Data-Enabled Predictive Control in Autonomous Energy Systems

Florian Dörfler

ETH Zürich

ILES Workshop 28/10/2022

Acknowledgements

Jeremy Coulson

Linbin Huang

John Lygeros

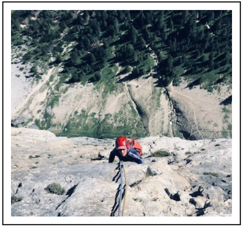

Ivan Markovsky

Further:

Ezzat Elokda,
Paul Beuchat,
Daniele Alpago,
Jianzhe (Trevor) Zhen,
Claudio de Persis,
Pietro Tesi,
Henk van Waarde,
Eduardo Prieto,
Saverio Bolognani,
Andrea Favato,
Paolo Carlet,
Andrea Martin,
Luca Furieri,
Giancarlo Ferrari-Trecate,
...
IfA DeePC team, &
many master students

Thoughts on data-driven control

- ***indirect data-driven control*** via models:
  data $\xrightarrow{\text{SysID}}$ model + uncertainty $\to$ control
- growing trend: ***direct data-driven control***
  by-passing models . . . (again) hyped, why ?

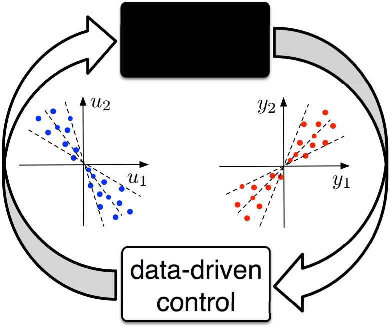

The direct approach is ***viable alternative***

- for some ***applications*** : model-based approach is too complex to be useful
  $\to$ too complex models, environments, sensing modalities, specifications (e.g., wind farm)
- due to (well-known) ***shortcomings of ID***
  $\to$ too cumbersome, models not identified for control, incompatible uncertainty estimates, ...
- when ***brute force*** data/compute available

Central promise: *It is often easier to learn a control policy from data rather than a model.*

Example: PID [Åström et al., '73]

$\to$ **theory trade-offs**: (non)modular + (in)tractable + (sub)optimal (?)

Today: tractable direct approach

1. behavioral system theory: ***fundamental lemma***
2. ***DeePC*** : data-enabled predictive control
3. robustification via salient ***regularizations***
4. cases studies from ***wind*** & ***power systems***

blooming literature (2-3 ArXiv / week)

→ survey & tutorial to get started:

DATA-DRIVEN CONTROL BASED ON BEHAVIORAL APPROACH: FROM THEORY TO APPLICATIONS IN POWER SYSTEMS

Ivan Markovsky, Linbin Huang, and Florian Dörfler
I. Markovsky is with ICREA, Pg. Lluis Companys 23, Barcelona, and CIMNE, Gran Capitàn, Barcelona, Spain (e-mail: imarkovsky@cimne.upc.edu),
L. Huang and F. Dörfler are with the Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland (e-mails: linhuang@ethz.ch, dorfler@ethz.ch).

Annual Reviews in Control 52 (2021) 42–64

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

ELSEVIER

Review article

Behavioral systems theory in data-driven analysis, signal processing, and control

Ivan Markovsky [a,*], Florian Dörfler [b]

[a] Department ELEC, Vrije Universiteit Brussel, Brussels, 1050, Belgium
[b] Automatic Control Laboratory (IfA), ETH Zürich, Zürich, 8092, Switzerland

Preview

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

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

control objective: oscillation damping $\boxed{\text{without model}}$

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

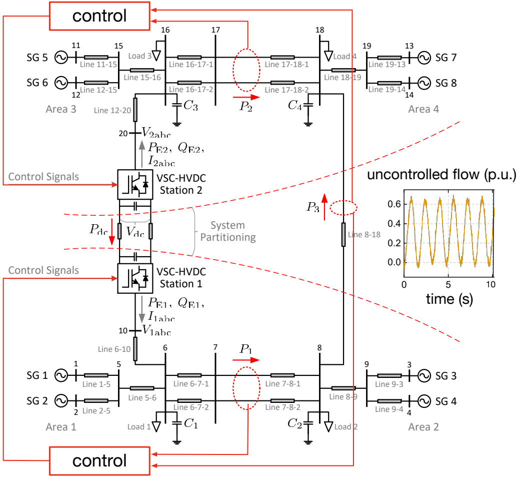

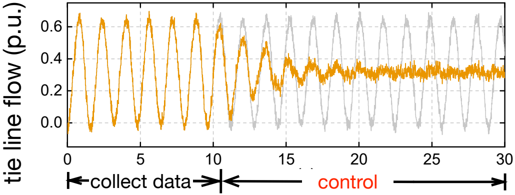

seek a method that ***works reliably***, can be ***efficiently*** implemented, & ***certifiable***

$\rightarrow$ automating ourselves

Reality check: magic or hoax ?

surely, nobody would put apply such a **shaky data-driven method**

- on the **world's most complex engineered system** (the electric grid),
- using the **world's biggest actuators** (Gigawatt-sized HVDC links),
- and subject to **real-time, safety, & stability constraints** . . . right?

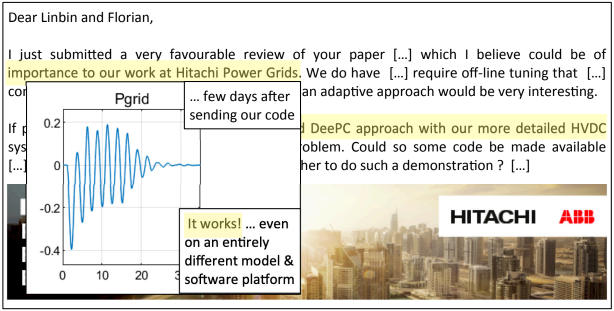
Dear Linbin and Florian,

I just submitted a very favourable review of your paper [...] which I believe could be of importance to our work at Hitachi Power Grids. We do have [...] require off-line tuning that [...] co[...] an adaptive approach would be very interesting.

If p[...] DeePC approach with our more detailed HVDC sys[...]oblem. Could so some code be made available [...] [...]her to do such a demonstration ? [...]

at least someone believes that DeePC is practically useful . . .

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*, &

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

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

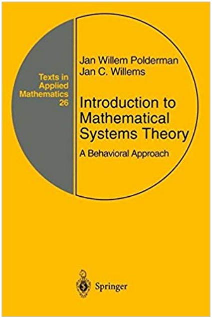

Definition: The dynamical system $(\mathbb{Z}_{\geq 0}, \mathbb{W}, \mathscr{B})$ is

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

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

LTI system $=$ shift-invariant subspace of trajectory space

$\longrightarrow$ abstract perspective suited for ***data-driven control***

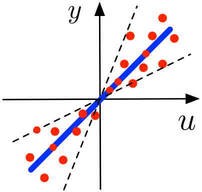

Fundamental Lemma [Willems et al. ’05 + many recent extensions]

Given: data $\begin{pmatrix} u_i^d \\ y_i^d \end{pmatrix} \in \mathbb{R}^{m+p}$ & LTI complexity parameters $\begin{cases} \text{lag } \ell \\ \text{order } n \end{cases}$

$$\underbrace{\begin{array}{l} \text{set of all } T\text{-length trajectories } = \\ \left\{ (u, y) \in \mathbb{R}^{(m+p)T} : \exists x \in \mathbb{R}^{nT} \text{ s.t.} \right. \\ \left. x^+ = Ax + Bu \,,\ y = Cx + Du \right\} \end{array}}_{\text{parametric state-space model}} \quad = \quad \underbrace{\text{colspan} \begin{bmatrix} \begin{pmatrix} u_{1,1}^d \\ y_{1,1}^d \end{pmatrix} & \begin{pmatrix} u_{1,2}^d \\ y_{1,2}^d \end{pmatrix} & \begin{pmatrix} u_{1,3}^d \\ y_{1,3}^d \end{pmatrix} & \cdots \\ \begin{pmatrix} u_{2,1}^d \\ y_{2,1}^d \end{pmatrix} & \begin{pmatrix} u_{2,2}^d \\ y_{2,2}^d \end{pmatrix} & \begin{pmatrix} u_{2,3}^d \\ y_{2,3}^d \end{pmatrix} & \cdots \\ \vdots & \vdots & \vdots & \vdots \\ \begin{pmatrix} u_{T,1}^d \\ y_{T,1}^d \end{pmatrix} & \begin{pmatrix} u_{T,2}^d \\ y_{T,2}^d \end{pmatrix} & \begin{pmatrix} u_{T,3}^d \\ y_{T,3}^d \end{pmatrix} & \cdots \end{bmatrix}}_{\text{raw data (every column is an experiment)}}$$

if and only if the trajectory matrix has rank $m \cdot T + n$ for all $T > \ell$

$$\text{set of all } T\text{-length trajectories} = \left\{ (u, y) \in \mathbb{R}^{(m+p)T} : \exists x \in \mathbb{R}^{nT} \text{ s.t. } x^+ = Ax + Bu\,,\ y = Cx + Du \right\} = \operatorname{colspan} \begin{bmatrix} \begin{pmatrix} u^d_{1,1} \\ y^d_{1,1} \end{pmatrix} & \begin{pmatrix} u^d_{1,2} \\ y^d_{1,2} \end{pmatrix} & \begin{pmatrix} u^d_{1,3} \\ y^d_{1,3} \end{pmatrix} & \cdots \\ \begin{pmatrix} u^d_{2,1} \\ y^d_{2,1} \end{pmatrix} & \begin{pmatrix} u^d_{2,2} \\ y^d_{2,2} \end{pmatrix} & \begin{pmatrix} u^d_{2,3} \\ y^d_{2,3} \end{pmatrix} & \cdots \\ \vdots & \vdots & \vdots & \vdots \\ \begin{pmatrix} u^d_{T,1} \\ y^d_{T,1} \end{pmatrix} & \begin{pmatrix} u^d_{T,2} \\ y^d_{T,2} \end{pmatrix} & \begin{pmatrix} u^d_{T,3} \\ y^d_{T,3} \end{pmatrix} & \cdots \end{bmatrix}$$

all trajectories constructible from finitely many previous trajectories

- ***standing on the shoulders of giants***: classic Willems' result was only "if" & required further assumptions: Hankel, persistency of excitation, controllability

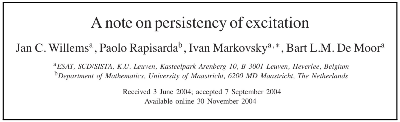
A note on persistency of excitation

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

- terminology ***fundamental*** is justified : motion primitives, subspace SysID, dictionary learning, (E)DMD, ... all implicitly rely on this equivalence

- many recent ***extensions*** to other ***system classes*** (bi-linear, descriptor, LPV, delay, Volterra series, Wiener-Hammerstein, ...), other ***matrix data structures*** (mosaic Hankel, Page, ...), & other ***proof methods***

- ***blooming literature*** (2-3 / week) on theory, applications, & computation

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

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

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

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

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

model for **estimation** with $k \in [-T_{\text{ini}} - 1, 0]$ & $T_{\text{ini}} \geq$ lag (many flavors)

hard operational or safety **constraints**

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

Elegance aside, for a deterministic LTI plant with known model, MPC is truly the ***gold standard of control***.

Data-Enabled Predictive Control

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

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

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

hard operational or safety **constraints**

- real-time measurements $(u_{\text{ini}}, y_{\text{ini}})$ for estimation — updated **online**
- trajectory matrix $\mathscr{H}\begin{pmatrix} u^{\mathsf{d}} \\ y^{\mathsf{d}} \end{pmatrix}$ from past experimental data — collected **offline** (could be adapted online)

→ ***equivalent to MPC*** in deterministic LTI case . . .
but needs to be robustified in case of noise / nonlinearity !

Regularizations counter-acting noise

$$\underset{g,u,y,\sigma}{\text{minimize}} \quad \sum_{k=1}^{T_{\text{future}}} \|y_k - r_k\|_Q^2 + \|u_k\|_R^2 + \lambda_y \|\sigma\|_p + \lambda_g h(g)$$

$$\text{subject to} \quad \mathscr{H}\begin{pmatrix} u^{\text{d}} \\ y^{\text{d}} \end{pmatrix} \cdot g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ \hline u \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma \\ 0 \\ 0 \end{bmatrix},$$

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

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

measurement noise
→ infeasible y_{ini} estimate
→ estimation slack σ
→ moving-horizon least-square filter

noisy or nonlinear (offline) **data matrix**
→ any $\binom{u}{y}$ feasible
→ add regularizer $h(g)$

Bayesian intuition: regularization ⇔ prior, e.g., $h(g) = \|g\|_1$ sparsely selects {trajectory matrix columns} = {motion primitives} ∼ low-order basis

Robustness intuition: regularization ⇔ robustifies, e.g., in a simple case

$$\min_x \max_{\|\Delta\| \leq \rho} \|(A+\Delta)x - b\| \overset{\text{tight}}{\leq} \min_x \max_{\|\Delta\| \leq \rho} \|Ax - b\| + \|\Delta x\| = \min_x \|Ax - b\| + \rho\|x\|$$

Regularization = relaxation of bi-level ID

$$\begin{aligned} &\text{minimize}_{u,y,g} && \text{control cost}\,(u,y) \\ &\text{subject to} && \begin{bmatrix} u \\ y \end{bmatrix} = \mathscr{H}\begin{pmatrix} \hat{u} \\ \hat{y} \end{pmatrix} g \\ &\text{where} && \begin{pmatrix} \hat{u} \\ \hat{y} \end{pmatrix} \in \text{argmin} \left\| \begin{pmatrix} \hat{u} \\ \hat{y} \end{pmatrix} - \begin{pmatrix} u^{\mathrm{d}} \\ y^{\mathrm{d}} \end{pmatrix} \right\| \\ &&& \text{subject to rank}\left(\mathscr{H}\begin{pmatrix} \hat{u} \\ \hat{y} \end{pmatrix}\right) = mL + n \end{aligned}$$

optimal control

system identification

↓ sequence of convex relaxations ↓

$$\begin{aligned} &\text{minimize}_{u,y,g} && \text{control cost}\,(u,y) + \lambda_g \cdot \|g\|_1 \\ &\text{subject to} && \begin{bmatrix} u \\ y \end{bmatrix} = \mathscr{H}\begin{pmatrix} u^d \\ y^d \end{pmatrix} g \end{aligned}$$

ℓ_1-regularization
= relaxation of id
smoothening order selection (no bias)

→ similar results for ID of ***parametric predictor via least squares***

$$y = K \cdot \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix}$$

where K is low-rank + causal sparsity + correlated

⇒ relaxations yields ***regularizer*** $\left\| \ker \mathscr{H}\begin{pmatrix} u^{\mathrm{d}} \\ y^{\mathrm{d}} \end{pmatrix} g \right\|$

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 by LTI on finite horizon

regularization singles out relevant features / basis functions in data

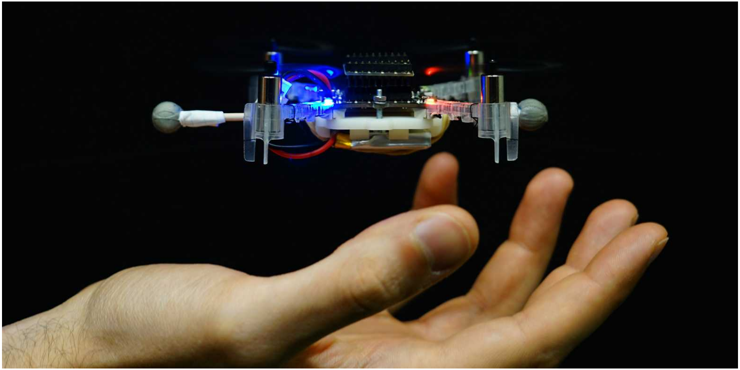

Works very well across case studies

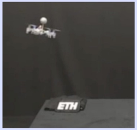

quad coptor fig-8 tracking

quadruped (by Fawcett, Afsari Amers, & Hamed)

greenhouse automation (by Automatoes)

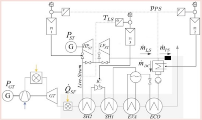

combined cycle power plant (by P Mahdavipour et. al)

robotic excavator

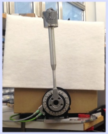

pendulum swing up

traffic coordination (by J. Wang et al.)

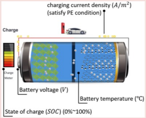

battery charging (by K. Chen et al.)

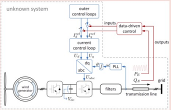

wind turbine control

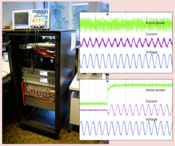

grid-connected converter

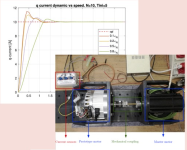

synchronous motor drive

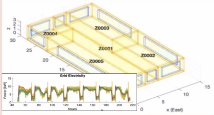

energy hub & building automation

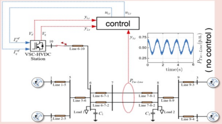

power system oscillation damping

Reason: distributional robustness

- ***problem abstraction***: $\min_{x\in\mathcal{X}} \; c\left(\widehat{\xi}, x\right)$ where $\widehat{\xi}$ is *measured* data
- ***distributionally robust*** formulation $\longrightarrow$ "$\min_{x\in\mathcal{X}} \max_{\xi} \; c\,(\xi, x)$"
  where $\max$ accounts for all stochastic processes (linear or nonlinear) that could have generated the data ... more precisely

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

where $\mathbb{B}_\epsilon(\widehat{\mathbb{P}})$ is an ***ϵ-Wasserstein ball*** centered at empirical sample distribution $\widehat{\mathbb{P}}$:

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

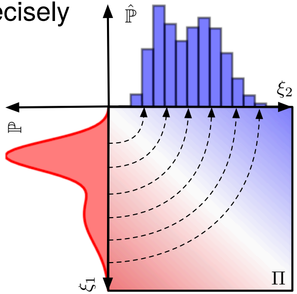

Theorem: $\underbrace{\inf_{x\in\mathcal{X}} \; \sup_{Q\in\mathbb{B}_\epsilon(\widehat{\mathbb{P}})} \; \mathbb{E}_Q\left[c\,(\xi, x)\right]}_{\text{distributional robust formulation}} \equiv \underbrace{\min_{x\in\mathcal{X}} \; c\left(\widehat{\xi}, x\right) + \epsilon\,\text{Lip}(c) \cdot \|x\|_p^\star}_{\text{previous regularized DeePC formulation}}$

Case study: wind turbine

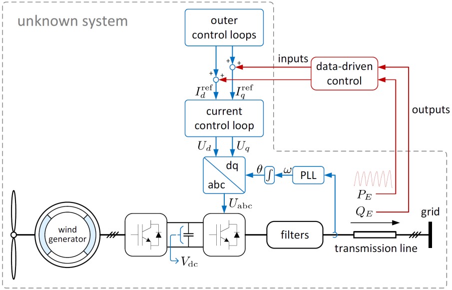

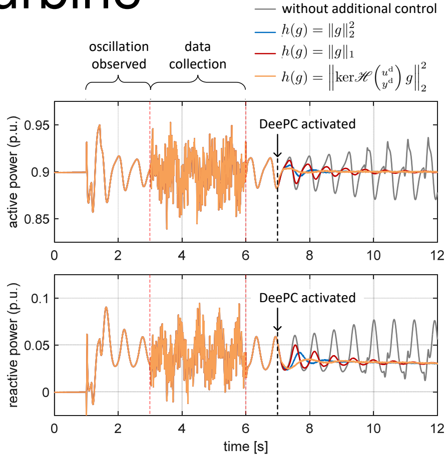

- detailed ***industrial model***: 37 states & highly nonlinear (abc $\leftrightarrow$ dq, MPTT, PLL, power specs, dynamics, etc.)
- turbine & grid model ***unknown*** to commissioning engineer & operator
- weak grid + PLL + fault $\rightarrow$ ***loss of sync***
- disturbance to be rejected by ***DeePC***

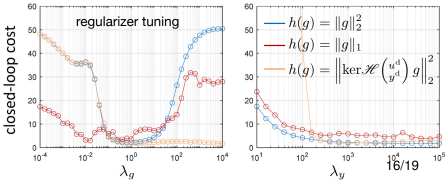

Case study ++ : wind farm

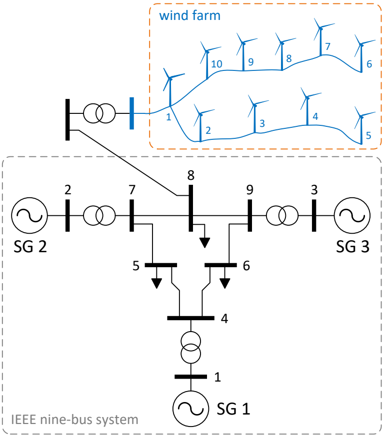

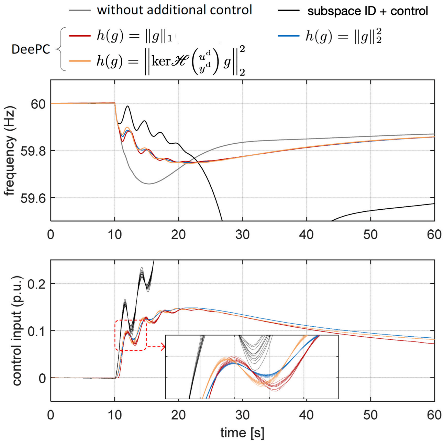

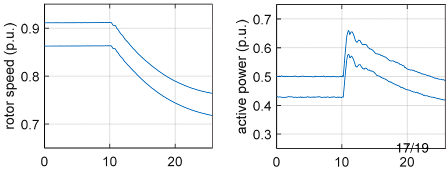

- ***high-fidelity models*** for turbines, machines, & IEEE-9-bus system
- ***fast frequency response*** via ***decentralized DeePC*** at turbines

Conclusions

main take-aways

- matrix time series as predictive model
- robustness & implicit ID via regularizations
- method that works in theory & practice for stochastic & weakly nonlinear systems
- illustrated via energy system case studies

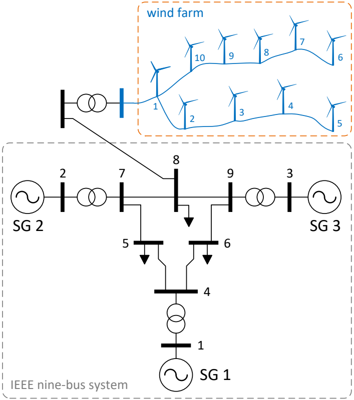

ongoing work

$\rightarrow$ certificates for truly nonlinear systems

$\rightarrow$ explicit policies & direct adaptive control

$\rightarrow$ applications with a true "business case"

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

$\rightarrow$ models are compressed, de-noised, & tidied-up representations

Thanks !

Florian Dörfler

mail: dorfler@ethz.ch

[link] to homepage

[link] to related publications

back-up slides

Performance of least-square-induced regularizer on stochastic LTI system

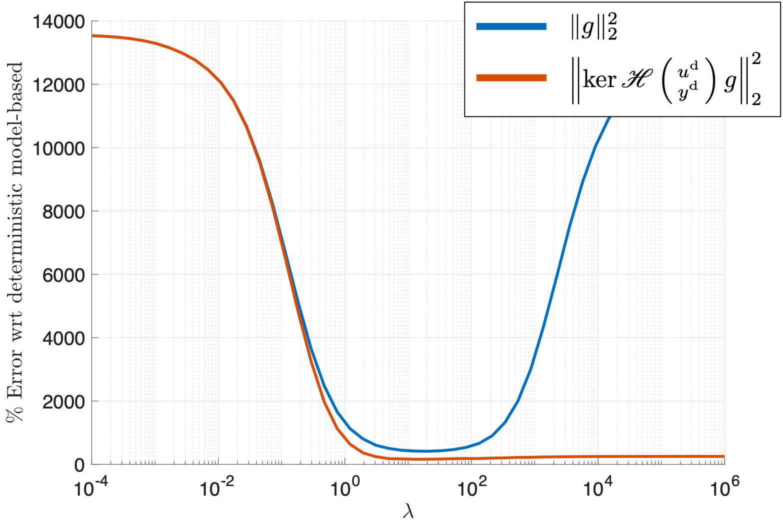

Further . . .

- ***measure concentration***: average matrix $\frac{1}{N}\sum_{i=1}^{N}\mathscr{H}_i(y^{\mathsf{d}})$ from i.i.d. experiments $\Longrightarrow$ ambiguity set $\mathbb{B}_\epsilon(\widehat{\mathbb{P}})$ includes true $\mathbb{P}$ with high confidence if $\epsilon \sim 1/N^{1/\dim(\xi)}$

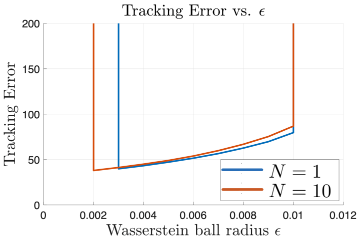

- ***distributionally robust constraints***

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

- more ***structured uncertainty sets*** : tractable reformulations (relaxations) & guarantees on realized performance

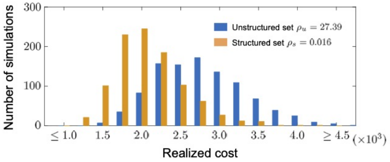

- replace (finite) moving horizon estimation via $\binom{u_{\mathrm{ini}}}{y_{\mathrm{ini}}}$ by ***recursive Kalman filtering*** based on explicit optimizer $g^\star$ as hidden state how does DeePC relate to sequential SysID + control ?

surprise: **DeePC consistently beats models** across all our case studies !

Abstraction reveals pros & cons

indirect (model-based) ***data-driven control***

minimize	control cost (u, x)	outer optimization
subject to	(u, x) 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 (black-box) ***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: account for ***uncertainty*** (hard to propagate in indirect approach)

Comparison: direct vs. indirect control

indirect* ID-based *data-driven control

minimize control cost (u, y)

subject to (u, y) satisfy parametric model

where model $\in$ argmin id cost (u^d, y^d)

subject to model $\in$ LTI(n, ℓ) class

$\Big\}$ ID

ID projects data on the set of LTI models

- with parameters (n, ℓ)
- removes noise & thus lowers variance error
- suffers bias error if plant is not LTI(n, ℓ)

direct* regularized *data-driven control

minimize control cost (u, y) + $\lambda \cdot$ regularizer

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

- ***regularization robustifies***
  $\rightarrow$ choosing λ makes it work
- ***no projection*** on LTI(n, ℓ)
  $\rightarrow$ no de-noising & no bias

hypothesis: ID wins in stochastic (variance) & DeePC in nonlinear (bias) case

Case study: direct vs. indirect control

stochastic LTI case $\rightarrow$ indirect ID wins

- LQR control of 5th order LTI system
- Gaussian noise with varying noise to signal ratio (100 rollouts each case)
- ℓ_1-regularized DeePC, SysID via N4SID, & judicious hyper-parameters

nonlinear case $\rightarrow$ direct DeePC wins

- Lotka-Volterra + control: $x^+ = f(x,u)$
- interpolated system $x^+ = \epsilon \cdot f_{\text{linearized}}(x,u) + (1-\epsilon)\cdot f(x,u)$
- same ID & DeePC as on the left & 100 initial x_0 rollouts for each ϵ