Distributed demand control in power grids
and ODEs for Markov decision processes

Second Conference on the Mathematics of Energy Markets
Wolfgang Pauli Institute, Vienna, 4-9 July 2017

Ana Bušić

Inria Paris
Département d'Informatique de l'ENS

Joint work with Prabir Barooah, Yue Chen, and Sean Meyn
University of Florida

INVENTEURS DU MONDE NUMÉRIQUE
ENS

COGNITION & CONTROL IN COMPLEX SYSTEMS
Electrical and Computer Engineering
University of Florida

Challenges of renewable power generation

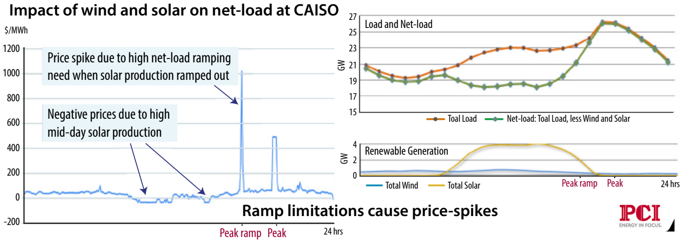

Challenges of renewable power generation

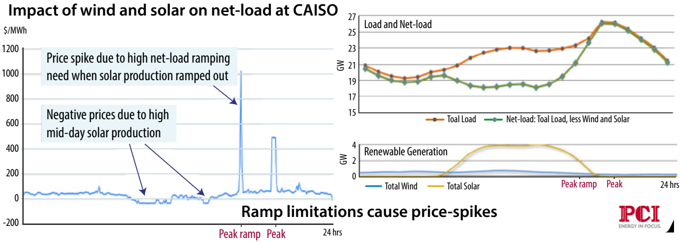

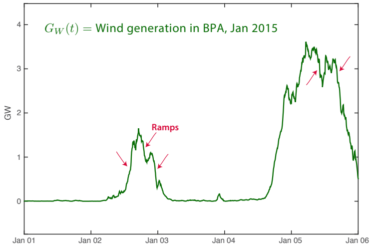

Challenges of renewable power generation

Balancing control loop

- wind and solar volatility seen as disturbance
- grid level measurements: scalar function of time (ACE) a linear combination of frequency deviation and the tie-line error (power missmatch between the sceduled and actual power out of the balancing region)
- compensation G_c designed by a balancing authority
- In many cases control loops are based on standard PI (proportional-integral) control design.

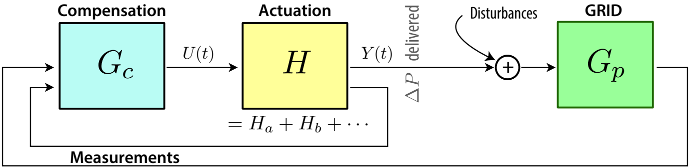

Challenges of renewable power generation

Increasing needs for ancillary services

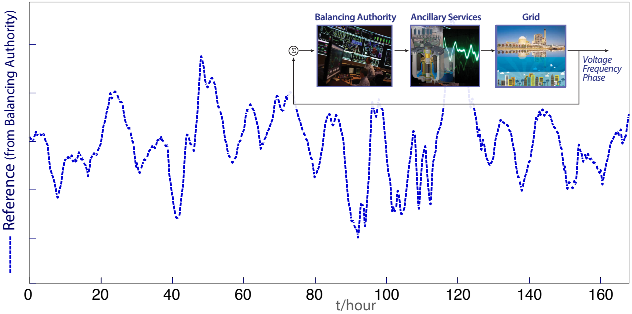

In the past, provided by the generators - high costs!

Tracking Grid Signal with Residential Loads

Tracking objective:

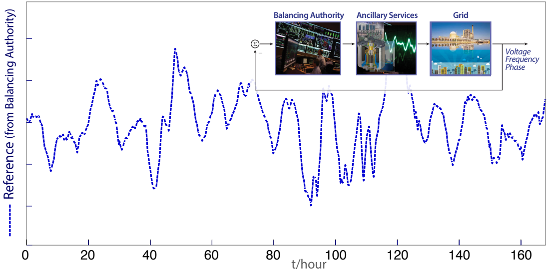

Prior work

- Deterministic centralized control:
  Sanandaji et al. 2014 [HICSS], Biegel et al. 2013 [IEEE TSG]
- Randomized control:
  Mathieu, Koch, Callaway 2013 [IEEE TPS] (decisions at the BA)
  Meyn, Barooah, B., Chen, Ehren 2015 [IEEE TAC]
  (local decisions, restricted load models)

Tracking Grid Signal with Residential Loads

Example: 20 pools, 20 kW max load

Each pool consumes 1kW when operating
12 hour cleaning cycle each 24 hours

Power Deviation:

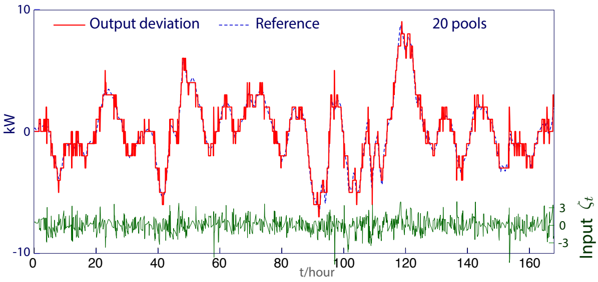

Nearly Perfect Service from Pools
Meyn, Barooah, B., Chen, Ehren 2015 [IEEE TAC]
using an extension/reinterpretation of Todorov 2007 [NIPS] (linearly solvable MDPs)

Tracking Grid Signal with Residential Loads

Example: 300,000 pools, 300 MW max load

Each pool consumes 1kW when operating
12 hour cleaning cycle each 24 hours

Power Deviation:

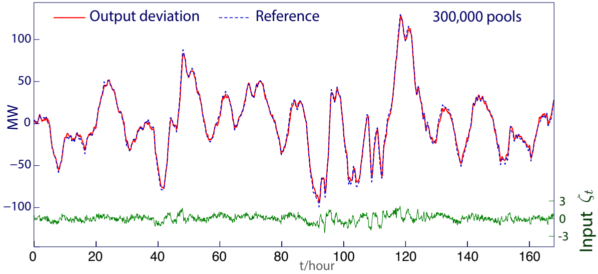

Nearly Perfect Service from Pools

Meyn, Barooah, B., Chen, Ehren 2015 [IEEE TAC]

using an extension/reinterpretation of Todorov 2007 [NIPS] (linearly solvable MDPs)

What About Other Loads?

Control Goals and Architecture

Macro control

High-level control layer: BA or a load aggregator.

The balancing challenges are of many different categories and time-scales:

- Automatic Generation Control (AGC); time scales of seconds to 20 minutes.
- Balancing reserves. In the Bonneville Power Authority, the balancing reserves include both AGC and balancing on timescales of many hours. Balancing on a slower time-scale is achieved through real time markets in some other regions of the U.S.
- Contingencies (e.g., a generator outage)
- Peak shaving
- Smoothing ramps from solar or wind generation

Control Goals and Architecture

Local Control: decision rules designed to respect needs of load and grid

Demand Dispatch: Power consumption from loads varies automatically to provide *service to the grid*, *without impacting QoS* to the consumer

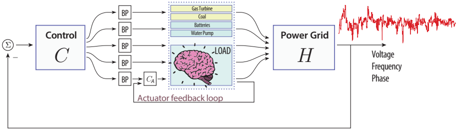

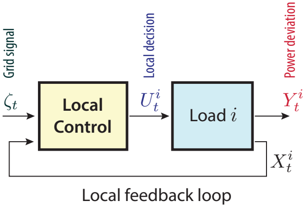

- Min. communication: each load monitors its state and a regulation signal from the grid.
- Aggregate must be controllable: **randomized policies** for finite-state loads.

Load Model

Controlled Markovian Dynamics

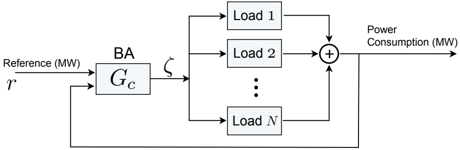

- Discrete time: ith load $X^i(t)$ evolves on finite state space X
- Each load is subject to *common* controlled Markovian dynamics.

Signal $\boldsymbol{\zeta} = \{\zeta_t\}$ is broadcast to all loads

- Controlled transition matrix $\{P_\zeta : \zeta \in \mathbb{R}\}$:

$$\mathsf{P}\{X^i_{t+1} = x' \mid X^i_t = x,\, \zeta_t = \zeta\} = P_\zeta(x, x')$$

Questions

- How to analyze aggregate of similar loads?
- Local control design?

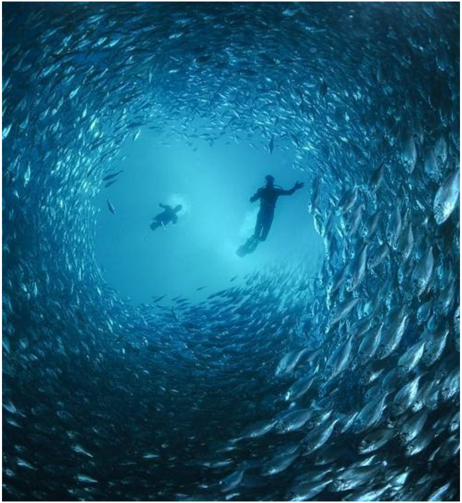

Aggregate model

How to analyze aggregate?

Mean field model

N loads running independently, each under the command ζ.

Empirical Distributions:

$$\mu_t^N(x) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}\{X^i(t) = x\}, \qquad x \in \mathsf{X}$$

$\mathcal{U}(x)$ power consumption in state x,

$$y_t^N = \frac{1}{N}\sum_{i=1}^{N} \mathcal{U}(X_t^i) = \sum_x \mu_t^N(x)\mathcal{U}(x)$$

Mean-field model:

via Law of Large Numbers for martingales

$$\mu_{t+1} = \mu_t P_{\zeta_t}, \qquad y_t = \langle \mu_t, \mathcal{U} \rangle$$

$$\zeta_t = f_t(y_0, \dots, y_t) \quad \textit{by design}$$

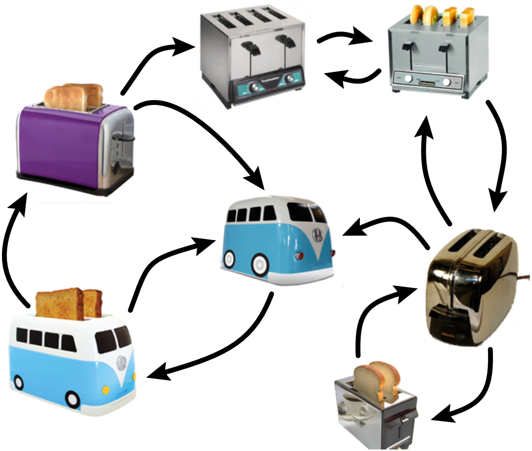

Local Control Design

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Myopic Design: $P_\zeta^{myop}(x,x') := P_0(x,x')\exp\big(\zeta\mathcal{U}(x') - \Lambda_\zeta(x)\big)$

with $\Lambda_\zeta(x) := \log\Big(\sum_{x'} P_0(x,x')\exp\big(\zeta\mathcal{U}(x')\big)\Big)$ the normalizing constant.

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Myopic Design: $P_\zeta^{myop}(x,x') := P_0(x,x')\exp\big(\zeta\mathcal{U}(x') - \Lambda_\zeta(x)\big)$
with $\Lambda_\zeta(x) := \log\Big(\sum_{x'} P_0(x,x')\exp\big(\zeta\mathcal{U}(x')\big)\Big)$ the normalizing constant.

Exponential family design: $P_\zeta(x,x') := P_0(x,x')\exp\big(h_\zeta(x,x') - \Lambda_{h_\zeta}(x)\big)$ with

$$h_\zeta(x,x') = \zeta H_0(x,x').$$

The choice of H_0 will typically correspond to the linearization of a more advanced design around the value $\zeta = 0$ (or some other fixed value of ζ).

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Individual Perspective Design

Consider a finite-time-horizon optimization problem: For a given terminal time T, let p_0 denote the pmf on strings of length T,

$$p_0(x_1, \dots, x_T) = \prod_{i=0}^{T-1} P_0(x_i, x_{i+1}),$$

where $x_0 \in \mathsf{X}$ is assumed to be given. The scalar $\zeta \in \mathbb{R}$ is interpreted as a weighting parameter in the following definition of total welfare. For any pmf p,

$$\mathcal{W}_T(p) = \zeta \mathsf{E}_p\Big[\sum_{t=1}^{T} \mathcal{U}(X_t)\Big] - D(p \| p_0)$$

where the expectation is with respect to p, and D denotes relative entropy:

$$D(p\|p_0) := \sum_{x_1,\dots,x_T} \log\Big(\frac{p(x_1,\dots,x_T)}{p_0(x_1,\dots,x_T)}\Big) p(x_1,\dots,x_T)$$

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

It is easy to check that the myopic design is an optimizer for the horizon $T = 1$,

$$P_\zeta^{myop}(x_0, \cdot) \in \arg\max_p \mathcal{W}_1(p).$$

The infinite-horizon mean welfare is denoted,

$$\eta_\zeta^* = \lim_{T \to \infty} \frac{1}{T} \mathcal{W}_T(p_T^*)$$

Explicit construction via eigenvector problem:

$$P_\zeta(x, y) = \frac{1}{\lambda} \frac{v(y)}{v(x)} \hat{P}_\zeta(x, y), \qquad x, y \in \mathsf{X},$$

where $\quad \hat{P}_\zeta v = \lambda v, \qquad \hat{P}_\zeta(x, y) = \exp(\zeta \mathcal{U}(x)) P_0(x, y)$

Extension/reinterpretation of [Todorov 2007] + [Kontoyiannis & Meyn 200X]

Example: pool pumps

How Pools Can Help Regulate The Grid

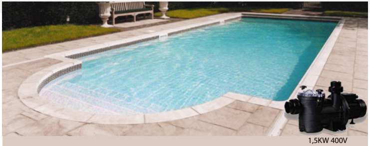

Needs of a single pool

- Filtration system circulates and cleans: Average pool pump uses 1.3kW and runs 6-12 hours per day, 7 days per week
- Pool owners are oblivious, until they see *frogs and algae*
- Pool owners do not trust anyone: *Privacy is a big concern*

Single pool dynamics: $\mathsf{X} = \{(m, j) : m \in \{0, 1\},\ j \in \{1, 2, \ldots, \mathcal{I}\}\}$.

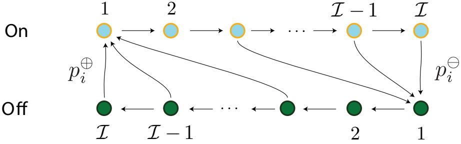

Tracking Grid Signal with Residential Loads

Example: 20 pools, 20 kW max load

Each pool consumes 1kW when operating
12 hour cleaning cycle each 24 hours

Power Deviation:

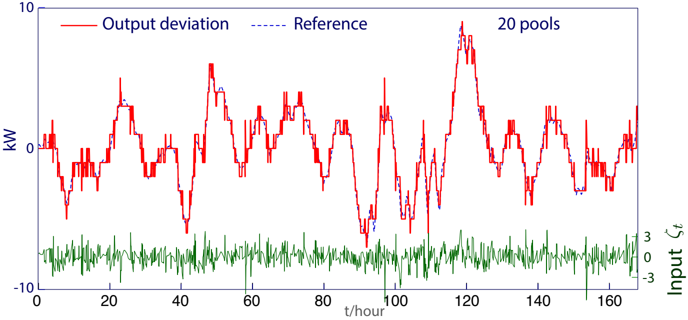

Nearly Perfect Service from Pools
Meyn et al. 2013 [CDC], Meyn et al. 2015 [IEEE TAC]

Tracking Grid Signal with Residential Loads

Example: 300,000 pools, 300 MW max load

Each pool consumes 1kW when operating
12 hour cleaning cycle each 24 hours

Power Deviation:

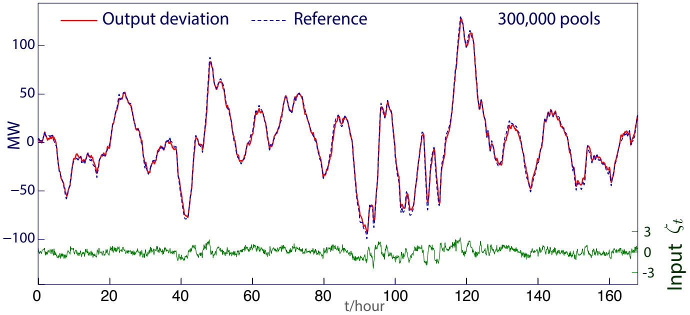

Nearly Perfect Service from Pools
Meyn et al. 2013 [CDC], Meyn et al. 2015 [IEEE TAC]

Range of services provided by pools

Example: 10,000 pools, 10 MW max load

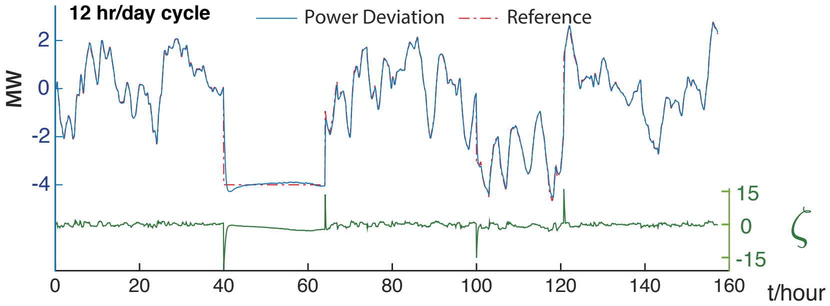

Local Design

Extending local control design to include exogenous disturbances

State space for a load model: $\mathsf{X} = \mathsf{X}_u \times \mathsf{X}_n$.

Components X_n are not subject to direct control
(e.g. impact of the weather on the climate of a building).

Local Design

Extending local control design to include exogenous disturbances

State space for a load model: $\mathsf{X} = \mathsf{X}_\mathsf{u} \times \mathsf{X}_\mathsf{n}$.

Components X_n are not subject to direct control
(e.g. impact of the weather on the climate of a building).

Conditional-independence structure of the local transition matrix

$$P(x, x') = R(x, x'_u) Q_0(x, x'_n), \quad x' = (x'_u, x'_n)$$

Q_0 models uncontroled load dynamics and exogenous disturbances.

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Nominal model

A Markovian model for an individual load, based on its typical behavior.

- Finite state space $\mathsf{X} = \{x^1, \dots, x^d\}$;
- Transition matrix P_0, with unique invariant pmf π_0.

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Nominal model

A Markovian model for an individual load, based on its typical behavior.

- Finite state space $\mathsf{X} = \{x^1, \dots, x^d\}$;
- Transition matrix P_0, with unique invariant pmf π_0.

Common structure for design

The family of transition matrices used for distributed control is of the form:

$$P_\zeta(x, x') := P_0(x, x') \exp\bigl(h_\zeta(x, x') - \Lambda_{h_\zeta}(x)\bigr)$$

with h_ζ continuously differentiable in ζ, and the normalizing constant

$$\Lambda_{h_\zeta}(x) := \log\Bigl(\sum_{x'} P_0(x, x') \exp\bigl(h_\zeta(x, x')\bigr)\Bigr)$$

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Nominal model

A Markovian model for an individual load, based on its typical behavior.

- Finite state space $\mathsf{X} = \{x^1, \dots, x^d\}$;
- Transition matrix P_0, with unique invariant pmf π_0.

Common structure for design

The family of transition matrices used for distributed control is of the form:

$$P_\zeta(x, x') := P_0(x, x') \exp\big(h_\zeta(x, x') - \Lambda_{h_\zeta}(x)\big)$$

with h_ζ continuously differentiable in ζ, and the normalizing constant

$$\Lambda_{h_\zeta}(x) := \log\Big(\sum_{x'} P_0(x, x') \exp\big(h_\zeta(x, x')\big)\Big)$$

Assumption: for all $x \in \mathsf{X}$, $x' = (x'_u, x'_n) \in \mathsf{X}$, $h_\zeta(x, x') = h_\zeta(x, x'_u)$.

Local Design

Goal: Construct a family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$

Construction of the family of functions $\{h_\zeta : \zeta \in \mathbb{R}\}$

Step 1: The specification of a function $\mathcal{H}$ that takes as input a transition matrix. $H = \mathcal{H}(P)$ is a real-valued function on $\mathsf{X} \times \mathsf{X}$.

Step 2: The families $\{P_\zeta\}$ and $\{h_\zeta\}$ are defined by the solution to the ODE:

$$\tfrac{d}{d\zeta} h_\zeta = \mathcal{H}(P_\zeta), \qquad \zeta \in \mathbb{R},$$

in which P_ζ is determined by h_ζ through:

$$P_\zeta(x, x') := P_0(x, x') \exp\big(h_\zeta(x, x') - \Lambda_{h_\zeta}(x)\big)$$

The boundary condition: $h_0 \equiv 0$.

Local Design

Extending local control design to include exogenous disturbances

For any function $H^\circ \colon \mathsf{X} \to \mathbb{R}$, one can define

$$H(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) H^\circ(x'_u, x'_n) \qquad (1)$$

Then functions $\{h_\zeta\}$ satisfy

$$h_\zeta(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) h^\circ_\zeta(x'_u, x'_n),$$

for some $h^\circ_\zeta \colon \mathsf{X} \to \mathbb{R}$. Moreover, these functions solve the d-dimensional ODE,

$$\tfrac{d}{d\zeta} h^\circ_\zeta = \mathcal{H}^\circ(P_\zeta), \qquad \zeta \in \mathbb{R},$$

with boundary condition $h^\circ_0 \equiv 0$.

Individual Perspective Design

- Local welfare function: $\mathcal{W}_\zeta(x,P) = \zeta\mathcal{U}(x) - D(P\|P_0)$,
  where D denotes relative entropy: $D(P\|P_0) = \sum_{x'} P(x,x') \log\big(\frac{P(x,x')}{P_0(x,x')}\big)$.
- Markov Decision Process: $\limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} E[\mathcal{W}_\zeta(X_t,P)]$
  Average reward optimization equation (AROE):

$$\max_P \Big\{ \mathcal{W}_\zeta(x,P) + \sum_{x'} P(x,x') h^*_\zeta(x') \Big\} = h^*_\zeta(x) + \eta^*_\zeta$$

where $P(x,x') = R(x,x'_u)Q_0(x,x'_n), \quad x' = (x'_u, x'_n)$

Individual Perspective Design

- ODE method for IPD design:

  Family $\{P_\zeta\}$: $P_\zeta(x,x') := P_0(x,x') \exp\bigl(h_\zeta(x,x') - \Lambda_{h_\zeta}(x)\bigr)$
  Functions $\{h_\zeta\}$: $h_\zeta(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) h^\circ_\zeta(x'_u, x'_n)$,
  for $h^\circ_\zeta \colon \mathsf{X} \to \mathbb{R}$ solutions of the d-dimensional ODE,

  $$\frac{d}{d\zeta} h^\circ_\zeta = \mathcal{H}^\circ(P_\zeta), \qquad \zeta \in \mathbb{R},$$

  with boundary condition $h^\circ_0 \equiv 0$.

  $H^\circ_\zeta(x) = \frac{d}{d\zeta} h^\circ_\zeta(x) = \sum_{x'} [Z_\zeta(x,x') - Z_\zeta(x^\circ, x')]\mathcal{U}(x'), \quad x \in \mathsf{X},$

  where $Z = [I - P + 1 \otimes \pi]^{-1} = \sum_{n=0}^\infty [P_\zeta - 1 \otimes \pi]^n$ is the fundamental matrix.

Example: Thermostatically Controlled Loads

- refrigerators, water heaters, air-conditioning . . .
- TCLs are already equipped with primitive "local intelligence" based on a *deadband* (or *hysteresis interval*)
- The state process for a TCL at time t:

$$X(t) = (X_u(t), X_n(t)) = (m(t), \Theta(t)) \,,$$

where $m(t) \in \{0, 1\}$ denotes the power mode ("1" indicating the unit is on), and $\Theta(t)$ the inside temperature of the load

Exogenous disturbances: ambient temperature, and usage

Example: Thermostatically Controlled Loads

The standard ODE model of a water heater is the first-order linear system,

$$\frac{d}{dt}\Theta(t) = -\lambda[\Theta(t) - \Theta^a(t)] + \gamma m(t) - \alpha[\Theta(t) - \Theta^{in}(t)]f(t)\,,$$

- $\Theta(t)$ temperature of the water in the tank
- $\Theta^{in}(t)$ temperature of the cold water entering the tank
- $f(t)$ flow rate of hot water from the WH
- $m(t)$ power mode of the WH ("on" indicated by $m(t) = 1$).

Deterministic deadband control: $\Theta(t) \in [\Theta_-, \Theta_+]$

Nominal model for local control design: based on the specification of two CDFs for the temperature at which the load turns on or turns off

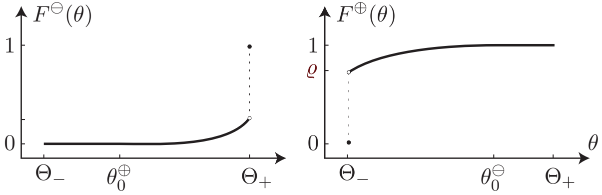

Example: Thermostatically Controlled Loads

Discrete-time control.

- At time instance k, if the water heater is on (i.e., $m(k) = 1$), then it turns off with probability,

$$p^{\ominus}(k+1) = \frac{[F^{\ominus}(\Theta(k+1)) - F^{\ominus}(\Theta(k))]_+}{1 - F^{\ominus}(\Theta(k))}$$

  where $[x]_+ := \max(0, x)$ for $x \in \mathbb{R}$;
- Similarly, if the load is off, then it turns on with probability

$$p^{\oplus}(k+1) = \frac{[F^{\oplus}(\Theta(k)) - F^{\oplus}(\Theta(k+1))]_+}{F^{\oplus}(\Theta(k))}$$

  The nominal behavior of the power mode can be expressed

$$\mathsf{P}\{m(k) = 1 \mid \theta(k-1), \theta(k), m(k-1) = 0\} = p^{\oplus}(k)$$
$$\mathsf{P}\{m(k) = 0 \mid \theta(k-1), \theta(k), m(k-1) = 1\} = p^{\ominus}(k)$$

Example: Thermostatically Controlled Loads

Myopic design - exponential tilting of these distributions:

$$p_\zeta^\oplus(k) := \mathsf{P}\{m(k) = 1 \mid \theta(k-1), \theta(k), m(k-1) = 0, \zeta(k-1) = \zeta\}$$
$$= \frac{p^\oplus(k)e^\zeta}{p^\oplus(k)e^\zeta + 1 - p^\oplus(k)}$$

$$p_\zeta^\ominus(k) = \mathsf{P}\{m(k) = 0 \mid \theta(k-1), \theta(k), m(k-1) = 1, \zeta(k-1) = \zeta\}$$
$$= \frac{p^\ominus(k)}{p^\ominus(k) + (1 - p^\ominus(k))e^\zeta}$$

If $p_0^\oplus(k) > 0$, then the probability $p_\zeta^\oplus(k)$ is strictly increasing in ζ, approaching 1 as $\zeta \to \infty$; it approaches 0 as $\zeta \to -\infty$, if $p_0^\oplus(k) < 1$.

Example: Thermostatically Controlled Loads

System identification

$$\frac{d}{dt}\Theta(t) = -\lambda[\Theta(t) - \Theta^a(t)] + \gamma m(t) - \alpha[\Theta(t) - \Theta^{in}(t)]f(t)\,,$$

- $\Theta(t)$ temperature of the water in the tank
- $\Theta^{in}(t)$ temperature of the cold water entering the tank
- $f(t)$ flow rate of hot water from the WH
- $m(t)$ power mode of the WH ("on" indicated by $m(t) = 1$).

Temp. Ranges	ODE Pars.	Loc. Control
$\Theta_+ \in [118, 122]$ F	$\lambda \in [8, 12.5] \times 10^{-6}$	$T_s = 15$ sec
$\Theta_- \in [108, 112]$ F	$\gamma \in [2.6, 2.8] \times 10^{-2}$	$\kappa = 4$
$\Theta^a \in [68, 72]$ F	$\alpha \in [6.5, 6.7] \times 10^{-2}$	$\varrho = 0.8$
$\Theta^{in} \in [68, 72]$ F	$P_{\text{on}} = 4.5$ kW	$\theta_0 = \Theta_-$

Heterogeneous population: 100 000 WHs simulated by uniform sampling of the values in the table

Usage data from Oakridge National Laboratory (35WHs over 50 days)

Tracking performance

and the controlled dynamics for an individual load

100,000 water-heaters
When on, individual load consumes $4, 5$ kW
With no usage, approx. 2% duty cycle, avg. power consumption 10MW.

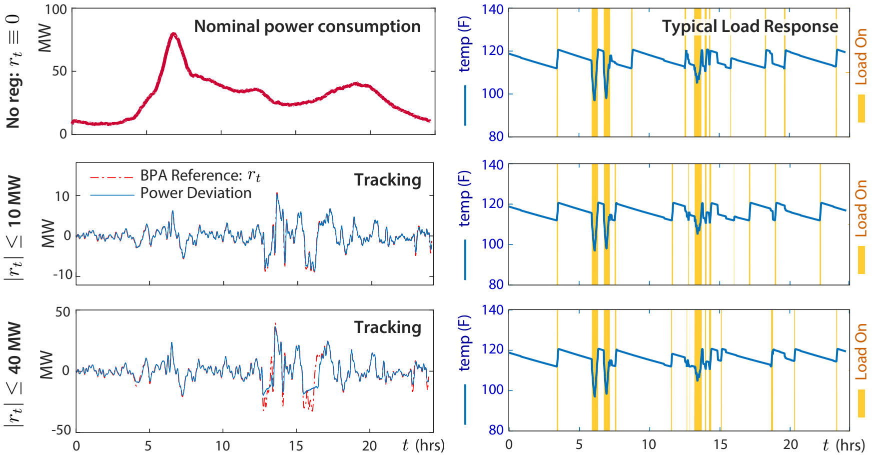

Tracking performance

Potential for contingency reserves and ramping

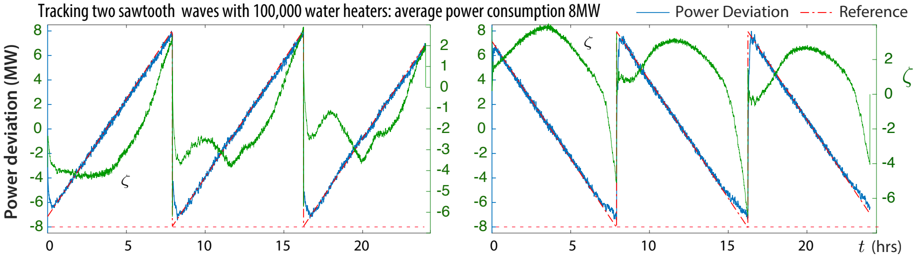

Tracking performance

and the controlled dynamics for an individual load

Heterogeneous setting:

- 40 000 loads per experiment;
- 20 different load types in each case

Lower plots show the on/off state for a typical load

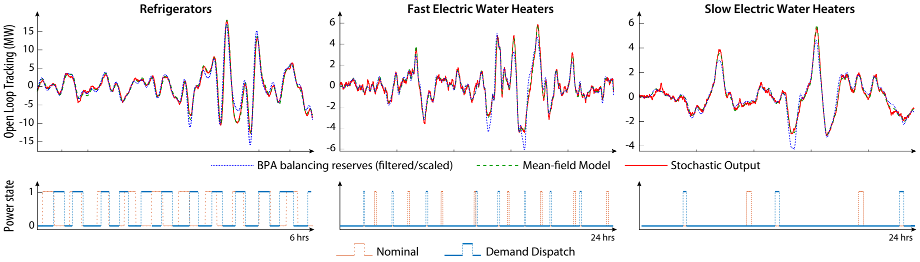

Unmodeled dynamics

Setting: 0.1% sampling, and

1. *Heterogeneous* population of loads
2. Load i overrides when QoS is out of bounds

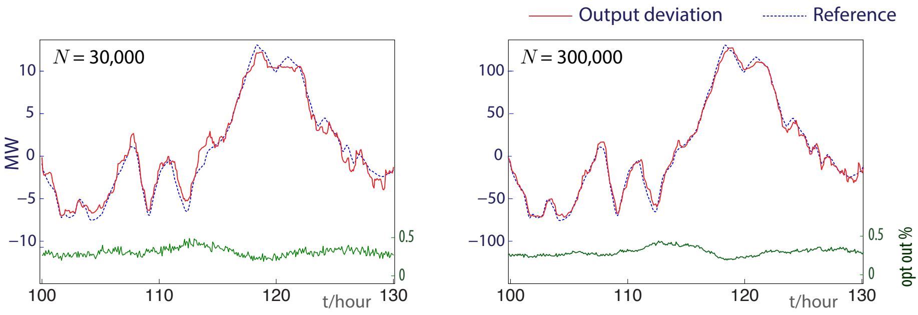

Closed-loop tracking

PI control: $\zeta_t = k_P e_t + k_I e^I_t, \quad e_t = r_t - y_t, \quad e^I_t = \sum_{s=0}^t e_s$

Control Architecture

Frequency *Allocation* for Demand Dispatch

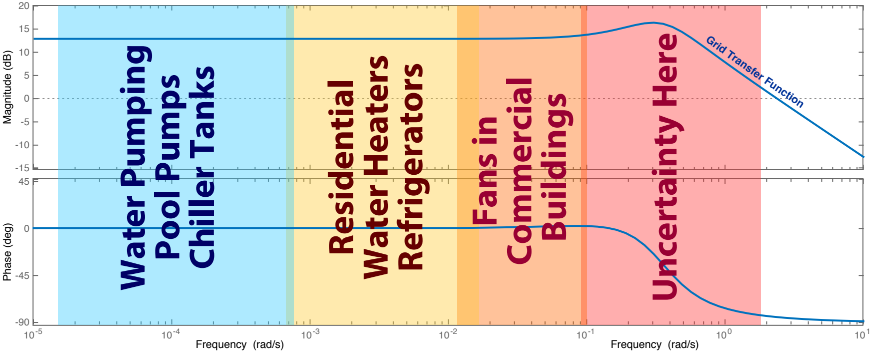

10,000 pools

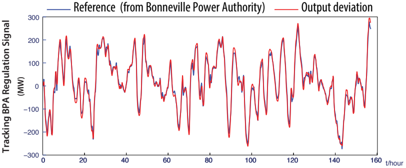

Bandwidth
centered around
its natural cycle

Conclusions

Virtual storage from flexible loads

Approach: creating Virtual Energy Storage through direct control of flexible loads - helping the grid while respecting user QoS

Conclusions

Virtual storage from flexible loads

Approach: creating Virtual Energy Storage through direct control of flexible loads - helping the grid while respecting user QoS

Challenges:

- Stability properties for IPD and myopic design?
- Information Architecture: $\zeta_t = f(?)$
  Different needs for communication, state estimation and forecast.
- Capacity estimation (time varying)
- Network constraints
- Resource optimization & learning
  Integrating VES with traditional generation and batteries.
- Economic issues
  Contract design, aggregators, markets ...

Conclusions

Thank You!

References: this talk

- A. Bušić and S. Meyn. Distributed randomized control for demand dispatch. *55th IEEE Conference on Decision and Control*, 2016.
- A. Bušić and S. Meyn. Ordinary Differential Equation Methods For Markov Decision Processes and Application to Kullback-Leibler Control Cost. arXiv:1605.04591v2. Oct 2016.
- S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren. Ancillary Service to the Grid Using Intelligent Deferrable Loads. *IEEE Trans. Automat. Contr.*, 60(11): 2847-2862, 2015.
- P. Barooah, A. Bušić, and S. Meyn. Spectral Decomposition of Demand-Side Flexibility for Reliable Ancillary Services in a Smart Grid. *48th Annual Hawaii International Conference on System Sciences (HICSS).* 2015.
- A. Bušić and S. Meyn. Passive dynamics in mean field control. *53rd IEEE Conf. on Decision and Control (CDC)* 2014.

References: related

Demand dispatch:

- Y. Chen, A. Bušić, and S. Meyn. Individual risk in mean-field control models for decentralized control, with application to automated demand response. *53rd IEEE Conf. on Decision and Control (CDC)*, 2014.
- Y. Chen, A. Bušić, and S. Meyn. State Estimation and Mean Field Control with Application to Demand Dispatch. *54rd IEEE Conference on Decision and Control (CDC)* 2015.
- J. L. Mathieu. Modeling, Analysis, and Control of Demand Response Resources. PhD thesis, Berkeley, 2012.
- J. L. Mathieu, S. Koch, D. S. Callaway, State Estimation and Control of Electric Loads to Manage Real-Time Energy Imbalance, *IEEE Transactions on Power Systems*, 28(1):430-440, 2013.

Markov processes:

- I. Kontoyiannis and S. P. Meyn. Spectral theory and limit theorems for geometrically ergodic Markov processes. *Ann. Appl. Probab.*, 13:304–362, 2003.
- I. Kontoyiannis and S. P. Meyn. Large deviations asymptotics and the spectral theory of multiplicatively regular Markov processes. *Electron. J. Probab.*, 10(3):61–123 (electronic), 2005.
- E. Todorov. Linearly-solvable Markov decision problems. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems*, (19) 1369–1376. MIT Press, Cambridge, MA, 2007.

Mean Field Model

Linearized Dynamics

Mean-field model: $$\mu_{t+1} = \mu_t P_{\zeta_t}, \qquad y_t = \langle \mu_t, \mathcal{U} \rangle$$

$$\zeta_t = f_t(y_0, \dots, y_t)$$

Linear state space model:

$$\Phi_{t+1} = A\Phi_t + B\zeta_t$$
$$\gamma_t = C\Phi_t$$

Interpretations: $|\zeta_t|$ is small, and π denotes invariant measure for P_0.

- $\Phi_t \in \mathbb{R}^{|\mathsf{X}|}$, a column vector with
  $$\Phi_t(x) \approx \mu_t(x) - \pi(x),\ x \in \mathsf{X}$$
- $\gamma_t \approx y_t - y^0$; deviation from nominal steady-state
- $A = P_0^{\intercal}$, $C = \mathcal{U}^{\intercal}$, and input dynamics linearized:

$$B^{\intercal} = \frac{d}{d\zeta} \pi P_\zeta \Big|_{\zeta=0}$$