Distributed control design for balancing the grid using flexible loads

caesars2018

5-7 Sep 2018, EDF Lab Paris-Saclay

Ana Bušić

Inria Paris
DI ENS, PSL Research University

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

inventors for the digital world

COGNITION & CONTROL IN COMPLEX SYSTEMS

Electrical and Computer Engineering
University of Florida

Challenges of renewable power generation

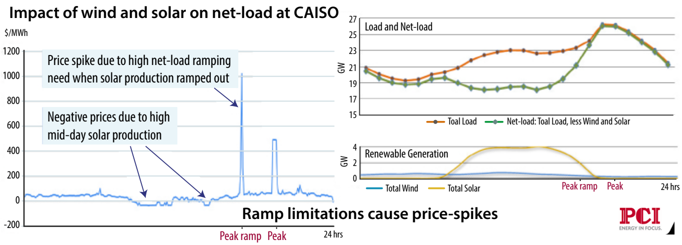

Challenges of renewable power generation

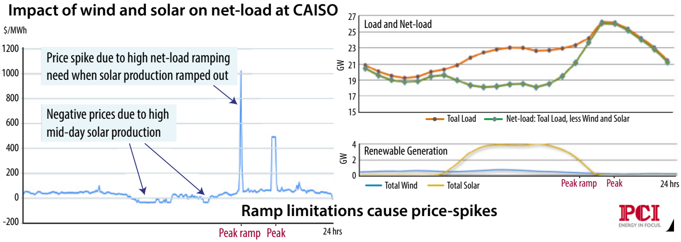

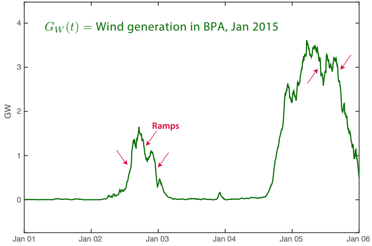

Challenges: ancillary services

... to compensate for energy imbalances ... Kirby 2013

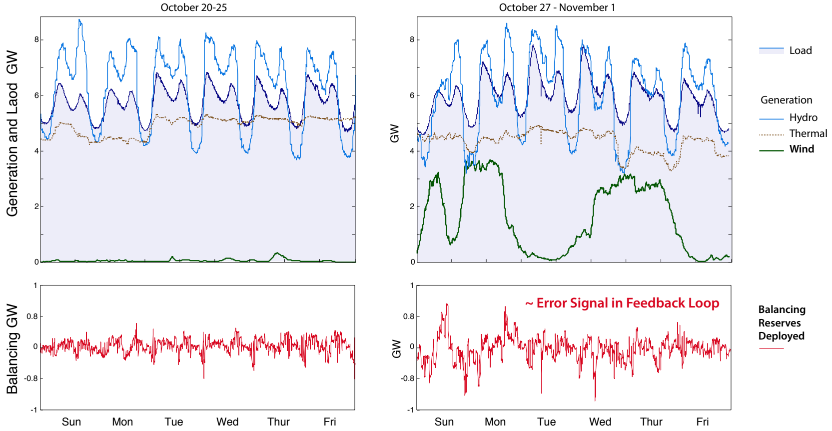

Challenges: ancillary services

... to compensate for energy imbalances ... Kirby 2013

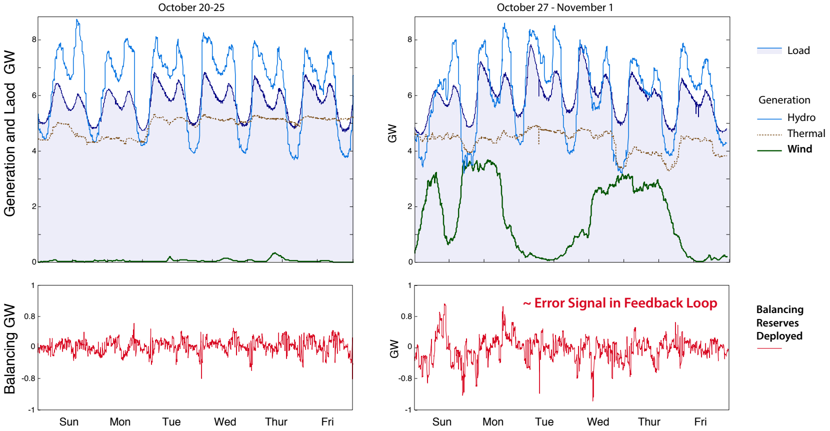

Time-scale of power deviations are similar to secondary reserves following a fault

Challenges: ancillary services

... to compensate for energy imbalances ... Kirby 2013

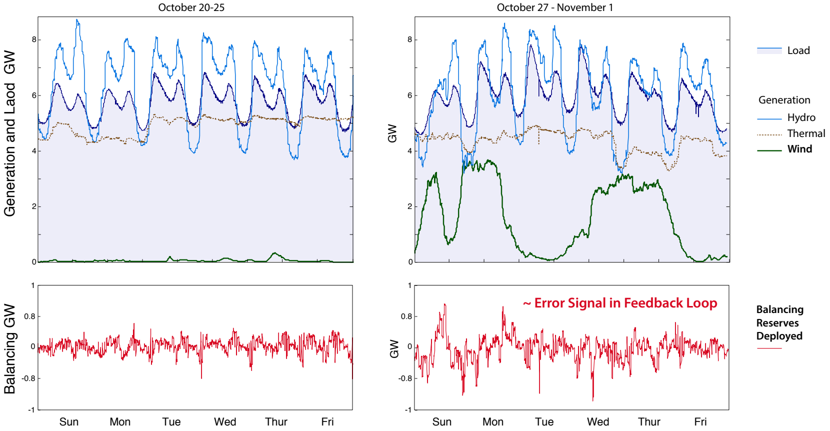

Time-scale of power deviations are similar to secondary reserves following a fault

The Balancing Reserves at BPA are a sum of many error signals in the grid

Secondary Control

Balancing Authority

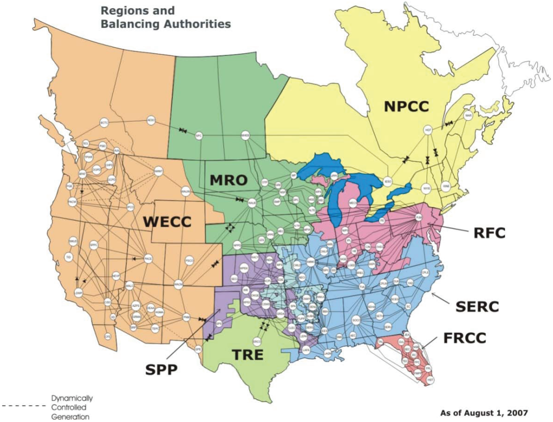

Figure 2 North American Balancing Authorities and Regions

Secondary Control

Balancing Authority

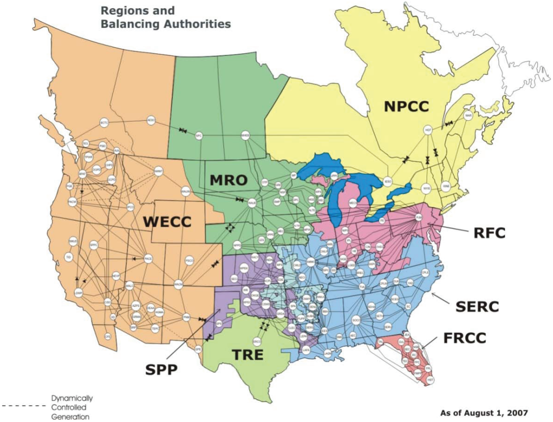

Figure 2 North American Balancing Authorities and Regions

- Transmission lines that join two areas are known as tie-lines.

Secondary Control

Balancing Authority

- Transmission lines that join two areas are known as tie-lines.
- The net power out of an area is the sum of the flow on its tie-lines.

Secondary Control

Balancing Authority

- Transmission lines that join two areas are known as tie-lines.
- The net power out of an area is the sum of the flow on its tie-lines.
- The flow out of an area is equal to

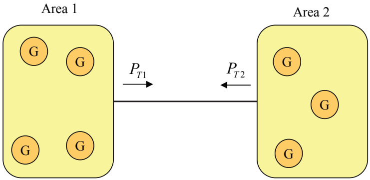

total gen - total load - total losses = tie-line flow

Secondary Control

Area Control Error

Area Control Error: combination of:

- Deviation of frequency from nominal, and

Secondary Control

Area Control Error

Area Control Error: combination of:

- Deviation of frequency from nominal, and
- the difference between the actual tie-line flow, and the scheduled flow (from economic dispatch).

Secondary Control

Area Control Error

Area Control Error: combination of:

- Deviation of frequency from nominal, and
- the difference between the actual tie-line flow, and the scheduled flow (from economic dispatch).

$$\text{ACE} = P_{\text{actual}} - P_{\text{scheduled}} + B\Delta\omega$$

B is the bias

Secondary Control

Area Control Error

Area Control Error: combination of:

- Deviation of frequency from nominal, and
- the difference between the actual tie-line flow, and the scheduled flow (from economic dispatch).

$$\text{ACE} = P_{\text{actual}} - P_{\text{scheduled}} + B\Delta\omega$$

B is the bias

Provides a measure of whether an area is producing more or less than it should to satisfy schedules and to contribute to controlling frequency.

AGC: control signal designed to bring ACE to zero.

Secondary Control

Area Control Error

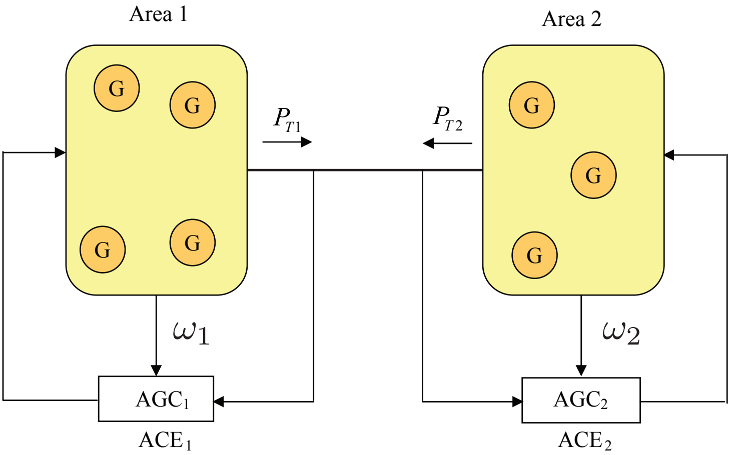

P_{T1} = Tie line power for Area 1 = $\sum_{j\in\Omega_1} P_{T1}^j$ = Sum over all tie lines

ETH Dynamics 2012

Primary and Secondary Control Loops

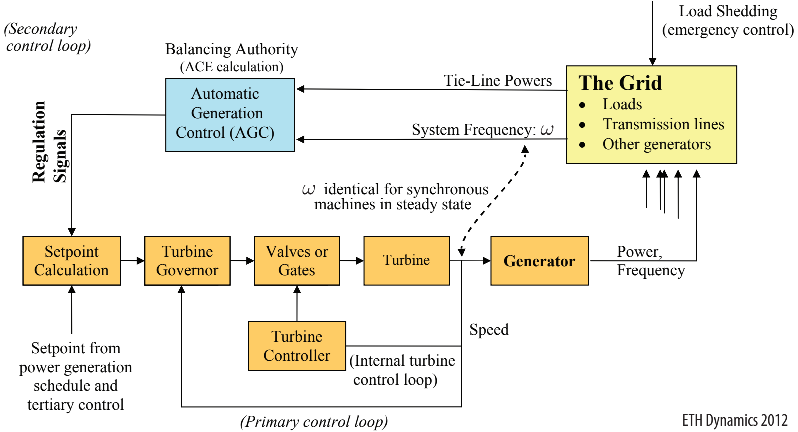

Balancing control loop

- wind and solar volatility seen as disturbance
- grid level measurements: scalar function of time (ACE)
- compensation G_c designed by a balancing authority
- In many cases control loops are based on standard PI (proportional-integral) control design.

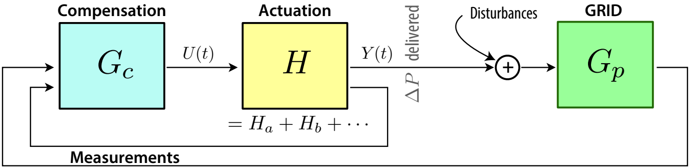

Secondary Control

Regulation signals

Regulation: on-line generation, responsive load and storage ... helps to maintain interconnection frequency, manage differences between actual and scheduled power flows between balancing areas, and match generation to load within the balancing area.

Automatic Generation Control (AGC): commands are typically sent about every four seconds.

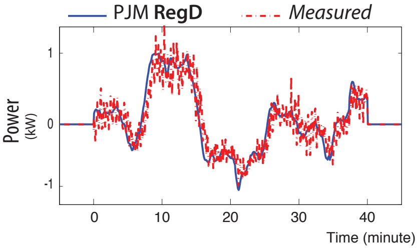

Secondary Control

Regulation signals

Regulation: on-line generation, responsive load and storage ... helps to maintain interconnection frequency, manage differences between actual and scheduled power flows between balancing areas, and match generation to load within the balancing area.

Automatic Generation Control (AGC): commands are typically sent about every four seconds.

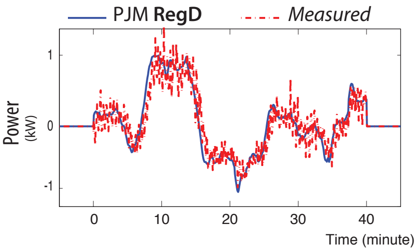

PJM regulation metrics: scores for correlation, precision, and performance. Faster and more accurate response is paid more.

Example: ACE and regulation signals at ERCOT

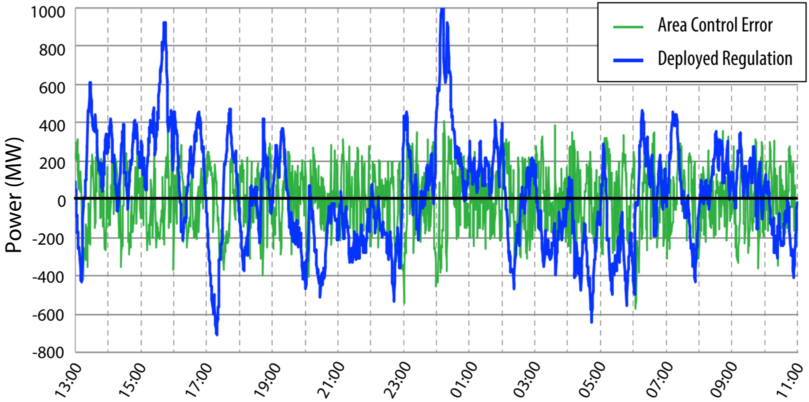

Does ACE Work?

Generators are not always good actuators

Fig. 10. Coal-fired generators do not follow regulation signals precisely....
Some do better than others

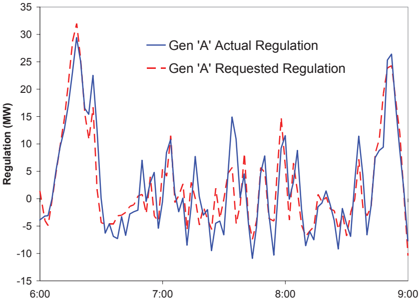

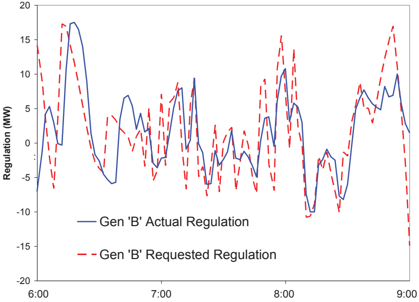

ODE method for MDPs

Markov Decision Process with:

- finite state space X; general action space U,
- $P(X(t+1) = x' \mid X(t) = x, U(t) = u) = P_u(x, x')$,
- one-step reward $w \colon \mathsf{X} \times \mathsf{U} \to \mathbb{R}$.

Two standard optimal control criteria are *finite-horizon*:

$$\mathcal{W}_T^*(x) = \max \sum_{t=0}^{T} \mathsf{E}[w(X(t), U(t)) \mid X(0) = x]$$

where $T \geq 0$ is fixed, and *average reward*:

$$\eta^*(x) = \max \Big\{ \liminf_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathsf{E}[w(X(t), U(t)) \mid X(0) = x] \Big\} .$$

The maximum is over all admissible input sequences $\boldsymbol{U} = \{U(t) : t \geq 0\}$; obtained as deterministic state feedback under general conditions.

Linearly solvable MDP

(Todorov 2007)

- The action space U: all probability mass functions on X,

$$\mathsf{P}\{X(t+1) = x' \mid X(t) = x, U(t) = \mu\} = \mu(x')\,, \quad x, x' \in \mathsf{X},\ \mu \in \mathsf{U}\,.$$

- The one-step reward is defined as the sum of two terms:

$$w(x,\mu) = \mathcal{U}(x) - c_{\mathsf{KL}}(x,\mu).$$

The second term is a "control cost", defined using Kullback–Leibler divergence from nominal (control-free) transition matrix P_0

$$c_{\mathsf{KL}}(x,\mu) = D(\mu \| P_0(x,\,\cdot\,)) := \sum_{x'} \mu(x') \log\Bigl(\frac{\mu(x')}{P_0(x,x')}\Bigr).$$

Linearly solvable MDP

The solution with respect to the average reward criterion is obtained as the solution to an eigenvector problem:

- let (λ, v) denote the Perron-Frobenius eigenvalue-eigenvector pair for the positive matrix

$$\widehat{P}(x, x') = \exp(\mathcal{U}(x)) P_0(x, x'), \quad x, x' \in \mathsf{X}.$$

  The "twisted" matrix

$$\check{P}(x, x') = \frac{1}{\lambda} \frac{v(x')}{v(x)} \widehat{P}(x, x'), \quad x, x' \in \mathsf{X},$$

  is a transition matrix on X. This matrix $\check{P}$ defines the dynamics of the model under optimal control.

Linearly solvable MDP

The solution with respect to the average reward criterion is obtained as the solution to an eigenvector problem:

- let (λ, v) denote the Perron-Frobenius eigenvalue-eigenvector pair for the positive matrix

$$\widehat{P}(x,x') = \exp(\mathcal{U}(x))P_0(x,x'), \quad x,x' \in \mathsf{X}.$$

The "twisted" matrix

$$\check{P}(x,x') = \frac{1}{\lambda}\frac{v(x')}{v(x)}\widehat{P}(x,x')\,, \quad x,x' \in \mathsf{X}\,,$$

is a transition matrix on X. This matrix $\check{P}$ defines the dynamics of the model under optimal control.

Limitations: in most applications, exogenous disturbances we cannot directly control!

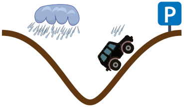

Main results

- New ODE approach for solving an entire family of MDP problems, parameterized by a scalar ζ:

$$w_\zeta(x,\mu) = \zeta\mathcal{U}(x) - c_{\textsf{KL}}(x,\mu).$$

Computational tool and tradeoff between reward and control effort.

Main results

- New ODE approach for solving an entire family of MDP problems, parameterized by a scalar ζ:

$$w_\zeta(x,\mu) = \zeta\mathcal{U}(x) - c_{\mathsf{KL}}(x,\mu).$$

Computational tool and tradeoff between reward and control effort.

- Extension to constrained actions:
State space: $\mathsf{X} = \mathsf{X}_\mathsf{u} \times \mathsf{X}_\mathsf{n}$ (control and exogenous dynamics);
$X(t) = (X_u(t), X_n(t))$.

Conditional-independence assumption: $X_n(t+1)$ is conditionally independent of the input at time t, given the value of $X(t)$:

$$P(x,x') = R(x,x'_u)Q_0(x,x'_n), \quad x,x' \in \mathsf{X}$$

  - R : randomized decision rule for $X_u(t+1)$ given $X(t) = x$,
  - Q_0 : distribution of $X_n(t+1)$ given $X(t) = x$.

Action constrained MDPs with KL-cost

The two optimal control problems can be transformed as:

$$\begin{aligned} \mathcal{W}_T^*(x,\zeta) &= \max \sum_{t=0}^{T} \mathsf{E}_x[w(X(t),R(t))] \\ \eta^*(\zeta) &= \max\Big\{ \liminf_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathsf{E}_x[w(X(t),R(t))] \Big\} \end{aligned}$$

where the maximum is over sequences of randomized decision rules $\{R(0),\dots,R(T)\}$,

$$w(x,R) := \zeta\mathcal{U}(x) - c_{\mathsf{KL}}(x,R)$$

and $$c_{\mathsf{KL}}(x,R) := \sum_{x'} P(x,x') \log\Big(\frac{P(x,x')}{P_0(x,x')}\Big) = \sum_{x'_u} R(x,x'_u) \log\Big(\frac{R(x,x'_u)}{R_0(x,x'_u)}\Big)$$

Average reward

Assumptions:

- P_0 irreducible and aperiodic;
- Action-constrained case: $P(x, x') = R(x, x'_u)Q_0(x, x'_n), \; x, x' \in \mathsf{X}$;
- Reward: $w(x, R) = \zeta\mathcal{U}(x) - c_{\text{KL}}(x, R)$.

Average reward fixed point equation:

$$\max_R \Big\{ w(x, R) + \sum_{x'} P(x, x') h^*_\zeta(x') \Big\} = h^*_\zeta(x) + \eta^*(\zeta)$$

Convention: $h^*_\zeta(x^\circ) = 0$, for $x^\circ \in \mathsf{X}$ is a fixed state.

The maximizer defines a transition matrix:

$$\check{P}_\zeta = \arg\max_P \big\{ \zeta\pi(\mathcal{U}) - K(P \| P_0) : \pi P = \pi \big\}$$

with $K(P \| P_0) = \sum_{x,x'} \pi(x) P(x, x') \log\Big(\frac{P(x,x')}{P_0(x,x')}\Big)$
(*Donsker-Varadhan rate function*).

Theorem

There exist optimizers $\{\check{\pi}_{0\zeta}, \check{P}_\zeta : \zeta \in \mathbb{R}\}$, *and solutions* $\{h^*_\zeta, \eta^*(\zeta) : \zeta \in \mathbb{R}\}$ *s.t.*

- $\check{P}_\zeta$ *can be obtained from the relative value function* h^*_ζ *as*

  $$\check{P}_\zeta(x, x') := P_0(x, x') \exp\big(h_\zeta(x'_u \mid x) - \Lambda_{h_\zeta}(x)\big)$$

  where $h_\zeta(x'_u \mid x) = \sum_{x'_n} Q_0(x, x'_n) h^*_\zeta(x'_u, x'_n)$, *and* $\Lambda_{h_\zeta}(x)$ *is the normalizing constant.*
- $\{\check{\pi}_{0\zeta}, \check{P}_\zeta, h^*_\zeta, \eta^*(\zeta) : \zeta \in \mathbb{R}\}$ *are continuously differentiable in the parameter* ζ.

Theorem

There exist optimizers $\{\check{\pi}_{0\zeta}, \check{P}_\zeta : \zeta \in \mathbb{R}\}$*, and solutions* $\{h^*_\zeta, \eta^*(\zeta) : \zeta \in \mathbb{R}\}$ *s.t.*

- $\check{P}_\zeta$ *can be obtained from the relative value function* h^*_ζ *as*

  $$\check{P}_\zeta(x,x') := P_0(x,x')\exp\bigl(h_\zeta(x'_u \mid x) - \Lambda_{h_\zeta}(x)\bigr)$$

  where $h_\zeta(x'_u \mid x) = \sum_{x'_n} Q_0(x,x'_n)h^*_\zeta(x'_u,x'_n)$*, and* $\Lambda_{h_\zeta}(x)$ *is the normalizing constant.*
- $\{\check{\pi}_{0\zeta}, \check{P}_\zeta, h^*_\zeta, \eta^*(\zeta) : \zeta \in \mathbb{R}\}$ *are continuously differentiable in the parameter* ζ.

A vector field $\mathcal{V}\colon \mathbb{R}^d \to \mathbb{R}^d$ (with $d = |\mathsf{X}|$) is constructed s.t.

$$\frac{d}{d\zeta}h^*_\zeta = \mathcal{V}(h^*_\zeta)\,, \qquad \text{with boundary condition } h^*_0 \equiv 0.$$

ODE for the average reward

For any function $h\colon \mathsf{X} \to \mathbb{R}$,

(i) Define a new transition matrix:

$$P_h(x,x') := P_0(x,x')\exp\bigl(h(x'_u \mid x) - \Lambda_h(x)\bigr), \quad x,x' \in \mathsf{X}, \tag{1}$$

with $h(x'_u \mid x) = \sum_{x'_n} Q_0(x,x'_n)h(x'_u,x'_n)$, and $\Lambda_h(x)$ is a normalizing constant.

(ii) Let $H_h = \mathcal{H}(P_h)$, be a solution to Poisson's equation,

$$P_h H_h = H_h - \mathcal{U} + \bar{u}_h\,, \quad \text{where } \bar{u}_h := \pi_h(\mathcal{U}) := \sum_x \pi_h(x)\mathcal{U}(x).$$

and define $\mathcal{V}(h) = H_h$.

Note: The functional $\mathcal{H}$ is constructed so that $H_h(x^\circ) = 0$ for any h.

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

Tracking objective

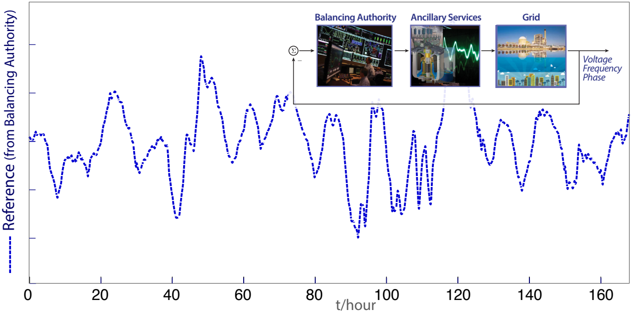

In the past, provided by the generators - high costs!

Control Goals and Architecture

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

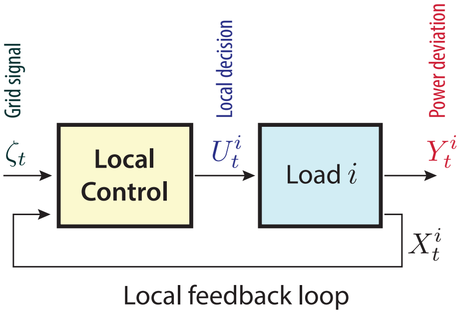

- Min. communication: each load monitors its state and a regulation signal from the grid.
- Aggregate must be controllable: **randomized policies** for finite-state loads.

Randomized Policies

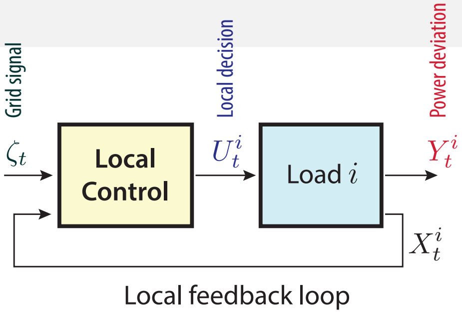

Local control architecture

For the ith load:

- Y_t^i power, U_t^i load setpoint, X_t^i local state.
  Signal ζ_t is from the grid operator – common to all loads of a certain class.

Randomized Policies

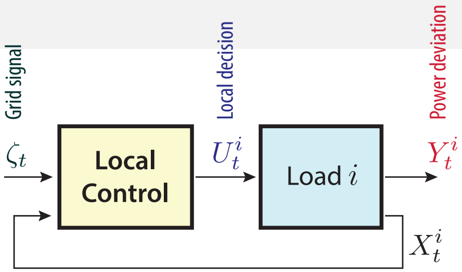

Local feedback loop

Local control architecture

For the ith load:

- Y_t^i power, U_t^i load setpoint, X_t^i local state.
  Signal ζ_t is from the grid operator – common to all loads of a certain class.
- Policy: Decision rule that maps (ζ_t, X_t^i) to the input U_t^i.
- Randomized Policy: Decision rule also depends on `rand`

(an intelligent coin-flip)

Load Model

Controlled Markovian Dynamics

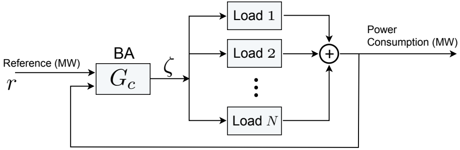

- Discrete time: ith load $X^i(t)$ evolves on finite state space X
- Each load is subject to *common* controlled Markovian dynamics.

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?

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 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:

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

with h_ζ continuously differentiable in ζ, and the normalizing constant

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

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

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

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 $\{h_\zeta : \zeta \in \mathbb{R}\}$

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 $\mathsf{X} \times \mathsf{X}$.

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 X × 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:

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

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

$$H(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) H^\circ(x'_u, x'_n) \tag{2}$$

Local Design

Extending local control design to include exogenous disturbances

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

$$H(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) H^\circ(x'_u, x'_n) \tag{2}$$

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,

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

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

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:

Individual Perspective Design

- ODE method for IPD design:

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

Individual Perspective Design

- ODE method for IPD design:

  Family $\{P_\zeta\}$: $P_\zeta(x,x') := P_0(x,x')\exp\big(h_\zeta(x,x') - \Lambda_{h_\zeta}(x)\big)$
  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 $h^\circ_0 \equiv 0$.

Individual Perspective Design

- ODE method for IPD design:

  Family $\{P_\zeta\}$: $P_\zeta(x,x') := P_0(x,x')\exp\big(h_\zeta(x,x') - \Lambda_{h_\zeta}(x)\big)$
  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 $Z_\zeta = [I - P_\zeta + 1\otimes\pi_\zeta]^{-1} = \sum_{n=0}^{\infty}[P_\zeta - 1\otimes\pi_\zeta]^n$ is the fundamental matrix.

Individual Perspective Design

Linearized dynamics

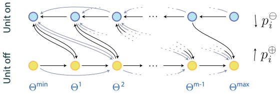

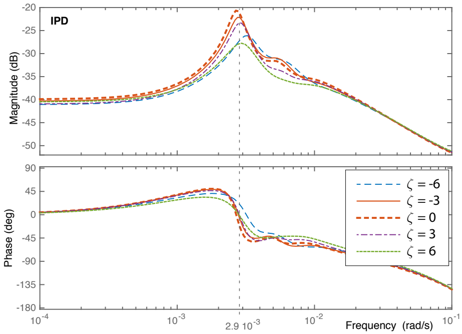

Bode plots for IPD: Linearizations at five values of ζ

Individual Perspective Design

Linearized dynamics

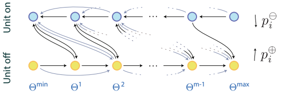

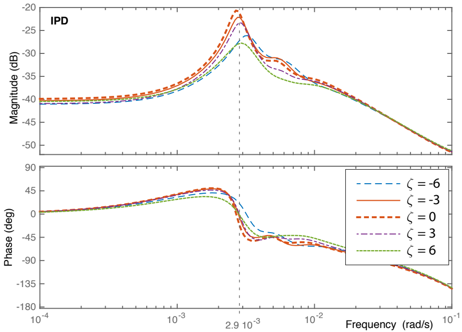

Bode plots for IPD: Linearizations at five values of ζ

Proof of positive real condition for reversible load dynamics.

System Perspective Design

Strictly positive real by design

Goal: The transfer function of the delay-free linearized aggregate model is passive:

$$\sum_{t=0}^{\infty} u_t y_{t+1} \geq 0, \ \forall \{u_t\}.$$

System Perspective Design

Strictly positive real by design

Goal: The transfer function of the delay-free linearized aggregate model is passive:

$$\sum_{t=0}^{\infty} u_t y_{t+1} \geq 0, \ \forall \{u_t\}.$$

Recall: The linearization at a particular value ζ is the state space model with transfer function,

$$G_\zeta(z) = C[Iz - A]^{-1} B$$

in which $A = P_\zeta^{\mathsf{T}}$, $C_i = \widetilde{\mathcal{U}}_\zeta(x^i)$ for each i, and

$$B_i = \sum_x \pi_\zeta(x) \mathcal{E}_\zeta(x, x^i), \qquad 1 \leq i \leq d$$

$\mathcal{E}_\zeta = \frac{d}{d\zeta} P_\zeta$ $\widetilde{\mathcal{U}}_\zeta = \mathcal{U} - \overline{\mathcal{U}}_\zeta$, with $\overline{\mathcal{U}}_\zeta = \pi_\zeta(\mathcal{U})$.

System Perspective Design

Strictly positive real by design

Goal: The transfer function of the delay-free linearized aggregate model is passive:

$$\sum_{t=0}^{\infty} u_t y_{t+1} \geq 0, \ \forall \{u_t\}.$$

Recall: The linearization at a particular value ζ is the state space model with transfer function,

$$G_\zeta(z) = C[Iz - A]^{-1} B$$

in which $A = P_\zeta^{\intercal}$, $C_i = \widetilde{\mathcal{U}}_\zeta(x^i)$ for each i, and

$$B_i = \sum_x \pi_\zeta(x)\mathcal{E}_\zeta(x, x^i), \qquad 1 \leq i \leq d$$

$\mathcal{E}_\zeta = \frac{d}{d\zeta} P_\zeta$ $\widetilde{\mathcal{U}}_\zeta = \mathcal{U} - \overline{\mathcal{U}}_\zeta$, with $\overline{\mathcal{U}}_\zeta = \pi_\zeta(\mathcal{U})$.

Sufficient condition: positive real.

A discrete-time transfer function F is *positive-real* if it is stable (all poles are strictly within the unit disk), and $F(e^{j\theta}) + F(e^{-j\theta}) \geq 0, \quad \theta \in \mathbb{R}$.

System Perspective Design

Strictly positive real by design

SPD design:

- $P^{\nabla} = P^{\Lsh} P$, with $P^{\Lsh}$ adjoint of P in $L_2(\pi)$:
  $P^{\Lsh}(x, x') = \frac{\pi(x')}{\pi(x)} P(x', x), \ x, x' \in \mathsf{X}$.

System Perspective Design

Strictly positive real by design

SPD design:

- $P^{\nabla} = P^{\dagger}P$, with $P^{\dagger}$ adjoint of P in $L_2(\pi)$:
  $P^{\dagger}(x,x') = \frac{\pi(x')}{\pi(x)} P(x',x),\ \ x,x' \in \mathsf{X}$.

- $H^{\circ}(x) = \sum_{x'} [Z^{\nabla}(x,x') - Z^{\nabla}(x^{\circ},x')]\mathcal{U}(x')\ x \in \mathsf{X}$
  where $Z^{\nabla} = [I - P^{\nabla} + 1 \otimes \pi]^{-1}$ the fundamental matrix for P^{∇}

System Perspective Design

Strictly positive real by design

SPD design:

- $P^{\triangledown} = P^{\dagger} P$, with $P^{\dagger}$ adjoint of P in $L_2(\pi)$:
  $P^{\dagger}(x,x') = \frac{\pi(x')}{\pi(x)} P(x',x), \; x,x' \in \mathsf{X}$.

- $H^{\circ}(x) = \sum_{x'} [Z^{\triangledown}(x,x') - Z^{\triangledown}(x^{\circ},x')]\mathcal{U}(x') \; x \in \mathsf{X}$
  where $Z^{\triangledown} = [I - P^{\triangledown} + 1 \otimes \pi]^{-1}$ the fundamental matrix for $P^{\triangledown}$

Thm. (SPD design) If $P_0^{\triangledown} = P_0^{\dagger} P_0$ is irreducible, and $P_0 = R_0$, then the linearized state-space model at any constant value ζ satisfies

$$G_\zeta^+(e^{j\theta}) + G_\zeta^+(e^{-j\theta}) \geq \sigma_\zeta^2, \qquad \theta \in \mathbb{R}$$

where σ_ζ^2 is the variance of $\mathcal{U}$ under π_ζ and $G^+(z) := zG(z)$.

Exponential family

Alternative to solving an ODE

For a function $H_{\mathrm{e}}^{\circ}\colon \mathsf{X} \to \mathbb{R}$, define for each x, x'_u and ζ,

$$h_\zeta(x, x'_u) = \zeta H_{\mathrm{e}}(x'_u \mid x)$$

$$\text{with} \quad H_{\mathrm{e}}(x'_u \mid x) := \sum_{x'_n} Q_0(x, x'_n) H_{\mathrm{e}}^{\circ}(x'_u, x'_n)$$

Exponential family

Alternative to solving an ODE

For a function $H_\mathrm{e}^\circ \colon \mathsf{X} \to \mathbb{R}$, define for each x, x'_u and ζ,

$$h_\zeta(x, x'_u) = \zeta H_\mathrm{e}(x'_u \mid x)$$

$$\text{with} \quad H_\mathrm{e}(x'_u \mid x) := \sum_{x'_n} Q_0(x, x'_n) H_\mathrm{e}^\circ(x'_u, x'_n)$$

- Myopic design: $H_\mathrm{e}^\circ = \mathcal{U}$.

Exponential family

Alternative to solving an ODE

For a function $H_{\mathrm{e}}^{\circ}\colon \mathsf{X} \to \mathbb{R}$, define for each x, x'_u and ζ,

$$h_\zeta(x, x'_u) = \zeta H_{\mathrm{e}}(x'_u \mid x)$$

$$\text{with} \quad H_{\mathrm{e}}(x'_u \mid x) := \sum_{x'_n} Q_0(x, x'_n) H_{\mathrm{e}}^{\circ}(x'_u, x'_n)$$

- Myopic design: $H_{\mathrm{e}}^{\circ} = \mathcal{U}$.
- Linear approximations to the IPD or SPD solutions, with $H_{\mathrm{e}}^{\circ} = \mathcal{H}^{\circ}(P_0)$.

Myopic Design

Linearized dynamics

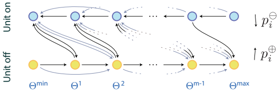

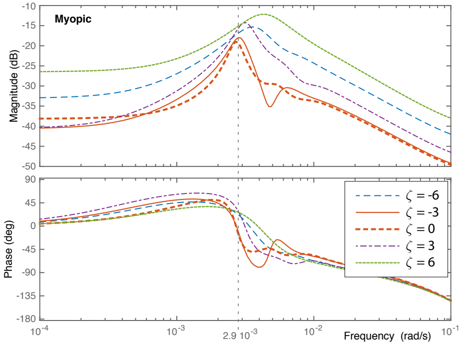

Bode plots for myopic design: Linearizations at five values of ζ

Myopic Design

Linearized dynamics

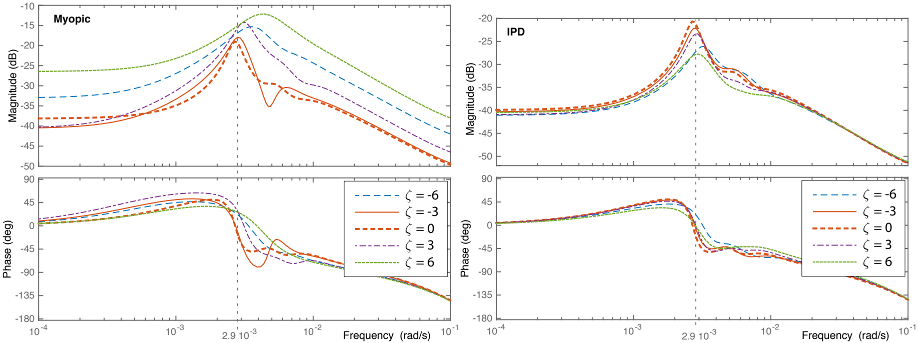

Example: pool pumps

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*

One Million Pools in Florida

Pools Service the Grid Today

On Call[1]: Utility controls residential pool pumps and other loads

[1] *Florida Power and Light*, Florida's largest utility.
www.fpl.com/residential/energysaving/programs/oncall.shtml

One Million Pools in Florida

Pools Service the Grid Today

On Call[1]: Utility controls residential pool pumps and other loads

Contract for services: no *price signals* involved

Used only in times of emergency — Activated only 3-4 times a year

[1] *Florida Power and Light*, Florida's largest utility.
www.fpl.com/residential/energysaving/programs/oncall.shtml

One Million Pools in Florida

Pools Service the Grid Today

On Call[1]: Utility controls residential pool pumps and other loads

Contract for services: no *price signals* involved

Used only in times of emergency — Activated only 3-4 times a year

Opportunity:
FP&L already has their hand on the switch of nearly one million pools!

[1] *Florida Power and Light*, Florida's largest utility.
www.fpl.com/residential/energysaving/programs/oncall.shtml

One Million Pools in Florida

Pools Service the Grid Today

On Call[1]: Utility controls residential pool pumps and other loads

Contract for services: no *price signals* involved

Used only in times of emergency — Activated only 3-4 times a year

Opportunity:

FP&L already has their hand on the switch of nearly one million pools!

Surely pools can provide much more service to the grid

[1] *Florida Power and Light*, Florida's largest utility.
www.fpl.com/residential/energysaving/programs/oncall.shtml

An Intelligent Pool

Local Control Architecture

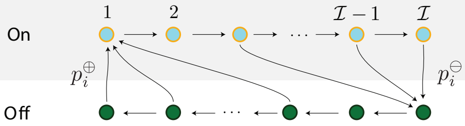

Local control architecture

For the ith load:

- $Y_t^i =$ power: 1kW when running,

An Intelligent Pool

Local Control Architecture

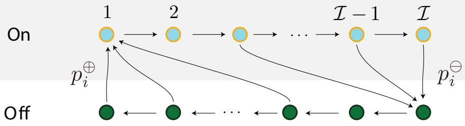

Local control architecture

For the ith load:

- Y_t^i = power: 1kW when running,
  $U_t^i = 1$ or 0 (pool pump is running or not)

An Intelligent Pool

Local Control Architecture

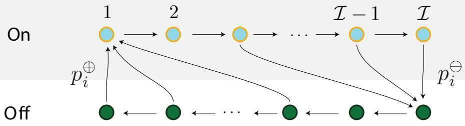

Local control architecture

For the ith load:

- Y_t^i = power: 1kW when running,
  $U_t^i = 1$ or 0 (pool pump is running or not)
  X_t^i local state: (U_t^i, I_t^i), with I_t^i the time in current power state.

An Intelligent Pool

Local Control Architecture

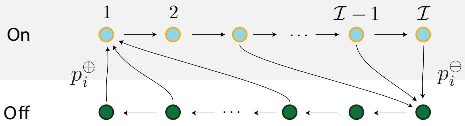

Local control architecture

For the ith load:

- Y_t^i = power: 1kW when running,
  $U_t^i = 1$ or 0 (pool pump is running or not)
  X_t^i local state: (U_t^i, I_t^i), with I_t^i the time in current power state.
- Randomized Policy: Decision rule that maps $(\zeta_t, X_t^i, \texttt{rand}_i^t)$ to the input U_t^i.

An Intelligent Pool

Local Control Architecture

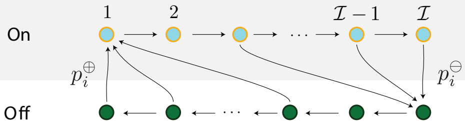

Local control architecture

For the ith load:

- Y_t^i = power: 1kW when running,
  $U_t^i = 1$ or 0 (pool pump is running or not)
  X_t^i local state: (U_t^i, I_t^i), with I_t^i the time in current power state.
- Randomized Policy: Decision rule that maps $(\zeta_t, X_t^i, \texttt{rand}_i^t)$ to the input U_t^i.

Randomized Policy: As ζ increases, probability of turning on increases:

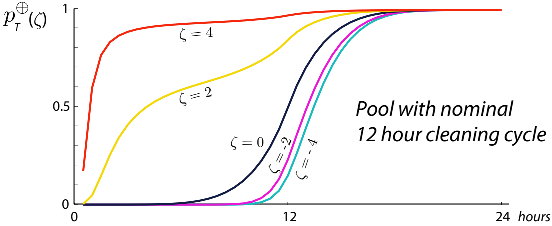

Tracking Grid Signal with Residential Loads

Example: 20 pools, 20 kW max load

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

Power Deviation:

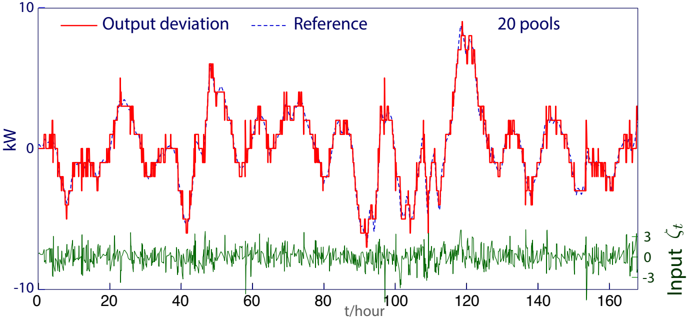

Nearly Perfect Service from Pools
Meyn et al. 2013 [CDC], Meyn et al. 2015 [IEEE TAC]

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:

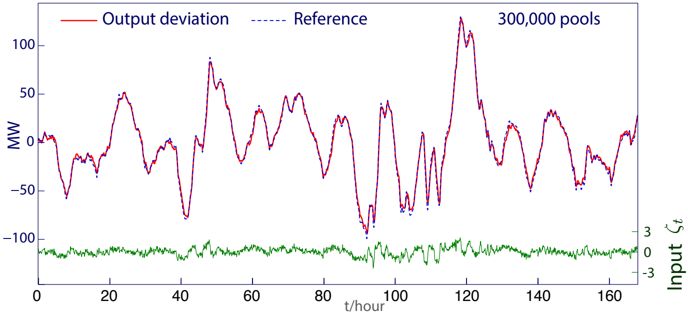

Nearly Perfect Service from Pools

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

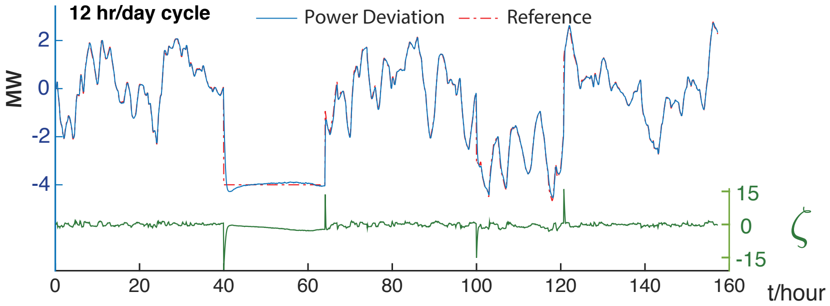

Example: Thermostatically Controlled Loads

- 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

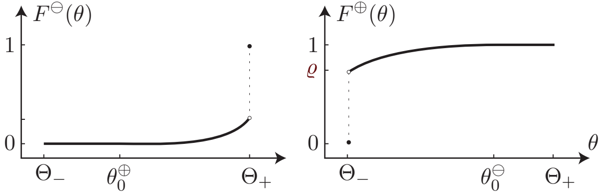

Example: Thermostatically Controlled Loads

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.

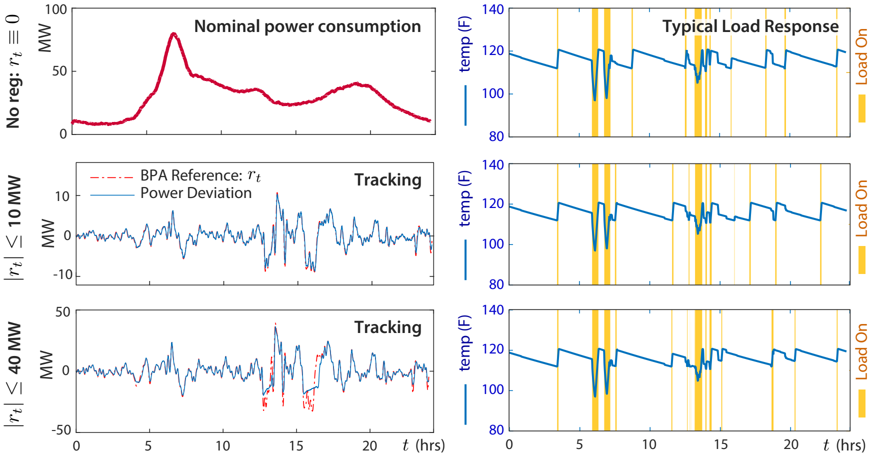

Tracking performance

Potential for contingency reserves and ramping

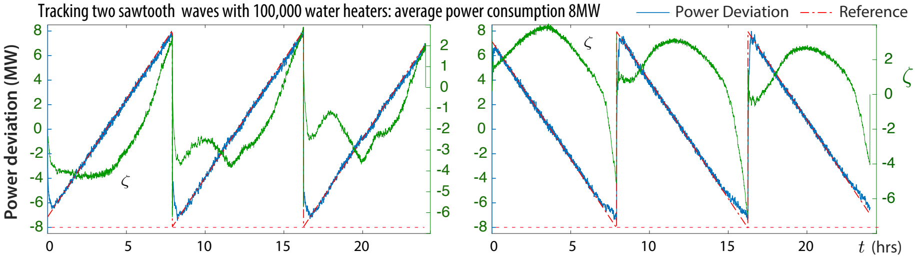

Tracking performance

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

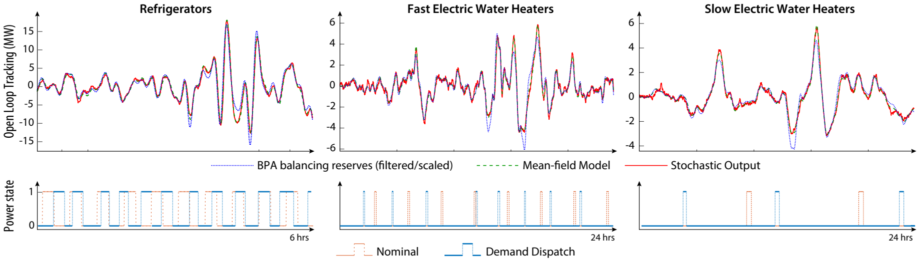

Example: fleet of batteries

[B. Hashmi, Meyn ACC'17]

State: $x = (m, s)$, where $m \in \{\text{ch, dis, id}\}$ denotes charging mode, and $s \in [0, 1]$ denotes the SoC.

The power delivery at state x depends only on charging mode:
$\mathcal{U}(\text{ch}, s) = \mathcal{U}_{\text{ch}} < 0$, $\mathcal{U}(\text{id}, s) = 0$, $\mathcal{U}(\text{dis}, s) = \mathcal{U}_{\text{dis}} > 0$.

Example: fleet of batteries

[B. Hashmi, Meyn ACC'17]

State: $x = (m, s)$, where $m \in \{\text{ch, dis, id}\}$ denotes charging mode, and $s \in [0, 1]$ denotes the SoC.

The power delivery at state x depends only on charging mode: $\mathcal{U}(\text{ch}, s) = \mathcal{U}_{\text{ch}} < 0$, $\mathcal{U}(\text{id}, s) = 0$, $\mathcal{U}(\text{dis}, s) = \mathcal{U}_{\text{dis}} > 0$.

Nominal model: $X_t^i = (M_t^i, S_t^i)$ denote the state of ith battery at time t. $S_{t+1}^i = S_t^i + h\delta_{\text{ch}}$, if $M_t^i = \text{ch}$, $S_{t+1}^i = S_t^i - h\delta_{\text{dis}}$, if $M_t^i = \text{dis}$, $S_{t+1}^i = S_t^i$, if $M_t^i = \text{id}$, where h is the time step length, and δ_{ch} and δ_{dis} charging and discharging rates.

The dynamics of the first component are governed by a "two coin-flip" randomized policy.
For example, in state (ch, s), the battery changes its mode to idle with probability $(1 - p_{\text{ch}}(s)) \times p_{\text{id}}(s)/(p_{\text{id}}(s) + p_{\text{dis}}(s))$

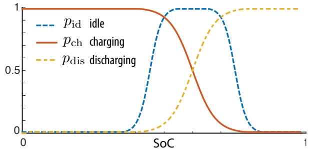

Example: fleet of batteries

1000 batteries, tracking PJM RegD test signal:

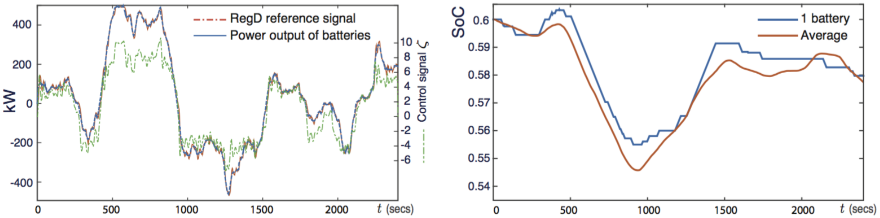

Unmodeled dynamics

[Chen, B., Meyn CDC'15, IEEE TAC'17]

Setting: 0.1% sampling, and

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

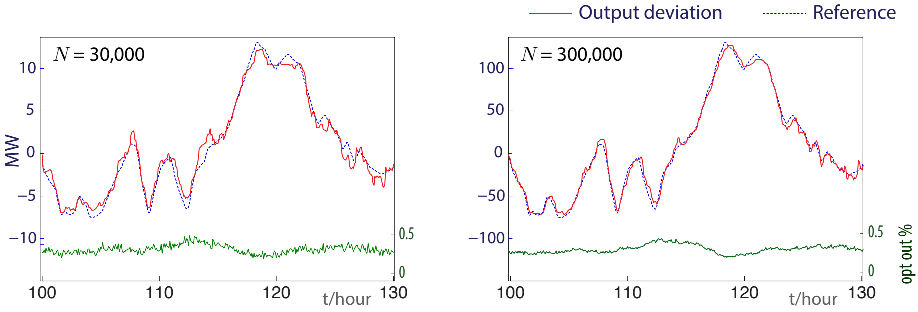

Closed-loop tracking

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

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:

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