.research-collection.ethz.ch/.

[click here] for related publications

Preview

complex 4-area power ***system***: large ($n = 208$), few sensors (8), nonlinear, noisy, stiff, input constraints, & decentralized control specifications

control objective: oscillation damping without model

(models are proprietary, grid has many owners, operation in flux, . . .)

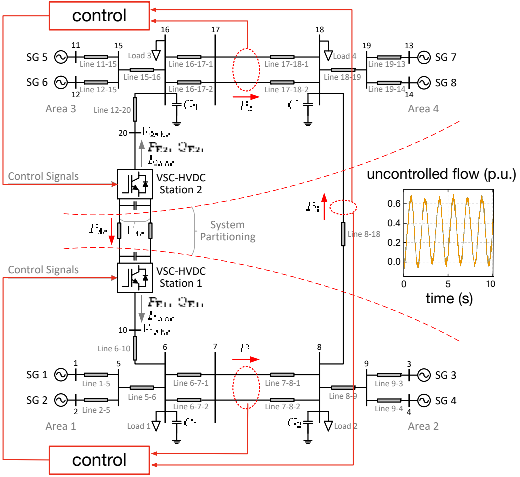

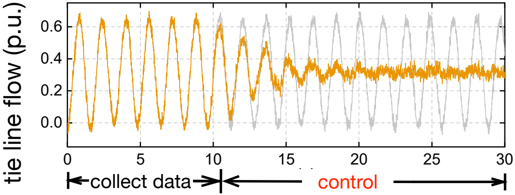

seek a method that ***works reliably***, can be ***efficiently*** implemented, & ***certifiable***

→ 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.

} $\mathscr{B}$ is the set of all trajectories

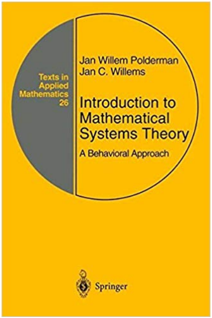

Definition: The dynamical system $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathscr{B})$ is

(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 $\sigma w_t = w_{t+1}$.

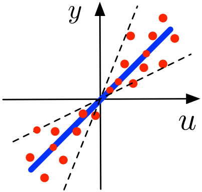

$\mathscr{B} =$ ***set of trajectories*** & $\mathscr{B}_T$ is ***restriction*** to $t \in [0, T]$

LTI systems and matrix time series

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

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

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

(ARX / kernel representation)

$\Leftarrow$

under assumptions

$\Rightarrow$

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

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

The Fundamental Lemma

Definition : The signal $u^d = \text{col}(u_1^d, \ldots, u_T^d) \in \mathbb{R}^{mT}$ is ***persistently exciting of order*** L if $\mathscr{H}_L(u) = \begin{bmatrix} u_1^d & \cdots & u_{T-L+1}^d \\ \vdots & \ddots & \vdots \\ u_L^d & \cdots & u_T^d \end{bmatrix}$ 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* $\text{col}(u^d, y^d) \in \mathscr{B}_T$, where
- u is *persistently exciting* of order $t + n$ (prediction span + # states).

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

Cartoon of Fundamental Lemma

persistently exciting | controllable LTI | sufficiently many samples

$$\text{set of trajectories} = \{(u, y) : \exists x \quad x^+ = Ax + Bu\,,\ y = Cx + Du\}$$

$$\Longleftrightarrow$$

$$\text{colspan} \begin{bmatrix} \begin{pmatrix} u_1^d \\ y_1^d \end{pmatrix} & \begin{pmatrix} u_2^d \\ y_2^d \end{pmatrix} & \begin{pmatrix} u_3^d \\ y_3^d \end{pmatrix} & \cdots \\ \begin{pmatrix} u_2^d \\ y_2^d \end{pmatrix} & \begin{pmatrix} u_3^d \\ y_3^d \end{pmatrix} & \begin{pmatrix} u_4^d \\ y_4^d \end{pmatrix} & \cdots \\ \begin{pmatrix} u_3^d \\ y_3^d \end{pmatrix} & \begin{pmatrix} u_4^d \\ y_4^d \end{pmatrix} & \begin{pmatrix} u_5^d \\ y_5^d \end{pmatrix} & \cdots \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix}$$

parametric state-space model | 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 $u \in \mathbb{R}^{m \cdot T_{\text{future}}}$ → to predict forward
- past data col$(u^{\text{d}}, y^{\text{d}}) \in \mathscr{B}_{T_{\text{data}}}$ → to form Hankel matrix

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

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

$$\mathscr{H}_{T_{\text{future}}}\begin{pmatrix} u^{\text{d}} \\ y^{\text{d}} \end{pmatrix} \boldsymbol{g} = \left[\begin{array}{cccc} u_1^{\text{d}} & u_2^{\text{d}} & \cdots & u_{T-N+1}^{\text{d}} \\ \vdots & \vdots & \ddots & \vdots \\ u_{T_{\text{future}}}^{\text{d}} & u_{T_{\text{future}}+1}^{\text{d}} & \cdots & u_T^{\text{d}} \\ \hline y_1^{\text{d}} & y_2^{\text{d}} & \cdots & y_{T-N+1}^{\text{d}} \\ \vdots & \vdots & \ddots & \vdots \\ y_{T_{\text{future}}}^{\text{d}} & y_{T_{\text{future}}+1}^{\text{d}} & \cdots & y_T^{\text{d}} \end{array}\right] \boldsymbol{g} = \left[\begin{array}{c} u_1 \\ \vdots \\ u_{T_{\text{future}}} \\ \hline y_1 \\ \vdots \\ y_{T_{\text{future}}} \end{array}\right]$$

Issue: predicted output is not unique → need to set initial conditions !

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

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

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

Solution: given u & col$(u_{\text{ini}}, y_{\text{ini}})$ → compute $\boldsymbol{g}$ & y from

$$\begin{bmatrix} \mathscr{H}_{T_{\text{ini}}}\begin{pmatrix} u^{\text{d}} \\ y^{\text{d}} \end{pmatrix} \\ \mathscr{H}_{T_{\text{future}}}\begin{pmatrix} u^{\text{d}} \\ y^{\text{d}} \end{pmatrix} \end{bmatrix} \boldsymbol{g} = \left[\begin{array}{ccc} u^{\text{d}}_1 & \cdots & u^{\text{d}}_{T-T_{\text{future}}-T_{\text{ini}}+1} \\ \vdots & \ddots & \vdots \\ u^{\text{d}}_{T_{\text{ini}}} & \cdots & u^{\text{d}}_{T-T_{\text{future}}} \\ \hline y^{\text{d}}_1 & \cdots & y^{\text{d}}_{T-T_{\text{future}}-T_{\text{ini}}+1} \\ \vdots & \ddots & \vdots \\ y^{\text{d}}_{T_{\text{ini}}} & \cdots & y^{\text{d}}_{T-T_{\text{future}}} \\ \hline u^{\text{d}}_{T_{\text{ini}}+1} & \cdots & u^{\text{d}}_{T-T_{\text{future}}+1} \\ \vdots & \ddots & \vdots \\ u^{\text{d}}_{T_{\text{ini}}+T_{\text{future}}} & \cdots & u^{\text{d}}_{T} \\ \hline y^{\text{d}}_{T_{\text{ini}}+1} & \cdots & y^{\text{d}}_{T-T_{\text{future}}+1} \\ \vdots & \ddots & \vdots \\ y^{\text{d}}_{T_{\text{ini}}+T_{\text{future}}} & \cdots & y^{\text{d}}_{T} \end{array}\right] \boldsymbol{g} = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix}$$

⇒ observability condition: if $T_{\text{ini}} \geq$ lag of system, then y is unique

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 & stochastic → Ivan Markovsky & ourselves

explicit & deterministic → Claudio de Persis & Pietro Tesi

→ ***lots of recent momentum*** (∼ 1 ArXiv / week) with contributions by
Scherer, Allgöwer, Camlibel, Trentelman, Pappas, Fischer, Pasqualetti, Goulart, Mesbahi, . . .

→ more classic ***subspace predictive control*** (De Moor) literature

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

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

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

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

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 Hankel matrix for receding-horizon prediction / estimation:

$$\begin{aligned} \underset{g,u,y}{\text{minimize}} \quad & \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 \\ \text{subject to} \quad & \mathscr{H}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix}, \\ & u_k \in \mathcal{U}, \quad \forall k \in \{0, \dots, T_{\text{future}} - 1\}, \\ & y_k \in \mathcal{Y}, \quad \forall k \in \{0, \dots, T_{\text{future}} - 1\} \end{aligned}$$

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

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

hard operational or safety **constraints**

- Hankel matrix $\mathscr{H}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix} = \begin{bmatrix} \mathscr{H}_{T_{\text{ini}}}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix} \\ \mathscr{H}_{T_{\text{future}}}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix} \end{bmatrix}$ from past data

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**

Consistency 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 :

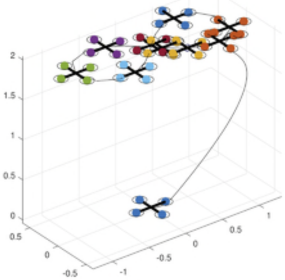

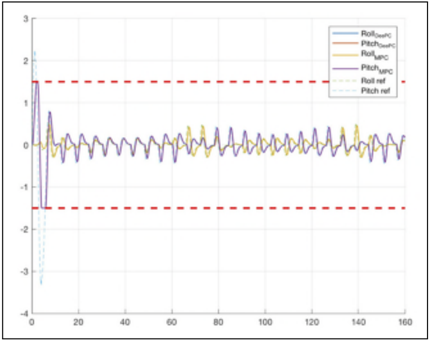

Thus, ***MPC carries over to DeePC***

…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 ?

…playing on certainty-equivalence will fail !
$\rightarrow$ need a robustified approach

Noisy real-time measurements

$$\begin{aligned}\underset{g,u,y}{\text{minimize}} \quad & \sum_{k=0}^{T_{\text{future}}-1} \|y_k - r_{t+k}\|_Q^2 + \|u_k\|_R^2 + \lambda_y \|\sigma_{\text{ini}}\|_p \\ \text{subject to} \quad & \mathscr{H}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_{\text{ini}} \\ 0 \\ 0 \end{bmatrix}, \\ & u_k \in \mathcal{U}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}, \\ & y_k \in \mathcal{Y}, \quad \forall k \in \{0, \ldots, T_{\text{future}} - 1\}\end{aligned}$$

Solution: add ℓ_p**-slack** σ_{ini} to ensure feasibility

→ receding-horizon least-square filter

→ for $\lambda_y \gg 1$: constraint is slack only if infeasible

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

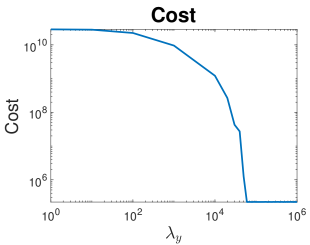

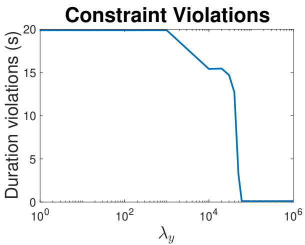

Hankel matrix corrupted by noise

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

$$\text{subject to} \quad \mathscr{H}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix},$$

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

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

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

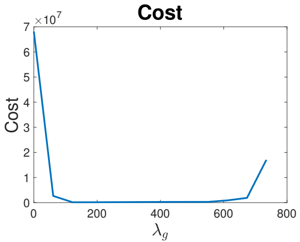

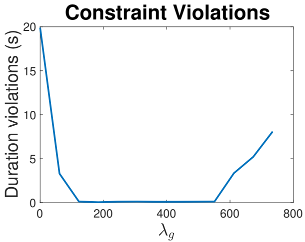

Towards nonlinear systems . . .

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 :

DeePC

+ σ_{ini} slack

+ $\|g\|_1$ regularizer

+ more columns in $\mathscr{H}\begin{pmatrix} u^{\mathrm{d}} \\ y^{\mathrm{d}} \end{pmatrix}$

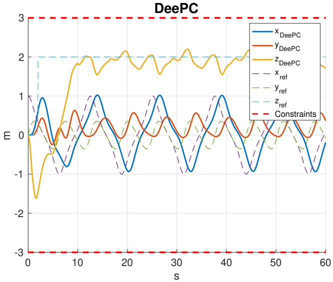

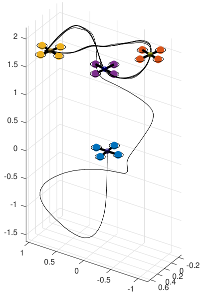

Experimental snippet

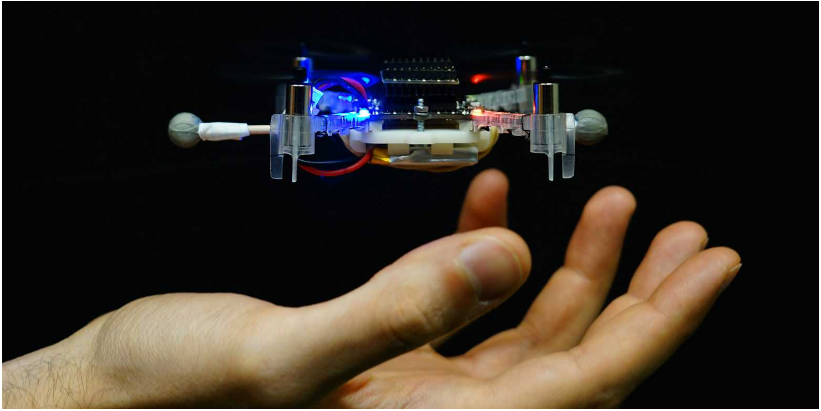

Consistent observations across case studies — more than a fluke

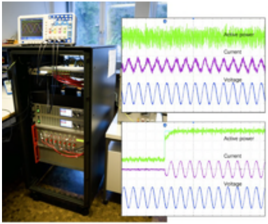

grid-connected converter

quad coptor fig-8 tracking

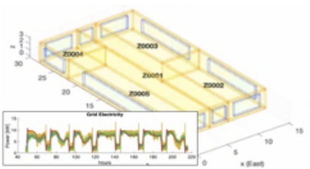
energy hub & building automation

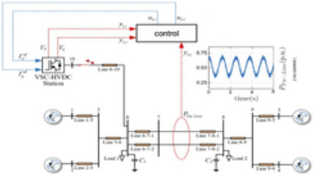

power system oscillation damping (see later)

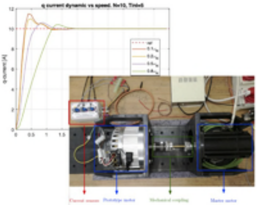
synchronous motor drive

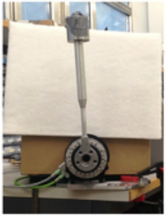
pendulum swing up

let’s try to put some theory behind all of this . . .

Distributionally robust formulation

- ***problem abstraction***: $\min_{x\in\mathcal{X}} \; c\left(\widehat{\xi}, x\right) \;=\; \min_{x\in\mathcal{X}} \; \mathbb{E}_{\widehat{\mathbb{P}}}\left[c\left(\xi, x\right)\right]$

  where $\widehat{\xi}$ denotes *measured* data (possibly not from deterministic LTI), and $\widehat{\mathbb{P}} = \delta_{\widehat{\xi}}$ denotes the *empirical distribution* of the data $\widehat{\xi}$

⇒ ***poor out-of-sample performance*** of above sample-average solution $x^\star$ for real problem: $\mathbb{E}_{\mathbb{P}}\left[c\left(\xi, x^\star\right)\right]$ where $\mathbb{P}$ is the *unknown distribution* of ξ

- ***distributionally robust*** formulation:

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

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

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

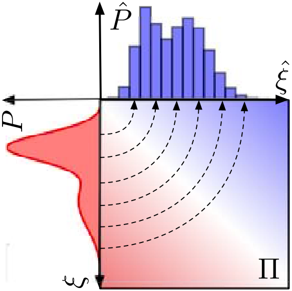

note: Wasserstein ball does not only include LTI systems with additive Gaussian noise but “everything” (integrable)

- ***distributionally robust*** formulation:

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

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

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

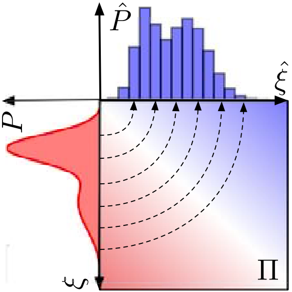

Theorem: Under minor technical conditions:

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

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

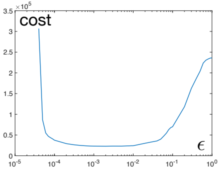

Proof uses methods by Esfahani & Kuhn: semi-infinite problem becomes tractable after marginalization, for discrete worst case, & with many convex conjugates.

Further ingredients & improvements

averaging & measure concentration

- multiple i.i.d. experiments → sample ***average Hankel matrix*** $\frac{1}{N}\sum_{i=1}^{n}\mathscr{H}_i(y^{\mathsf{d}})$
- ***measure concentration***: Wasserstein ball $\mathbb{B}_\epsilon(\widehat{P})$ includes true distribution $\mathbb{P}$ with high confidence if $\epsilon \sim 1/N^{1/\dim(\xi)}$

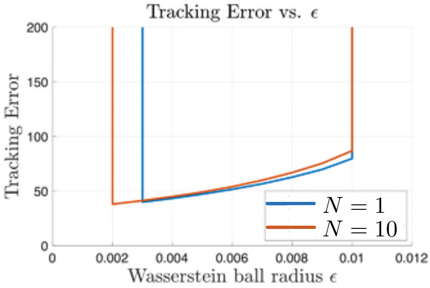

distributionally robust probabilistic constraints

$\sup_{Q\in\mathbb{B}_\epsilon(\widehat{P})} \mathrm{CVaR}^{\mathbb{Q}}_{1-\alpha}$ $\Leftrightarrow$ averaging + regularization + tightening

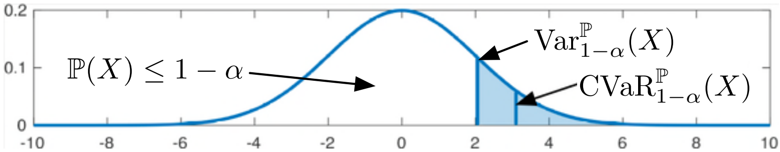

change predictor structure from Hankel to ***Chinese page matrix***

$$\mathscr{H}\begin{pmatrix} u^{d} \\ y^{d} \end{pmatrix} = \begin{bmatrix} \begin{pmatrix} u_1^d \\ y_1^d \end{pmatrix} & \begin{pmatrix} u_2^d \\ y_2^d \end{pmatrix} & \cdots \\ \begin{pmatrix} u_2^d \\ y_2^d \end{pmatrix} & \begin{pmatrix} u_3^d \\ y_3^d \end{pmatrix} & \iddots \\ \begin{pmatrix} u_3^d \\ y_3^d \end{pmatrix} & \begin{pmatrix} u_4^d \\ y_4^d \end{pmatrix} & \iddots \\ \vdots & \iddots & \iddots \\ \begin{pmatrix} u_L^d \\ y_L^d \end{pmatrix} & \begin{pmatrix} u_{L+1}^d \\ y_{L+1}^d \end{pmatrix} & \cdots \end{bmatrix} \rightarrow \mathscr{P}\begin{pmatrix} u^{d} \\ y^{d} \end{pmatrix} = \begin{bmatrix} \begin{pmatrix} u_1^d \\ y_1^d \end{pmatrix} & \begin{pmatrix} u_{L+1}^d \\ y_{L+1}^d \end{pmatrix} & \cdots \\ \begin{pmatrix} u_2^d \\ y_2^d \end{pmatrix} & \begin{pmatrix} u_{L+2}^d \\ y_{L+2}^d \end{pmatrix} & \vdots \\ \begin{pmatrix} u_3^d \\ y_3^d \end{pmatrix} & \begin{pmatrix} u_{L+3}^d \\ y_{L+3}^d \end{pmatrix} & \vdots \\ \vdots & \vdots & \vdots \\ \begin{pmatrix} u_L^d \\ y_L^d \end{pmatrix} & \begin{pmatrix} u_{2L}^d \\ y_{2L}^d \end{pmatrix} & \cdots \end{bmatrix}$$

$\rightarrow$ more data but independent entries $\rightarrow$ statistical & algorithmic pros
e.g. distr. robust. estimates tight & SVD-rank-reduction etc.

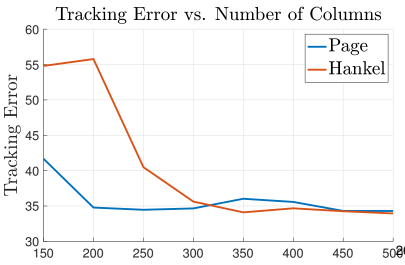

All together in action for nonlinear & stochastic quadcoptor setup

case study :

distr. robust objective

+ Page matrix predictor
+ averaging
+ CVaR constraints
+ σ_{ini} slack

→ DeePC works much better than it should !

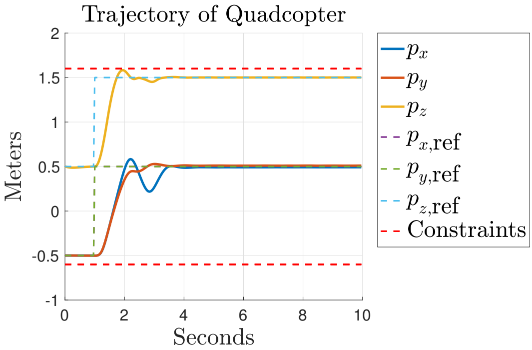

main catch : optimization problems become large (no-free-lunch)

→ models are compressed, de-noised, & tidied-up representations recall the ***central promise*** :

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

Comparison: DeePC vs. ID + MPC

DeePC with ℓ_1-regularizer

certainty-equivalence ***MPC*** based on prediction error ***ID***

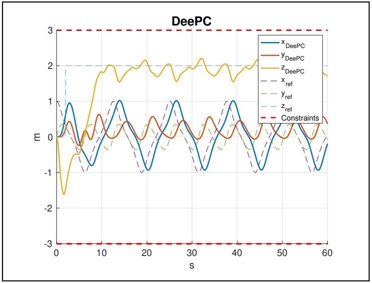

single fig-8 run

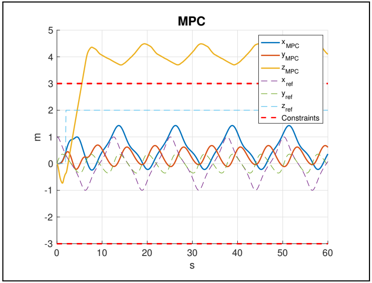

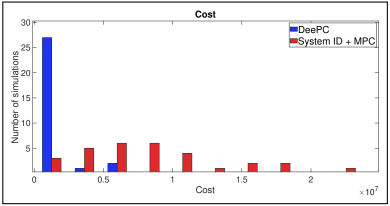

random sims

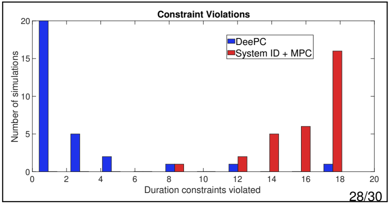

More to it than a single case study ?

consistent across all nonlinear case studies : DeePC always wins

reason (?) : DeePC is robust, whereas certainty-equivalence control is based on identified model with a bias error

stochastic LTI comparison (no bias) show certainty-equivalence vs. robust control trade-offs (mean vs. median)

link : DeePC includes implicit sys ID though biased by control objective & robustified through regularizations

$\rightarrow$ lot more to be understood ...

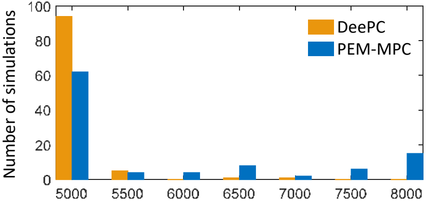

measured closed-loop cost $= \sum_k \|y_k - r_k\|_Q^2 + \|u_k\|_R^2$

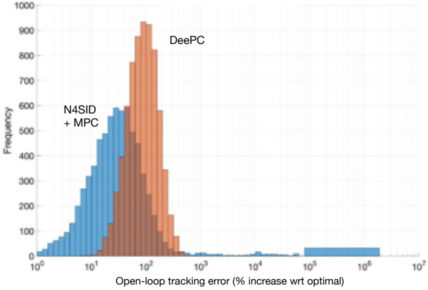

Summary & conclusions

main take-aways

- matrix time series serves as predictive model
- data-enabled predictive control (DeePC)

✓ consistent for deterministic LTI systems

✓ distributional robustness via regularizations

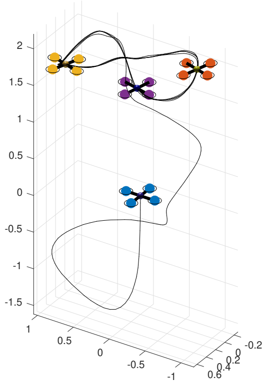

future work

→ tighter certificates for nonlinear systems

→ explicit policies & direct adaptive control

→ seek application with a "business case"

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"*

Thanks !

Florian Dörfler

mail: dorfler@ethz.ch

[link] to homepage

[link] to related publications

appendix

end-to-end automation case study in power systems

Power system case study

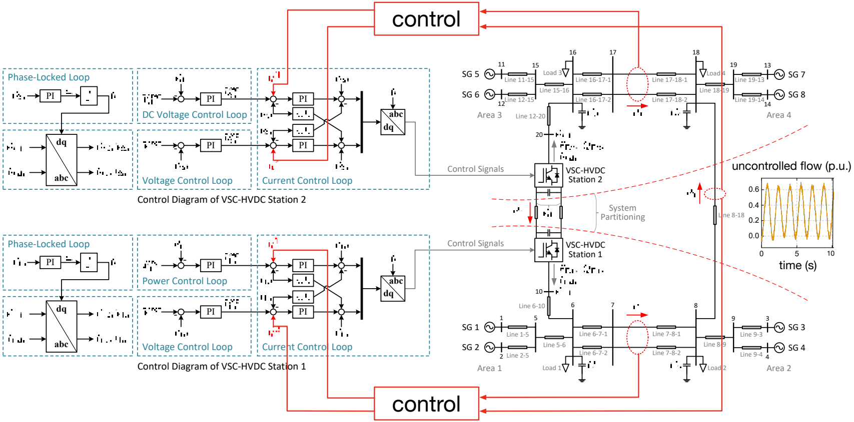

- ***complex*** 4-area power ***system***: large ($n = 208$), few measurements (8), nonlinear, noisy, stiff, input constraints, & decentralized control
- ***control objective:*** damping of inter-area oscillations via HVDC link
- ***real-time*** MPC & DeePC prohibitive $\rightarrow$ choose T, T_{ini},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

Centralized control

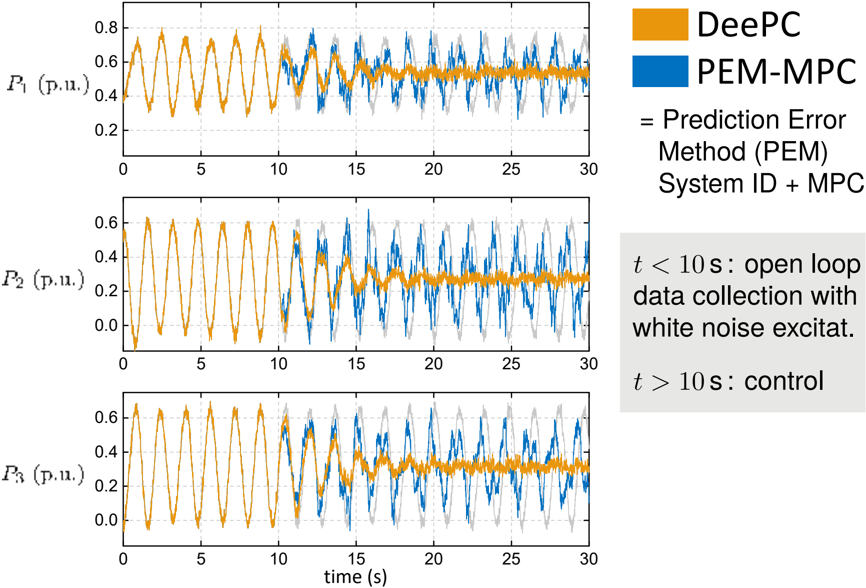

= Prediction Error Method (PEM) System ID + MPC

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

$t > 10$ s : control

Performance: DeePC wins (clearly!)

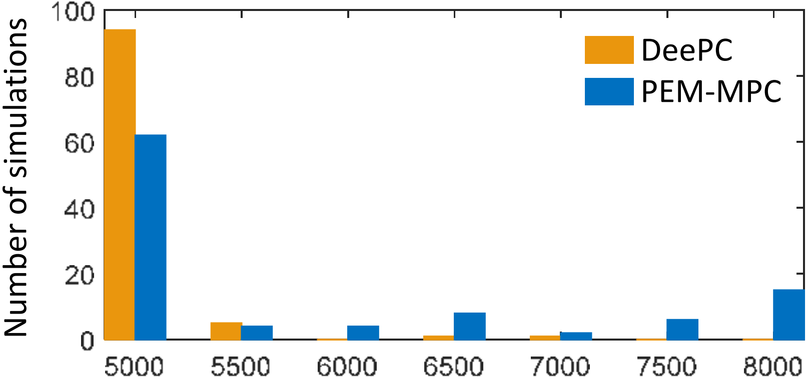

Measured closed-loop cost $= \sum_k \|y_k - r_k\|_Q^2 + \|u_k\|_R^2$

DeePC hyper-parameter tuning

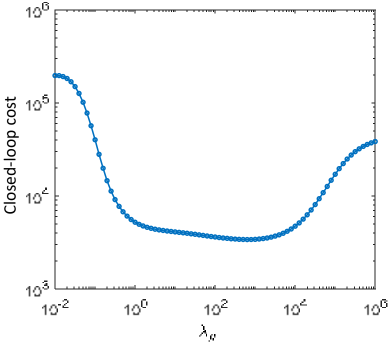

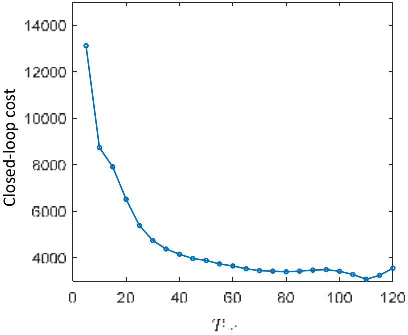

regularizer λ_g

- 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 $T_{\text{ini}} = 60$

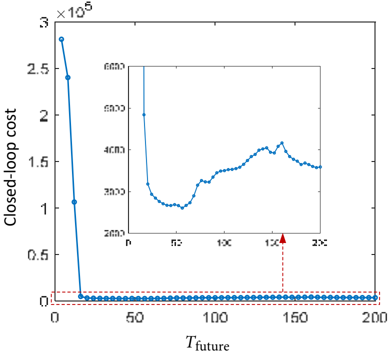

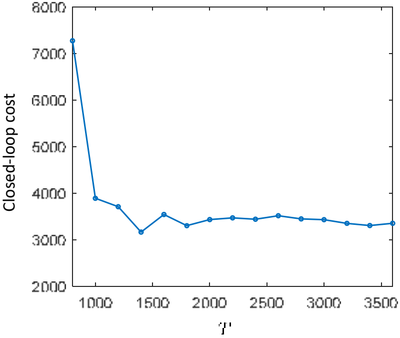

prediction horizon T_{future}

- 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

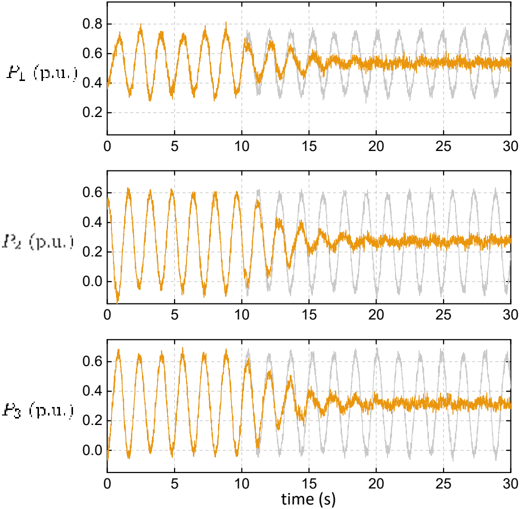

- $T = 1500$
- $\lambda_g = 20$
- $T_{\text{ini}} = 60$
- $T_{\text{future}} = 120$ and apply first 60 input steps
- sampling time = 0.02 s
- solver (OSQP) time = 1 s (on Intel Core i5 7200U)

$\Rightarrow$ **implementable**

Comparison: Hankel & Page matrix

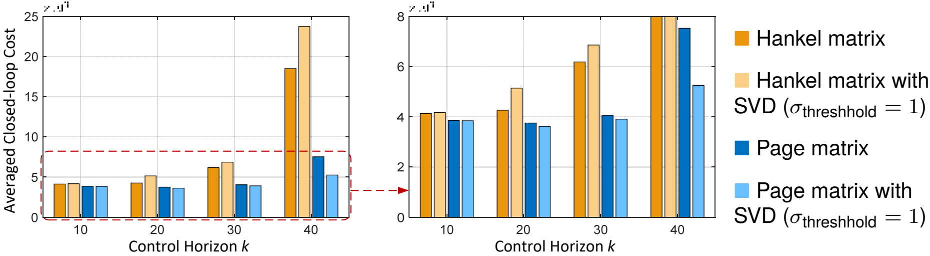

- comparison baseline: Hankel and Page matrices of ***same size***
- ***perfomance***: Page consistency beats Hankel matrix predictors
- offline ***denoising via SVD threshholding*** works wonderfully for Page though obviously not for Hankel (entries are constrained)
- effects very pronounced for ***longer horizon*** (= open-loop time)
- ***price-to-be-paid***: Page matrix predictor requires more data

Decentralized implementation

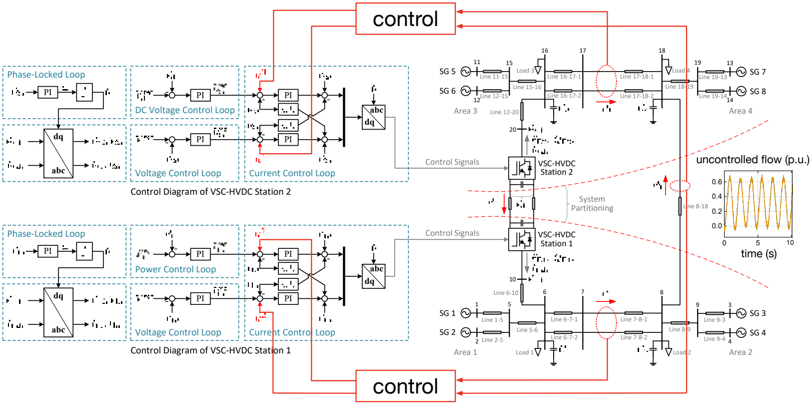

- ***plug'n'play MPC:*** treat interconnection P_3 as disturbance variable w with past disturbance w_{ini} measurable & future $w_{\text{future}} \in \mathcal{W}$ uncertain
- for each controller ***augment Hankel matrix*** with data W_p and W_f
- decentralized ***robust min-max DeePC:*** $\min_{g,u,y}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

Decentralized control performance

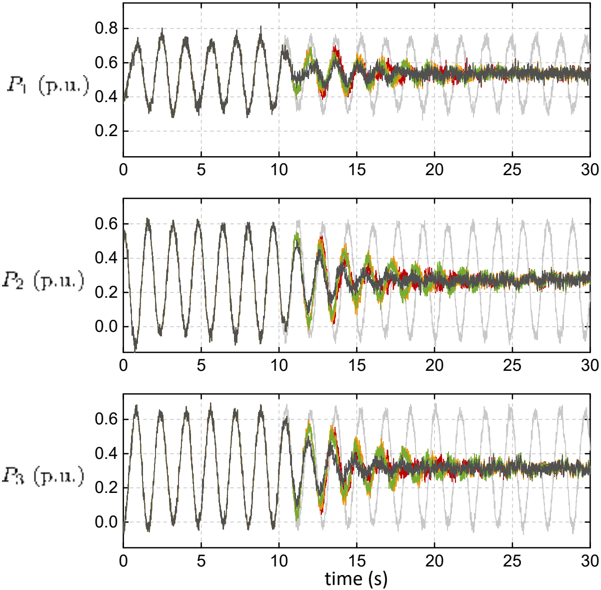

- colors correspond to different hyper-parameter settings (not discernible)
- ambiguity set $\mathcal{W}$ is ∞-ball (box)
- for computational efficiency $\mathcal{W}$ is downsampled (piece-wise linear)
- solver time ≈ 2.6 s

$\Rightarrow$ **implementable**