HARVARD
John A. Paulson
School of Engineering
and Applied Sciences

The Role of Prediction in Online Optimal Control

Na (Lina) Li
Electrical Engineering and Applied Mathematics


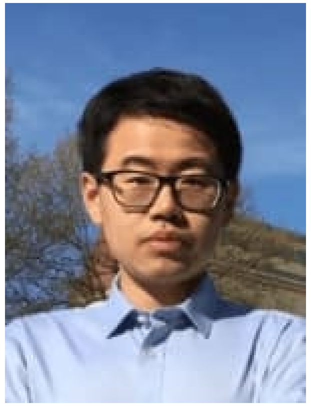


Acknowledgement:
Harvard University: Yingying Li, Xin Chen, Guannan Qu,
Funding Agencies: NSF, ARPA-E NODES, AFOSR, ONR

ACC workshop, July 09, 2019

Uncertainties in decision-making problems are ubiquitous

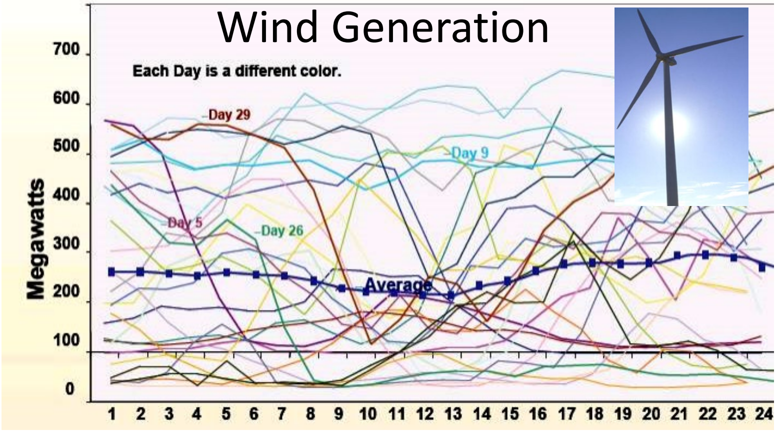

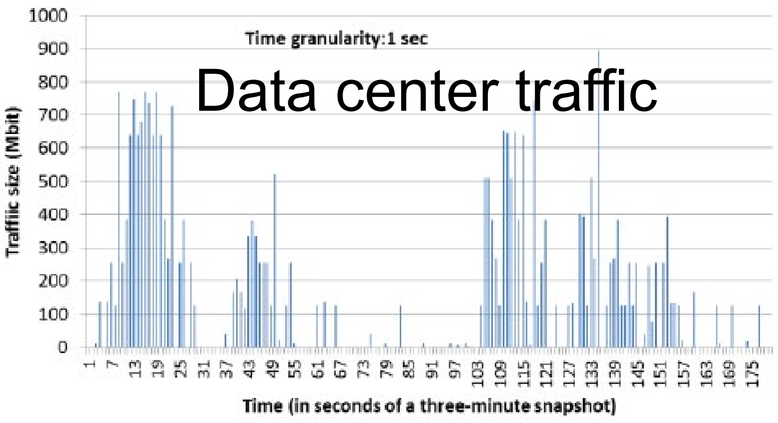

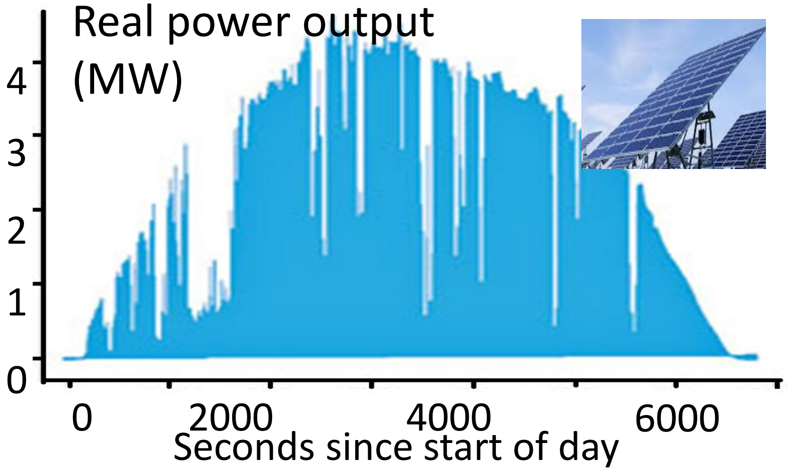

Online decision-making under a time-varying environment

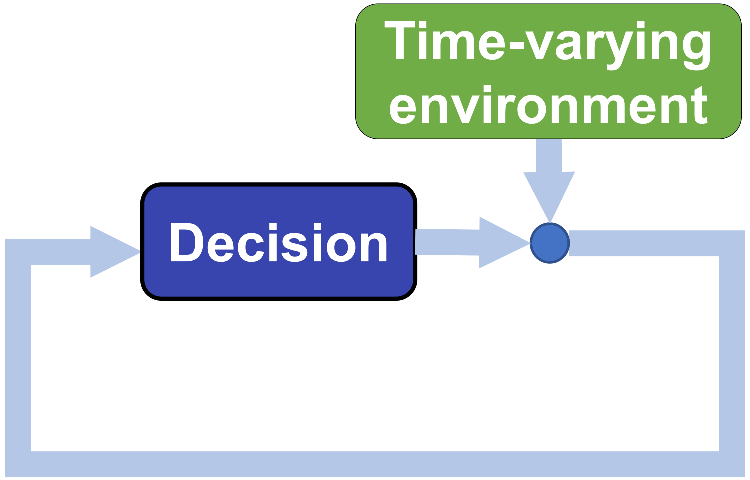

Online optimization

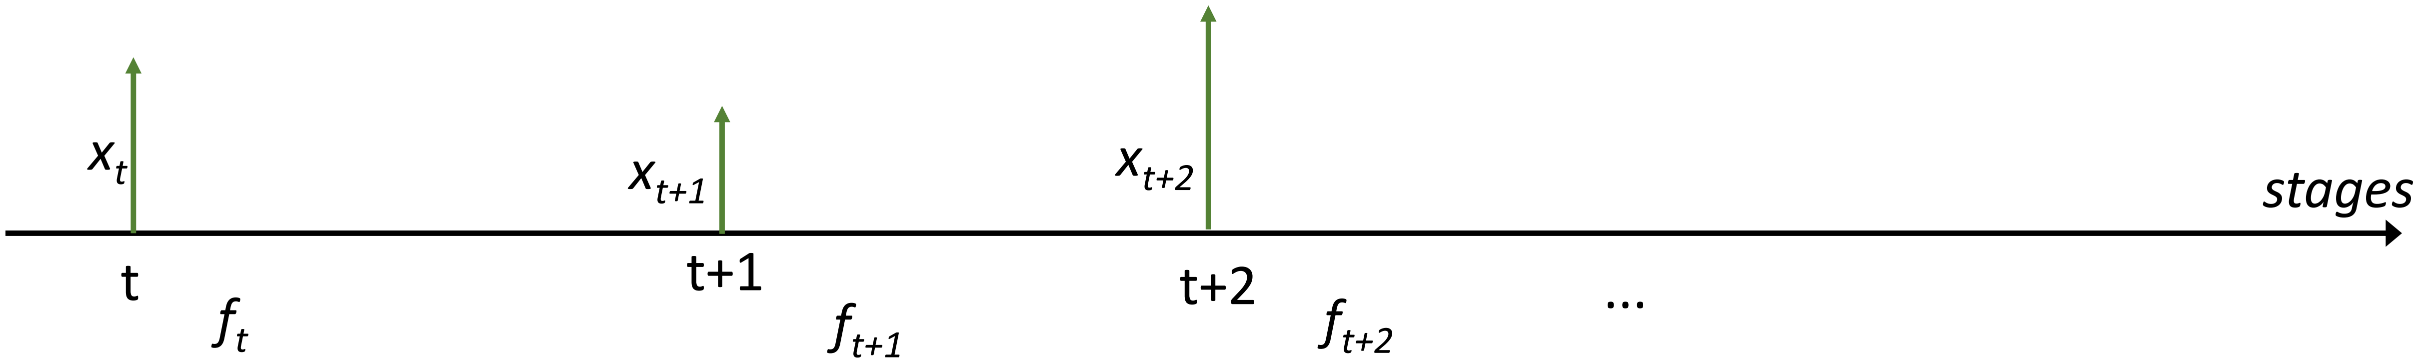

Input: A convex set $X \subset \mathbb{R}^n$
for: $t = 1,2, \ldots$
  make a decision $x_t \in X$
  receive a convex cost function $f_t : X \rightarrow \mathbb{R}$
  suffer cost $f_t(x_t)$
End

Goal: minimize total costs by making decisions without knowing the future

System dynamics: Switching/control costs

System Load from New England ISO

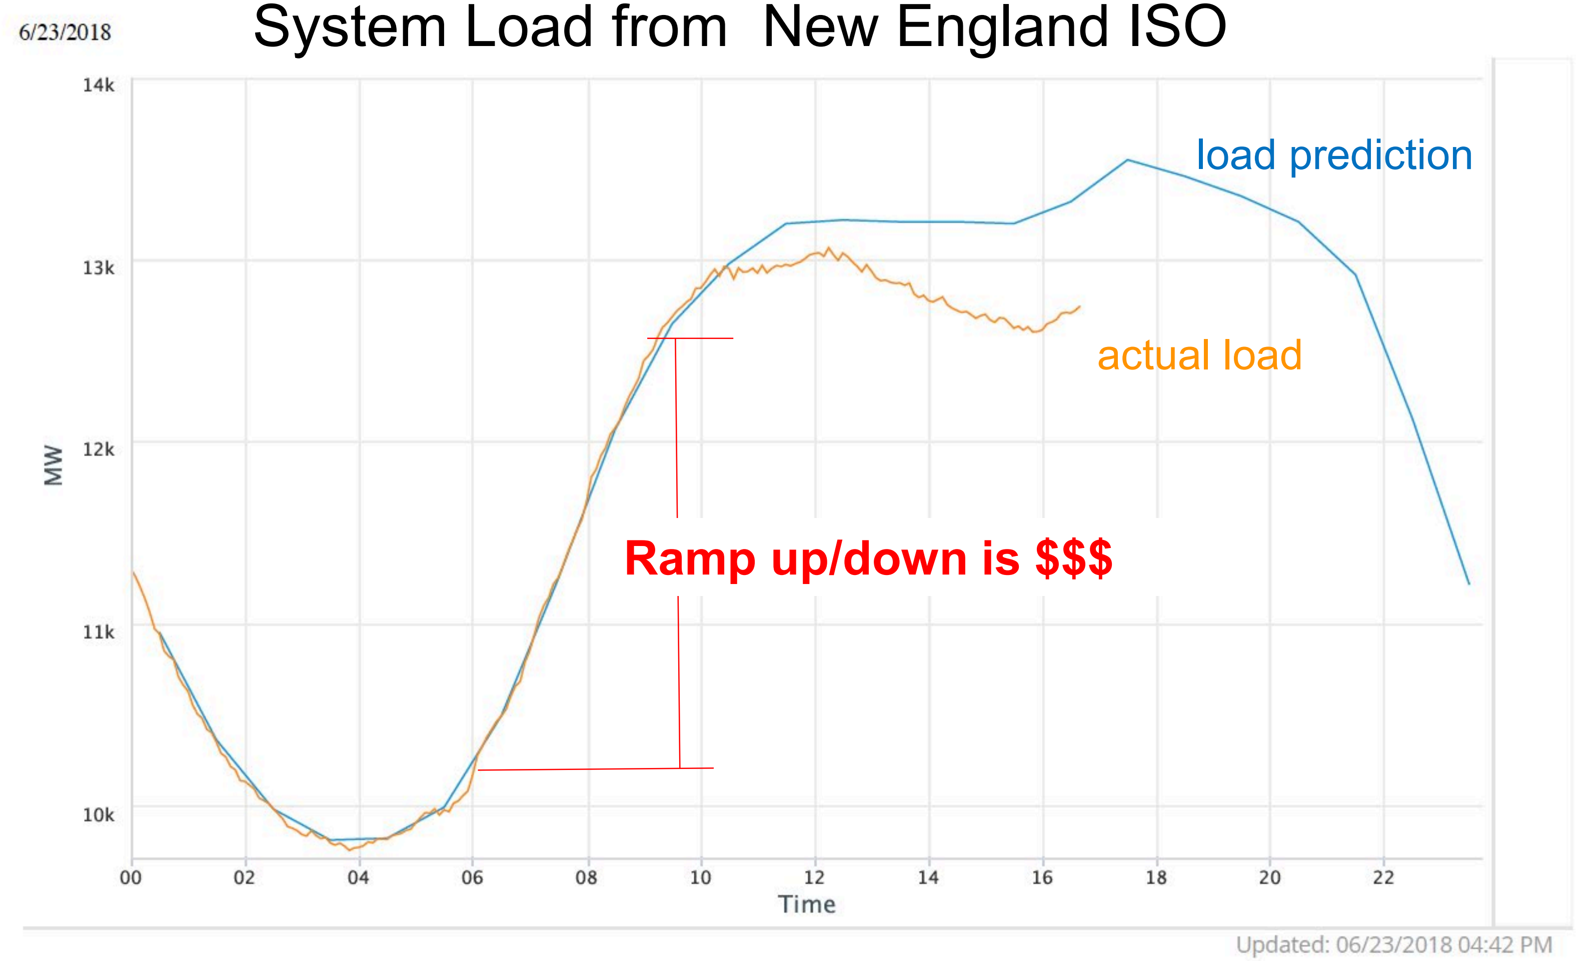

Economic Dispatch:
ramping cost

Vehicle racking:
energy cost

Source: New England ISO

Short-term prediction is possible

System Load from New England ISO

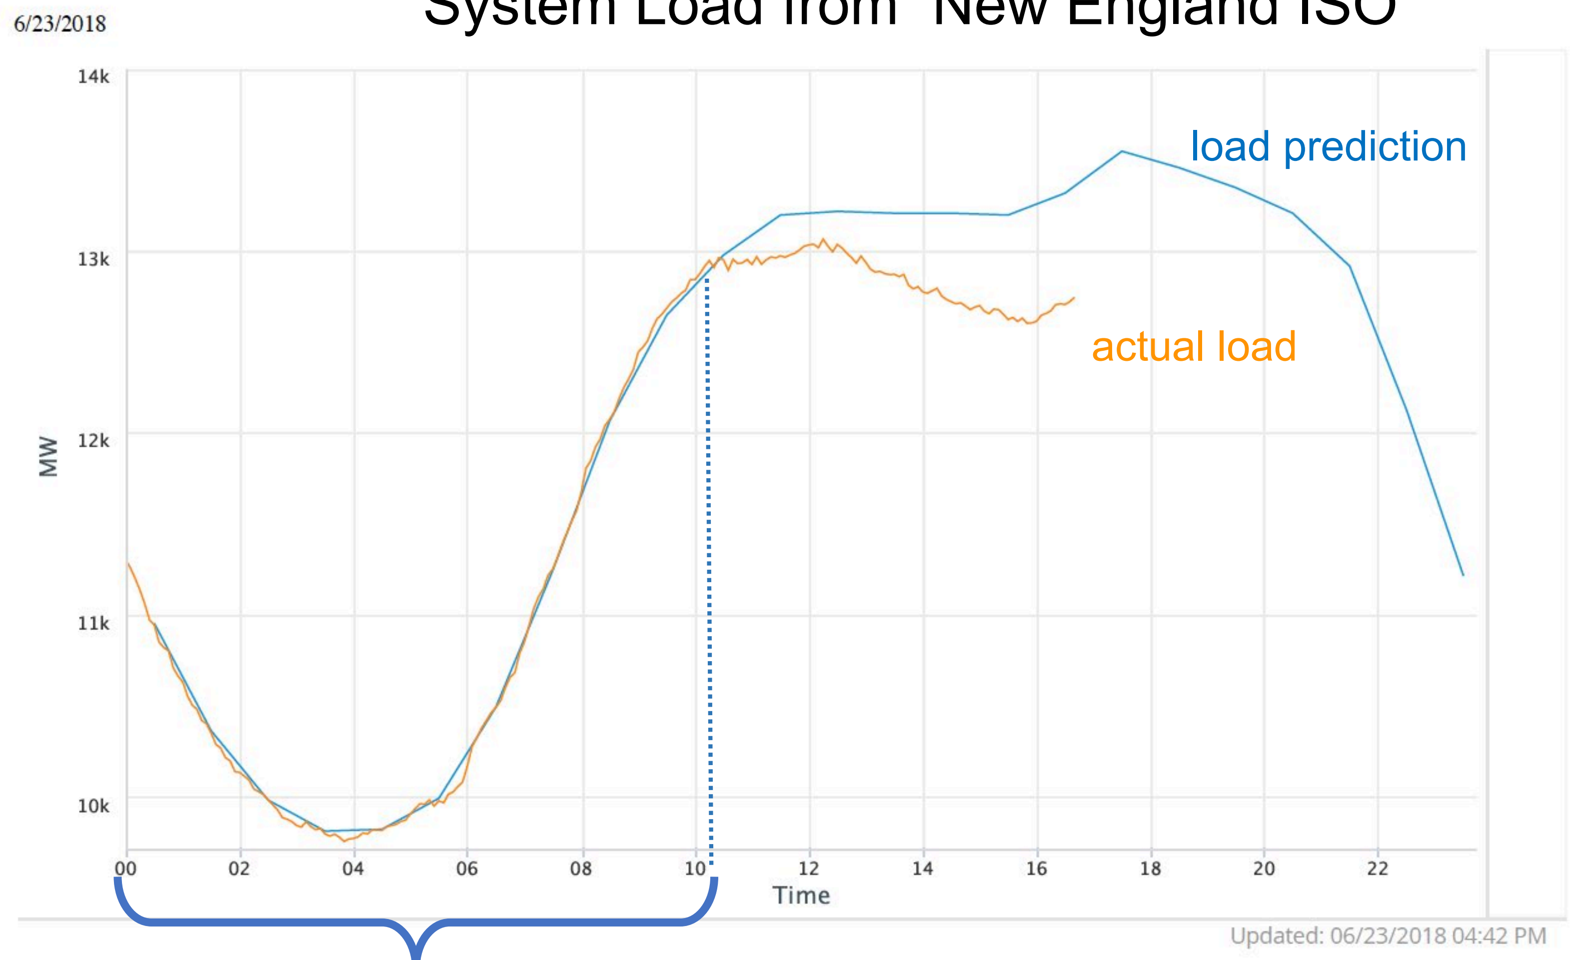

Source: New England ISO

Good short term prediction

Vehicle racking

Online control with linear system dynamics

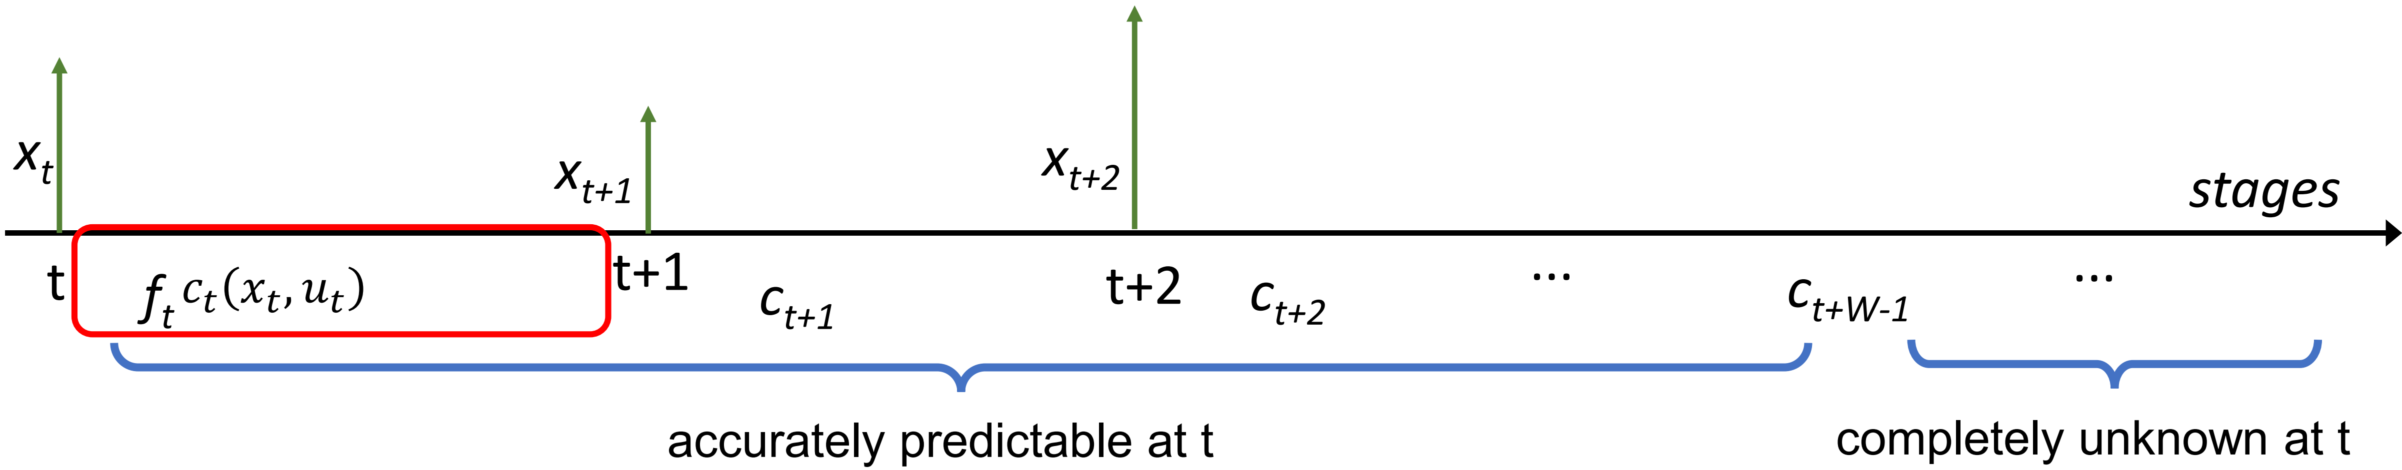

- System follows dynamics $x_t = g(x_{t-1}, u_t) \rightarrow x_t = Ax_{t-1} + Bu_t$
- Each time step $t = 1, \ldots, T$ is involved with a stage cost $c_t(x_t, u_t) = f_t(x_t) + g_t(u_t)$

  Example: LQR: $\frac{1}{2}(x_t - \theta_t)^T Q_t (x_t - \theta_t) + u_t^T R_t u_t$
- At time t, know accurate W-step prediction $c_t, \ldots, c_{t+W-1}$, $W \geq 0$

Online algorithm $\mathcal{A}$:

Pick action $u_t^{\mathcal{A}} \in U$ using prediction and history to minimize the total cost $C = \sum_t^T c_t$.

- Prior information: A, B, x_0 are known

Online control with linear system dynamics

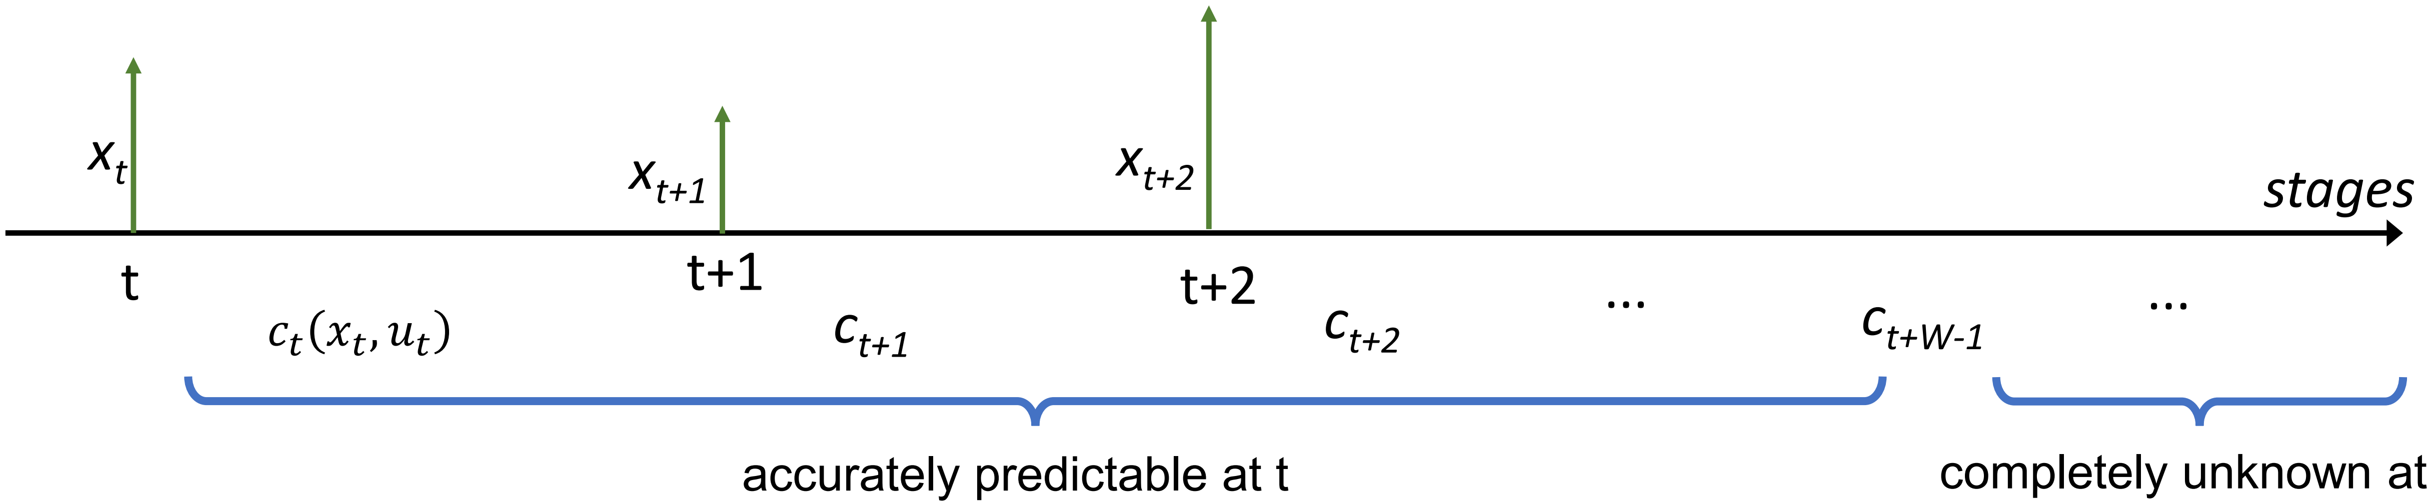

Performance metric: **dynamic** regret (competitive difference)

$$\text{Regret} := J(\mathbf{x}, \mathbf{u}) - \min_{\mathbf{x},\mathbf{u}} J(\mathbf{x}, \mathbf{u})$$

Online cost　　Optimal offline cost

where $J(\mathbf{x}, \mathbf{u})$ denotes the total cost function $J(\mathbf{x}, \mathbf{u}) = \sum_{t=1}^{T} c_t(x_t, u_t)$.

Literature review

MPC and economic MPC

[Mayne et al. 1999] [Morari et al. 1999] [Rawlings et al. 2012]...

- **Transient optimality** for **time-invariant** economic MPC
  [Grüne et al 2014] [Angeli et al 2012]...

- **Asymptotic** analysis for economic MPC with **time-varying costs**
  [Angeli et al 2015] [Ellis et al 2013] [Ferramosca, et al. 2010] [Ferramosca et al 2014] [Alessandretti et al 2016] [Angeli et al 2016]

Nonasymptotic analysis for time-varying costs?

What is achievable and what is not?—Role of prediction

Computational efficient way to achieve the performance limit?

Literature review

Online convex optimization with switching costs

[Lin et al. 2012] [Lin et al. 2013] [Chen et al. 2015] [Chen et al. 2016] [Li et al 2018] ...

- A **special** case of control problems: $c_t(x_t, u_t) = f(x_t) + \frac{\beta}{2}||u_t||^2$
  $$x_{t+1} = x_t + u_t$$
- Tools are developed for **nonasymptotic** regret analysis

Can we borrow ideas from online optimization to design control algorithms?
Can we borrow techniques to prove nonasymptotic bounds?

Brief introduction of online optimization with switching cost

Special case of online control

$$\min_{x_t, u_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|u_t\|^2)$$

s.t. $x_{t+1} = x_t + u_t$

$\Rightarrow$

Online opt with switching costs

$$\min_{x_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|x_{t+1} - x_t\|^2)$$

Main Messages:

Given any online optimization algorithm **without** using prediction, **additional W steps of gradient calculation** can **further** reduce the regret by ***O*****(*r*^{W}).**

Offline gradient descent

Offline optimization: $\min_{x \in X \times \cdots \times X} C_1^T(x)$ where $C_1^T(x) := \sum_{t=1}^{T} \left(f_t(x_t) + \frac{\beta}{2} \|x_t - x_{t-1}\|^2 \right)$

Offline gradient descent:

#iter

$$x_t^{(k)} = \Pi_X \left[x_t^{(k-1)} - \eta \left(\underbrace{\nabla f_t(x_t^{(k-1)}) + \beta(x_t^{(k-1)} - x_{t-1}^{(k-1)}) - \beta(x_{t+1}^{(k-1)} - x_t^{(k-1)})}_{\frac{\partial C_1^T}{\partial x_t}(x^{(k-1)})} \right) \right]$$

stage t=1...T

k=0 (initialization) $\cdots$ $\underbrace{x_{t+W-2}^{(0)}, x_{t+W-1}^{(0)}, x_{t+W}^{(0)}}$, $\cdots$ f_{t+W-1} OCO without prediction

k=1 $x_{t+W-1}^{(1)}$

$\vdots$

k=W−2 $\underbrace{x_t^{(W-2)}, x_{t+1}^{(W-2)}, x_{t+2}^{(W-2)}},$ f_{t+1}

k=W−1 $\underbrace{x_{t-1}^{(W-1)}, x_t^{(W-1)}, x_{t+1}^{(W-1)}},$ f_t

k=W $x_t^{(W)}$

Key observation: Predictable at t

To compute $x_t^{(W)}$, need f_t ... f_{t+W-1}.

Initialization: any OCO without prediction

RHGD

Offline gradient descent:

$$x_t^{(k)} = \Pi_X \left[x_t^{(k-1)} - \eta \left(\nabla f_t(x_t^{(k-1)}) + \beta(x_t^{(k-1)} - x_{t-1}^{(k-1)}) - \beta(x_{t+1}^{(k-1)} - x_t^{(k-1)}) \right) \right]$$

Additional Computation per stage:

- $(W+1)$ projected gradient evaluations

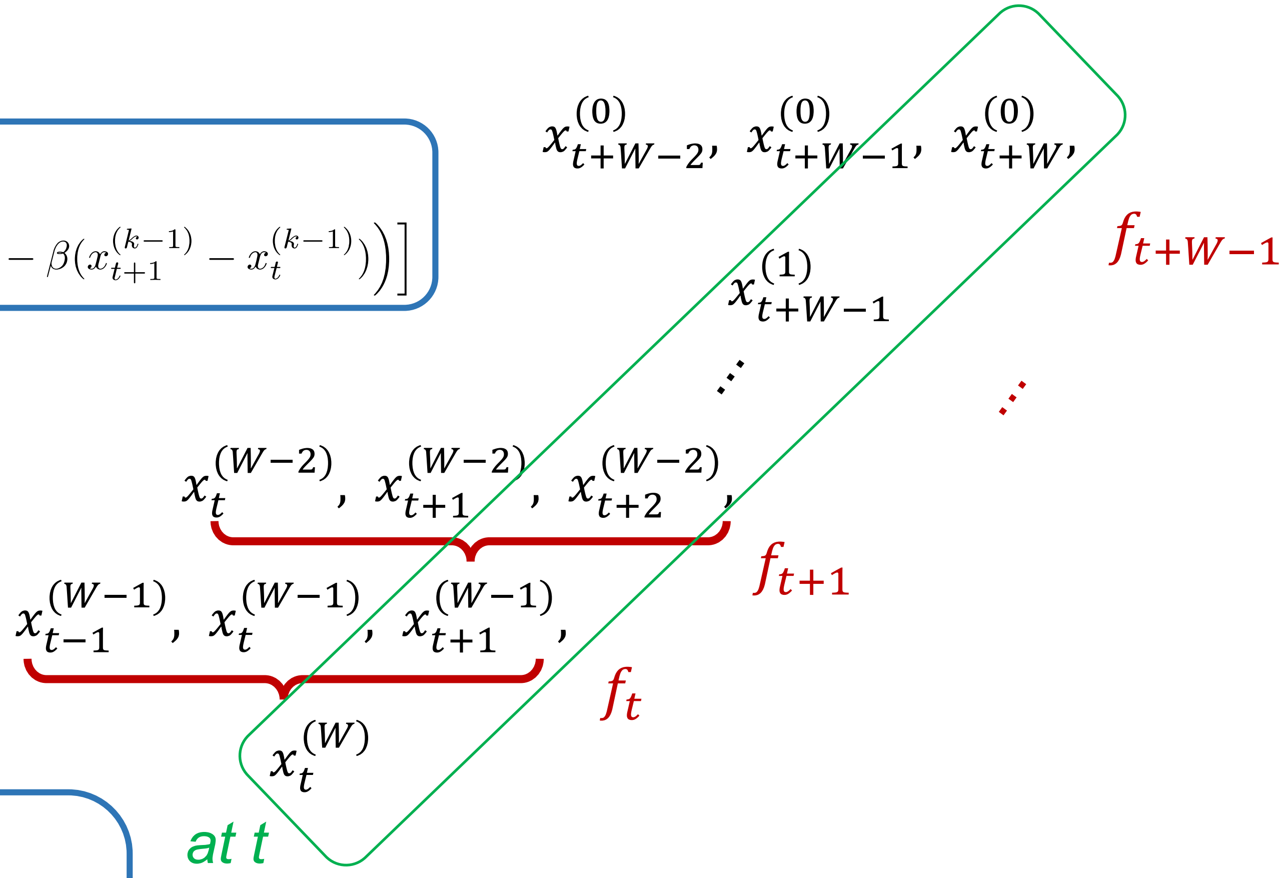

RHGD:

For t = 1, 2 ...

Receive prediction f_t ... f_{t+W-1}

Initialize $x_{t+W}^{(0)}$,

Backward compute $x_{t+W-1}^{(1)}$ ... , $x_t^{(W)}$

Output $x_t^{(W)}$

Key observation: Predictable at t

To compute $x_t^{(W)}$, need f_t ... f_{t+W-1}.

Initialization: any OCO without prediction

RHAG: Nesterov's accelerated gradient

RHGD:

Prediction $f_t, \dots, f_{t+W-1}$ allows W iterations of offline Nesterov's accelerated gradient

Offline optimization: $\min_{x \in X \times \cdots \times X} C_1^T(x)$ where $C_1^T(x) := \sum_{t=1}^{T} \left(f_t(x_t) + \frac{\beta}{2} \|x_t - x_{t-1}\|^2 \right)$

Offline Nesterov's accelerated gradient :

$$x_t^{(k)} = \Pi_X \left[y_t^{(k-1)} - \eta \left(\nabla f_t(y_t^{(k-1)}) + \beta(y_t^{(k-1)} - y_{t-1}^{(k-1)}) - \beta(y_{t+1}^{(k-1)} - y_t^{(k-1)}) \right) \right]$$

$$y_t^{(k)} = (1+\lambda) x_t^{(k)} - \lambda x_t^{(k-1)}$$

k=W−1 $\quad \underbrace{\left(x_{t-1}^{(W-1)}, y_{t-1}^{(W-1)}\right), \left(x_t^{(W-1)}, y_t^{(W-1)}\right), \left(x_{t+1}^{(W-1)}, y_{t+1}^{(W-1)}\right)}$,

k=W $\quad \left(x_t^{(W)}, y_t^{(W)}\right)$

RHAG: Nesterov's accelerated gradient

RHGD:

Prediction $f_t, \dots, f_{t+W-1}$ allows W iterations of offline Nesterov's accelerated gradient

Offline optimization: $\min_{x \in X \times \cdots \times X} C_1^T(x)$ where $C_1^T(x) := \sum_{t=1}^{T} \left(f_t(x_t) + \frac{\beta}{2} \|x_t - x_{t-1}\|^2 \right)$

Offline Nesterov's accelerated gradient :

$$x_t^{(k)} = \Pi_X \left[y_t^{(k-1)} - \eta \left(\nabla f_t(y_t^{(k-1)}) + \beta(y_t^{(k-1)} - y_{t-1}^{(k-1)}) - \beta(y_{t+1}^{(k-1)} - y_t^{(k-1)}) \right) \right]$$

$$y_t^{(k)} = (1+\lambda) x_t^{(k)} - \lambda x_t^{(k-1)}$$

Similar for other fast gradient methods, such as Triple Momentum (TM) method

(Scoy,Freeman,Lych, 2018)

Regret bounds: exponential decay with prediction length W

Theorem 1. (informal)

For any initialization method ψ, under proper stepsizes

$$\text{Regret}(RHGD) \leq O\left(\left(\frac{Q_f-1}{Q_f}\right)^W \text{Regret}(\psi)\right)$$

$$\text{Regret}(RHAG) \leq O\left(\left(\frac{\sqrt{Q_f}-1}{\sqrt{Q_f}}\right)^W \text{Regret}(\psi)\right)$$

$$\text{Regret}(RHTM) \leq O\left(\left(\frac{\sqrt{Q_f}-1}{\sqrt{Q_f}}\right)^{2W} \text{Regret}(\psi)\right)$$

where Q_f is the condition number of the total cost.

Exp. decay with W

General cases?

Simple Case of Online Control

$$\min_{x_t, u_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|u_t\|^2)$$

$$\text{s.t.} \quad x_{t+1} = x_t + u_t$$

Online Opt with Switching Costs

$$\min_{x_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|x_{t+1} - x_t\|^2)$$

General linear dynamics

$$\min_{x_t, u_t} \sum_t (f_t(x_t) + g_t(u_t))$$

$$\text{s.t.} \quad x_{t+1} = Ax_t + Bu_t$$

?

Canonical form

- Any controllable system can be converted to the canonical form

$$
A = \begin{bmatrix}
0 & 1 & & 0 & & & & & & & & & \\
\vdots & \ddots & \ddots & & & & & & & & & & \\
\vdots & & 0 & 1 & & & & & & & & & \\
* & * & \cdots & * & * & * & \cdots & * & & & \cdots & & * \\
 & & & & 0 & 1 & & 0 & & & & & \\
 & & & & \vdots & \ddots & \ddots & & & & & & \\
 & & & & \vdots & & 0 & 1 & & & & & \\
* & * & \cdots & * & * & * & \cdots & * & \cdots & * & \cdots & & * \\
\vdots & & & & & & & & \ddots & & & & \\
 & & & & & & & & & 0 & 1 & \cdots & 0 \\
 & & & & & & & & & \vdots & \ddots & \ddots & \\
 & & & & & & & & & & & 0 & 1 \\
* & * & \cdots & * & * & * & \cdots & * & \cdots & * & * & \cdots & *
\end{bmatrix}
\quad
B = \begin{bmatrix}
0 & 0 & \cdots \\
\vdots & \vdots & \vdots \\
0 & \vdots & \vdots \\
1 & 0 & \cdots \\
0 & 0 & \\
\vdots & \vdots & \vdots \\
0 & 1 & \cdots \\
\vdots & & \ddots \\
0 & \vdots & \cdots \\
\vdots & \ddots & \\
0 & 0 & \cdots 1
\end{bmatrix}
$$

$$\min_{x_t, u_t} \sum_t (f_t(x_t) + g_t(u_t))$$

- **Example:**

$$\begin{bmatrix} x^1_{t+1} \\ x^2_{t+1} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ a_1 & a_2 \end{bmatrix} \begin{bmatrix} x^1_t \\ x^2_t \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_t$$

$$x_t = (x^2_{t-1}, x^2_t)'$$

$$u_t = x^2_{t+1} - a_1 x^2_{t-1} - a_2 x^2_t$$

Free variables: $\ldots, x^2_{t-1}, x^2_t, x^2_{t+1}, \ldots$

Finite-coupling costs

$$\min_{x^2_t} \sum_t \underbrace{f_t(x^2_{t-1}, x^2_t) + g_t(x^2_{t+1} - a_1 x^2_{t-1} - a_2 x^2_t)}_{\tilde{c}_t(x^2_{t-1}, x^2_t, x^2_{t+1})}$$

Finite-coupling costs have finite-coupling gradient

Finite-coupling costs: $$C(\mathbf{x}^2) := \sum_t \tilde{c}_t(x_{t-1}^2, x_t^2, x_{t+1}^2)$$

Finite-coupling partial derivative:

$$\frac{\partial C}{\partial x_t^2}(\mathbf{x}^2) = \frac{\partial(\tilde{c}_{t-1} + \cdots + \tilde{c}_{t+1})}{\partial x_t^2}(x_{t-2}^2(k), \ldots, x_{t+2}^2(k))$$

RHGD, RHAG, RHTM can be extended to the general cases

Regret bounds

Theorem 1. (informal)

For any initialization method φ, under proper stepsizes

$$\text{Regret}(RHGD) \le O\left(\left(\frac{Q_f-1}{Q_f}\right)^K \text{Regret}(\psi)\right)\ ;\ \text{Regret}(RHAG) \le O\left(\left(\frac{\sqrt{Q_f}-1}{\sqrt{Q_f}}\right)^K \text{Regret}(\psi)\right)$$

where $K = \lfloor \frac{W}{p} \rfloor$, p is maximum block size, Q_f is the conditional number.

- p is the max. block size of A: $p = \max(p_1, \ldots, p_m)$
- p=1: response time = 1

$$\begin{bmatrix} x_{t+1}^1 \\ x_{t+1}^2 \end{bmatrix} = \begin{bmatrix} a_1 & a_2 \\ a_3 & a_4 \end{bmatrix} \begin{bmatrix} x_t^1 \\ x_t^2 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} u_t$$

- p=2: response time =2

$$\begin{bmatrix} x_{t+1}^1 \\ x_{t+1}^2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ a_1 & a_2 \end{bmatrix} \begin{bmatrix} x_t^1 \\ x_t^2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_t$$

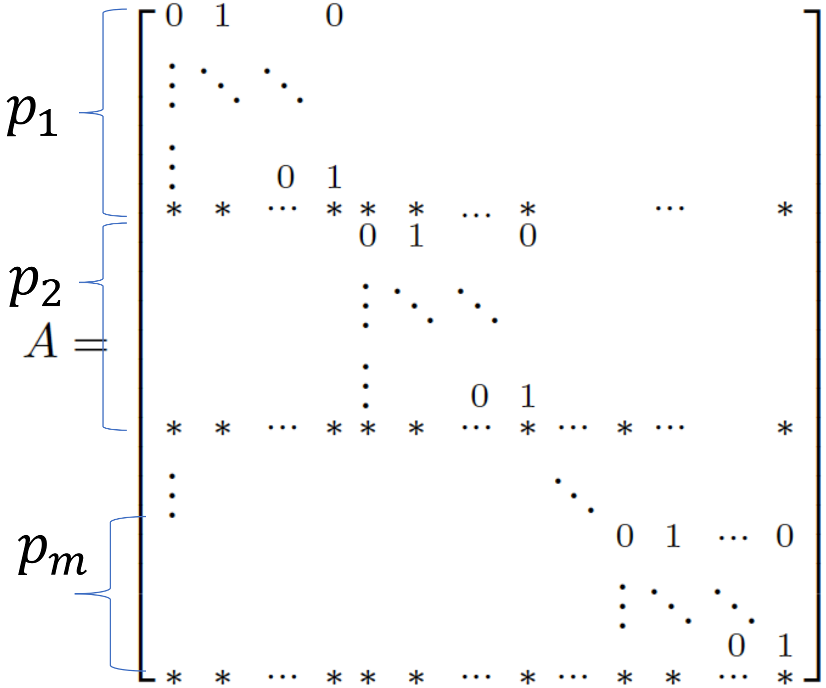

An initialization method: follow optimal steady states

Predictable

t

t+W-2 t+W-1 t+W

$x_{t+W}^2(0)$ Initialize?

Initialization algorithm:

- Solve steady state optimization at *t+W-1*:

$$(\bar{x}_{t+W-1}, \bar{u}_{t+W-1}) = \operatorname*{arg\,min}_{x=Ax+Bu} c_{t+W-1}(x,u)$$

- Initialize by *following the optimal steady state*:

$$x_{t+W}^2(0) = \bar{x}_{t+W-1}$$

Regret bound of the initialization method

Theorem 2. (informal)

variation of costs:
path length

For the online LQR problem, applying the optimal steady state at each time,

$$\text{Regret}(\psi) \le O\left(\sum_t \left(\|P_t^* - P_{t-1}^*\|^2 + \|\bar{x}_t^* - \bar{x}_{t-1}^*\|^2 + \|\beta_t^* - \beta_{t-1}^*\|^2\right)\right)$$

where P_t^* is the solution to the algebraic Riccati equation defined by Q_t, R_t; β_t^* is the drifting term in the optimal controller $u = K_t^* x + K_t' \beta_t^*$ to the LQR problem defined by Q_t, R_t, θ_t.

Small variation of costs ➡ Easy.
Future is similar; thus well-estimated by history

Large variation of costs ➡ Difficult!

Other initial. alg: online control without prediction [Abbasi-Yadkori et al 2014], [Goel et al. 2019]; online gradient [Li et al 2019]

Regret bound of RHTM with opt-steady-state initialization

Corollary 1. (informal)

For the online LQR problem, applying RHTM with the optimal steady state as initialization,

$$\text{Regret}(RHTM) \leq O\left(\left(\frac{\sqrt{Q_f}-1}{\sqrt{Q_f}}\right)^{2K} \sum_t \left(\|P_t^* - P_{t-1}^*\|^2 + \|\bar{x}_t^* - \bar{x}_{t-1}^*\|^2 + \|\beta_t^* - \beta_{t-1}^*\|^2\right)\right)$$

where P_t^* is the solution to the algebraic Riccati equation defined by Q_t, R_t; β_t^* is the drifting term in the optimal controller $u = K_t^* x + K_t' \beta_t^*$ to the LQR problem defined by Q_t, R_t, θ_t.

Special case: Q and R not change

Theorem. (informal)

For LQR tracking where stage cost is $c_t(x_t, u_t) := \frac{1}{2}(x_t - \theta_t)^T Q (x_t - \theta_t) + u_t^T R u_t$,

$$C(\text{RHTM}) - C(\text{Offline Optimal}) \leq O\left(\left(\frac{\sqrt{Q_f} - 1}{\sqrt{Q_f}}\right)^{2K} \sum_t \|\theta_t - \theta_{t-1}\|_2\right)$$

variation of costs
path length

where $K = \lfloor \frac{W}{p} \rfloor$, p is maximum block size, Q_f is the conditional number.

Tracking cost: $\| x_t - \theta_t \|^2$

tracker's location at t target's location at t

Small path length:

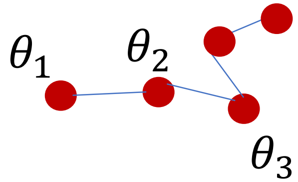

Easy!

Large path length:

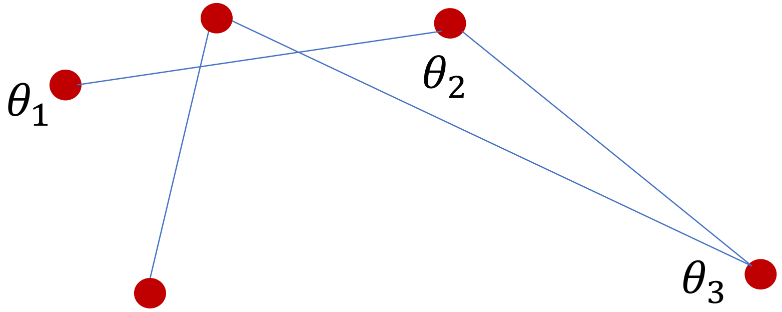

Difficult!

Larger regret

Dependence on W is nearly optimal

Theorem. (informal)

For LQR tracking where stage cost is $c_t(x_t, u_t) := \frac{1}{2}(x_t - \theta_t)^T Q(x_t - \theta_t) + u_t^T R u_t$,

$$C(\text{RHTM}) - C(\text{Offline Optimal}) \leq O\left(\left(\frac{\sqrt{Q_f} - 1}{\sqrt{Q_f}}\right)^{2K} \sum_t \|\theta_t - \theta_{t-1}\|_2\right)$$

where $K = \lfloor \frac{W}{p} \rfloor$, p is maximum block size, Q_f is the conditional number.

nearly optimal

Theorem 3 [Fundamental lower bound] (informal)

For any online algorithm $\mathcal{A}$, given any $1 < \text{W} < T$ and p>=1, there exists $A, B, \{\theta_t\}_{t=1}^T$ s.t.

$$C(\mathcal{A}) - C(\text{Offline Optimal}) \geq \Omega(\left(\frac{\sqrt{Q_f} - 1}{\sqrt{Q_f} + 1}\right)^{2K} \sum_t \|\theta_t - \theta_{t-1}\|_2)$$

where $K = \lfloor \frac{W}{p} \rfloor$, p is maximum block size, Q_f is the conditional number.

Proof sketch for the special case.

Special case of online control

$$\min_{x_t, u_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|u_t\|^2)$$

s.t. $x_{t+1} = x_t + u_t$

Online opt with switching costs

$$\min_{x_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|x_{t+1} - x_t\|^2)$$

Proof sketch: $\min_{x_t} \sum_t (f_t(x_t) + \frac{\beta}{2}\|x_{t+1} - x_t\|^2)$ & when path length is T.

- Construct $f_t(x_t, u_t) = \| x_t - \theta_t \|^2$, with **random** $\theta_1, \ldots, \theta_n$ i.i.d. $E\theta_i = 0$, $\text{var}(\theta_i) = 1$, bounded.

Path length $= O(T)$

- **Optimal offline** action

$$x_t^* = \underbrace{a_{t,1}\theta_1 + \cdots + a_{t,t}\theta_t + a_{t,t+1}\theta_{t+1} + \cdots + a_{t,t+W-1}\theta_{t+W-1}}_{x_t^t} + \underbrace{a_{t,t+W}\theta_{t+W} + \cdots + a_{t,T}\theta_T}_{x_t^{T-t}}$$

- **Online** action: only know $\theta_1, \ldots, \theta_{t+W-1}$,

Key 1: x_t^t is a function of $\theta_1, \ldots, \theta_{t+W-1}$

Error: $$E|x_t - x_t^*|^2 = E\left|x_t - x_t^t - x_t^{T-t}\right|^2$$
$$= E|x_t - x_t^t|^2 + E\left|x_t^{T-t}\right|^2$$
$$\geq a_{t,t+W}^2 + \cdots a_{t,T}^2$$

Key 2: $a_{t,t+W}, \ldots$ **decays exponentially**

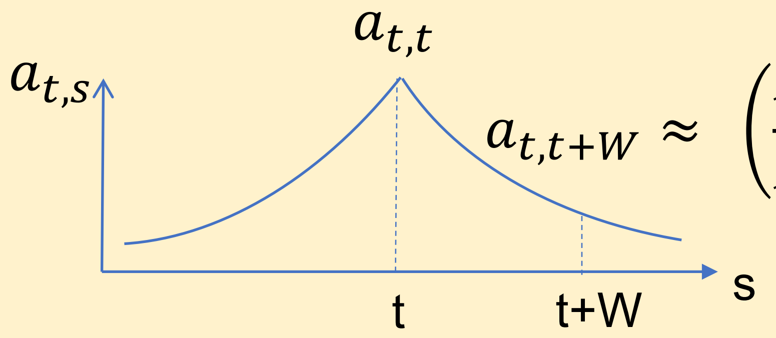

$$a_{t,t+W} \approx \left(\frac{\sqrt{Q_f} - 1}{\sqrt{Q_f} + 1}\right)^W$$

- **Regret**: By strong convexity,

$$E(C_1^T(x) - C_1^T(x^*)) \geq \alpha/2 \sum_{t=1}^T E\left||x_t - x_t^*|\right|^2 \geq \frac{a}{2}\left(\frac{\sqrt{Q_f} - 1}{\sqrt{Q_f} + 1}\right)^{2W} T \quad \text{Path length}$$

Numerical studies: Comparison with sub-MPC

- LTI system: two states
- Single input
- p=2

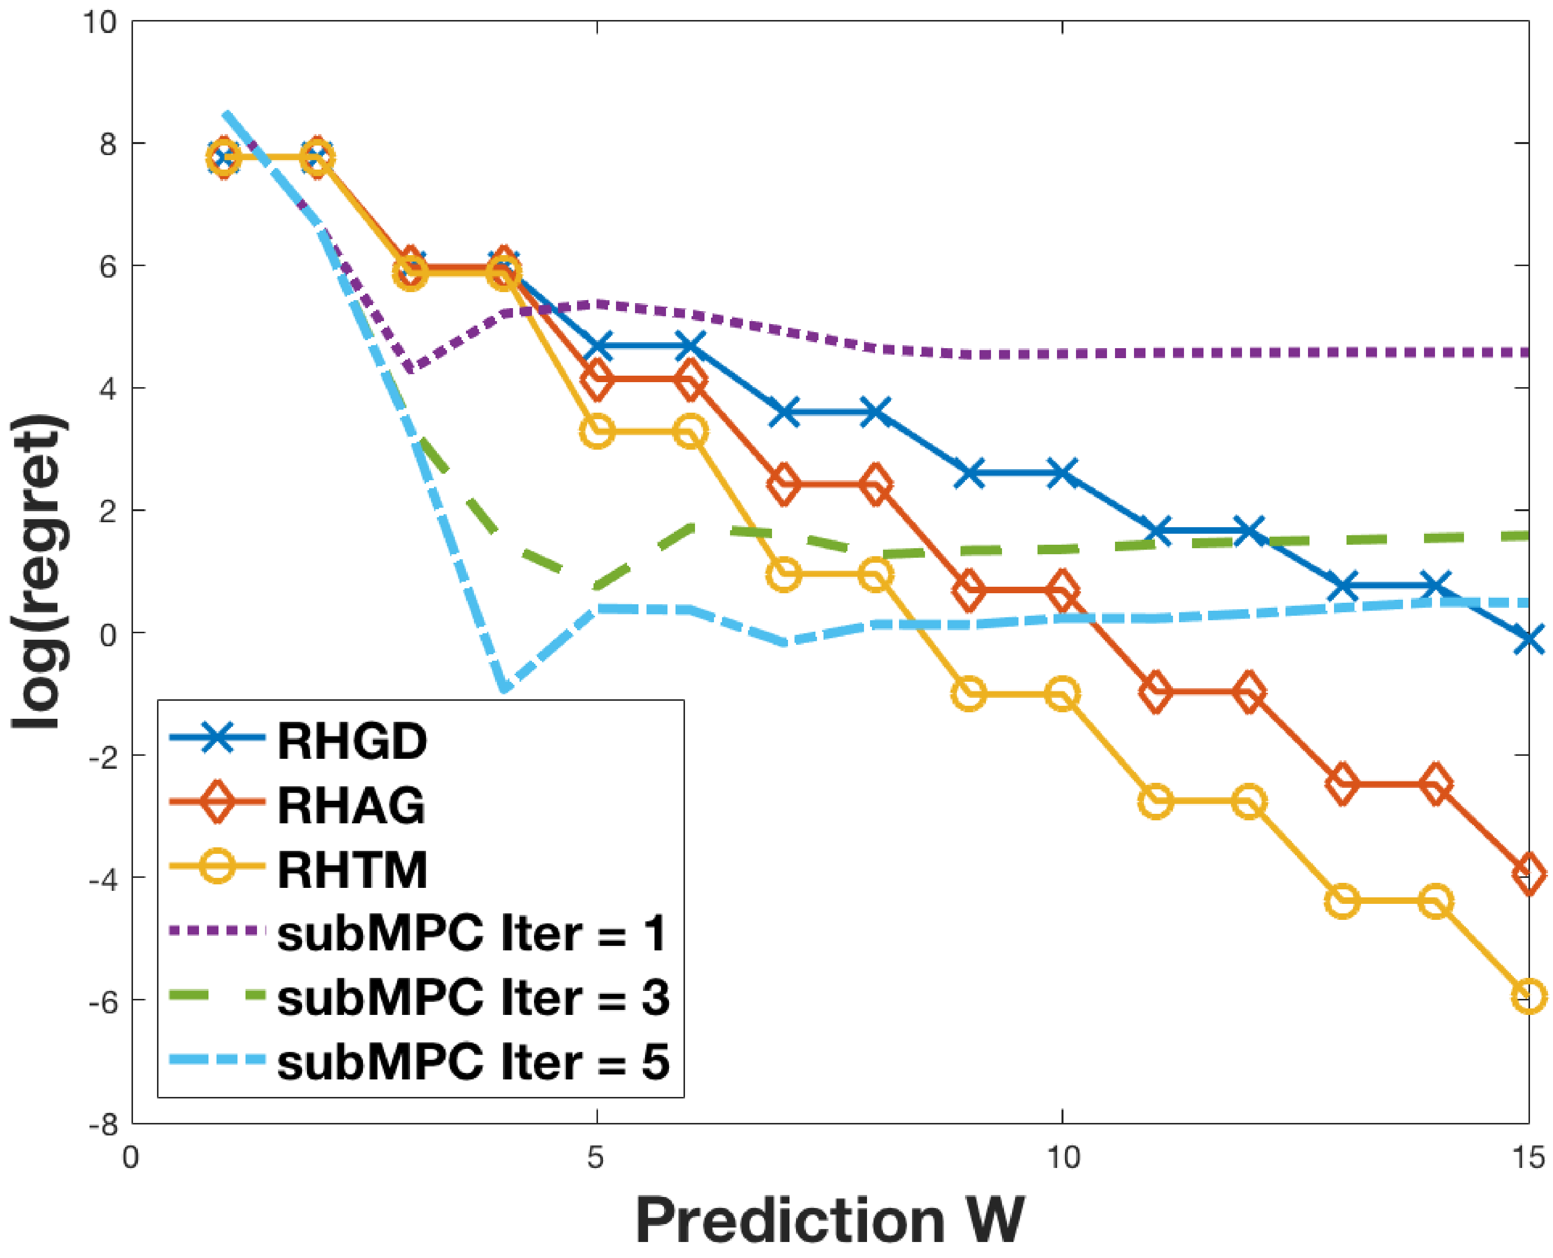

Numerical example: Two wheel robot dynamic tracking

$$\dot{x} = v \cos \delta$$

$$\dot{y} = v \sin \delta$$

$$\dot{\delta} = \omega$$

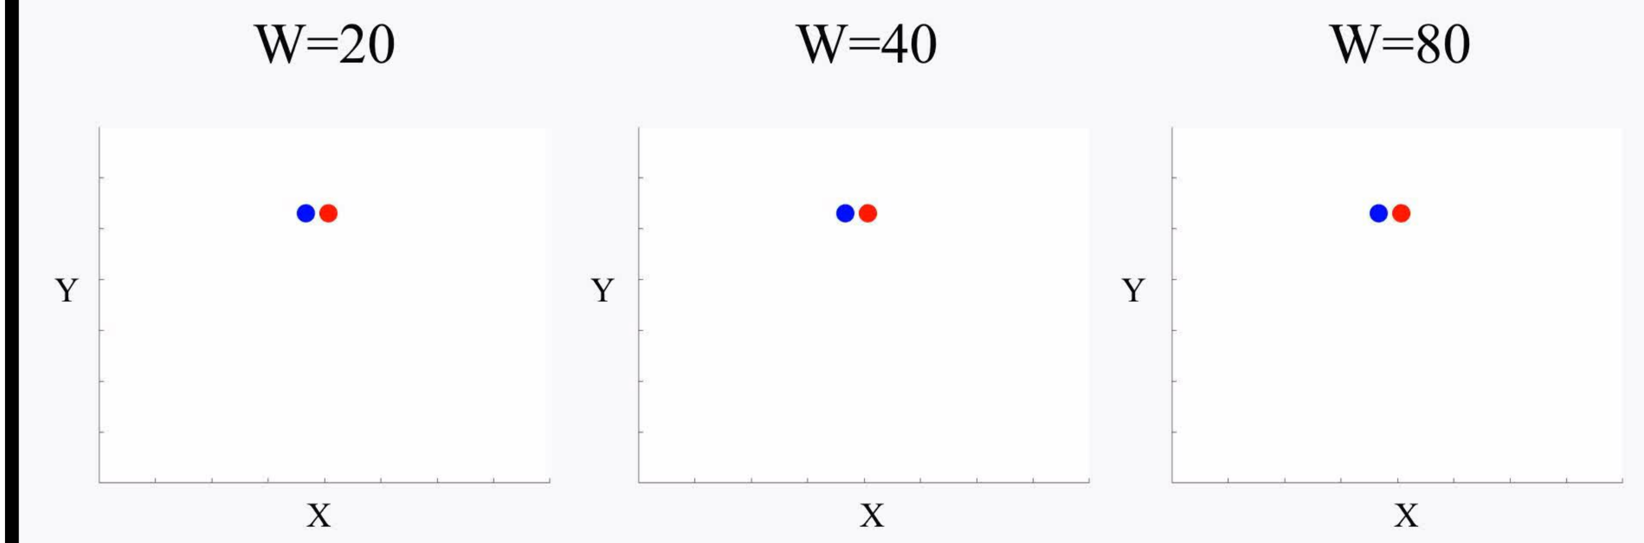

Two-wheel robot trajectory tracking

Summary

- Given any online optimal control algorithm without using prediction, additional W step of gradient calculations can further reduce the regret by $O(r^K)$, $K = \lfloor \frac{W}{p} \rfloor$
- Fundamental lower bound for any online algorithm: Dynamic regret **exponential** decays with $K = \lfloor \frac{W}{p} \rfloor$
- The dependence on W is nearly optimal

Thank you!

Ongoing and Future work:

- Stochastic/Noisy prediction
- Study optimal control with constraints; the dynamic regret bounds of MPC
- Online control problem for general dynamical systems including Markov decision models
- Other meaningful metrics for online performance analysis

Simulation 1: economic dispatch

(a) Power demand

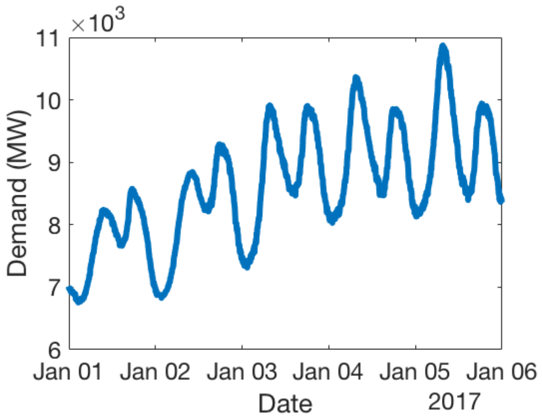

(b) Wind generation

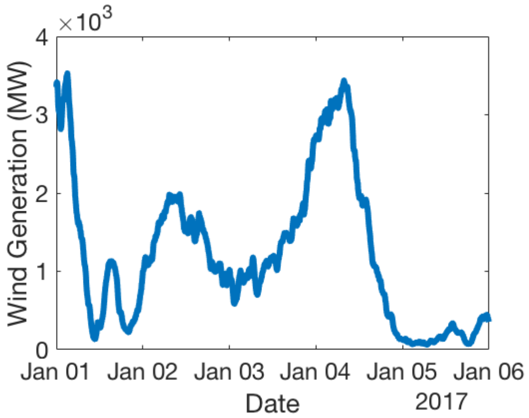

(c) Dynamic regrets

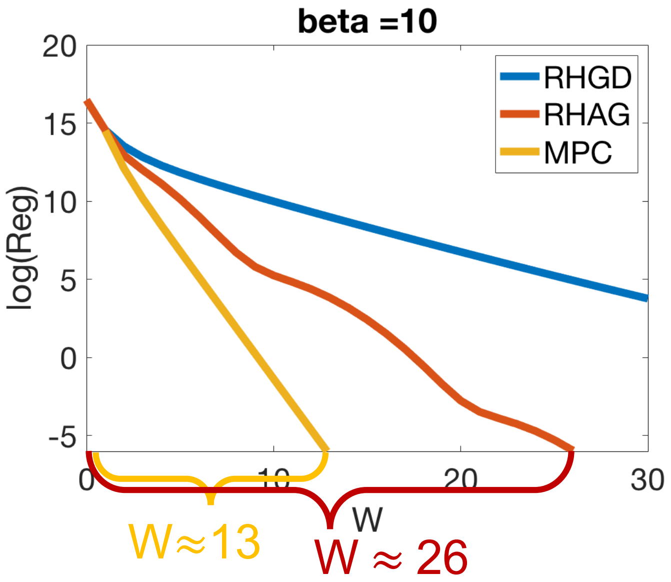

(d) Running time /second

TABLE I
RUNNING TIME OF RHGD, RHAG AND MPC

W \ $\mathscr{A}$	5	10
RHGD	8.8781×10^{-5}	1.4923×10^{-4}
RHAG	1.0416×10^{-4}	1.9052×10^{-4}
MPC	0.0105	0.0110

Setup:

- stage: 5 min
- horizon: 5 days (T=1440)
- data source: Bonneville Power Administration (BPA) controlled area
- n=3 generators
- Q_f = 11.75
- stage cost:

1) cost of generation
2) cost of uncleared demand

- Matlab (quadprog())

https://transmission.bpa.gov/business/operations/wind/

Simulation 2

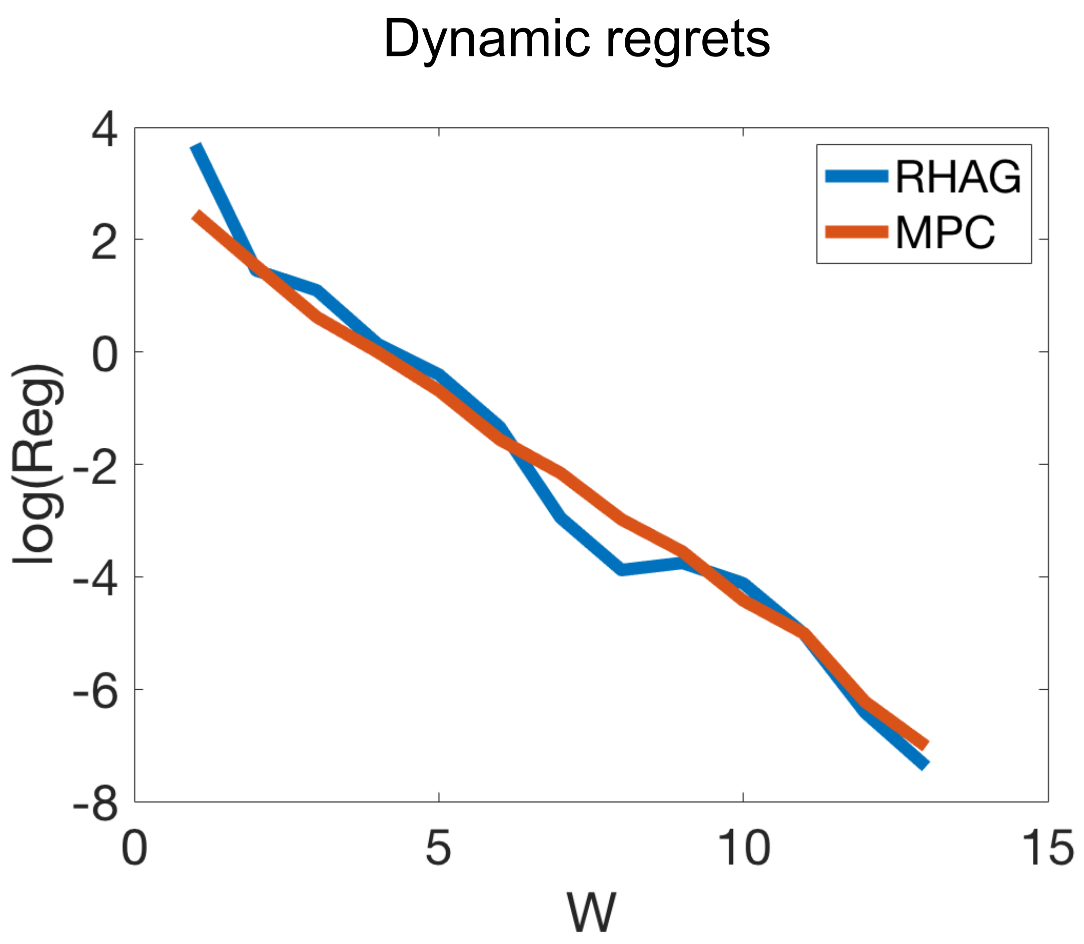

Setup:

- horizon: T=16
- Scalar decision variable
- $Q_f \approx 47$
- stage cost:
$$\frac{1}{2}(x_t - \theta_t)^2$$
- θ_t picked from {0,4}

Summary: Online Learning and Decision Making

Part II

Online decision-making under a time-varying environment

Part I

Online learning and control to learn & handle unknown uncertainties

???

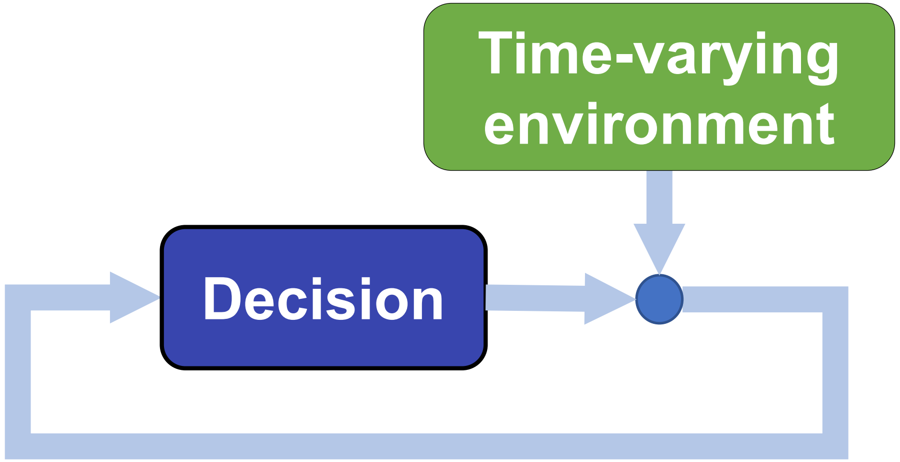

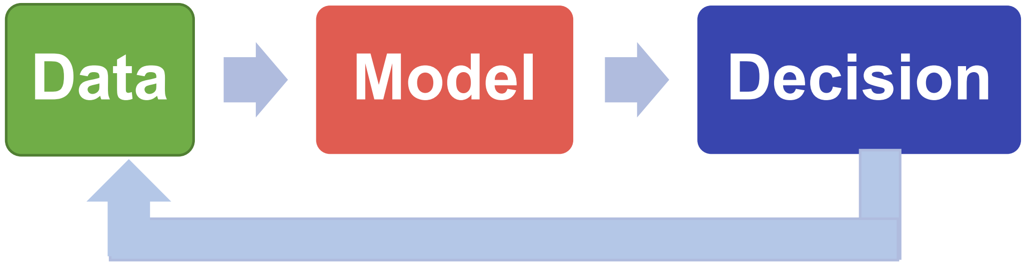

Thank you!

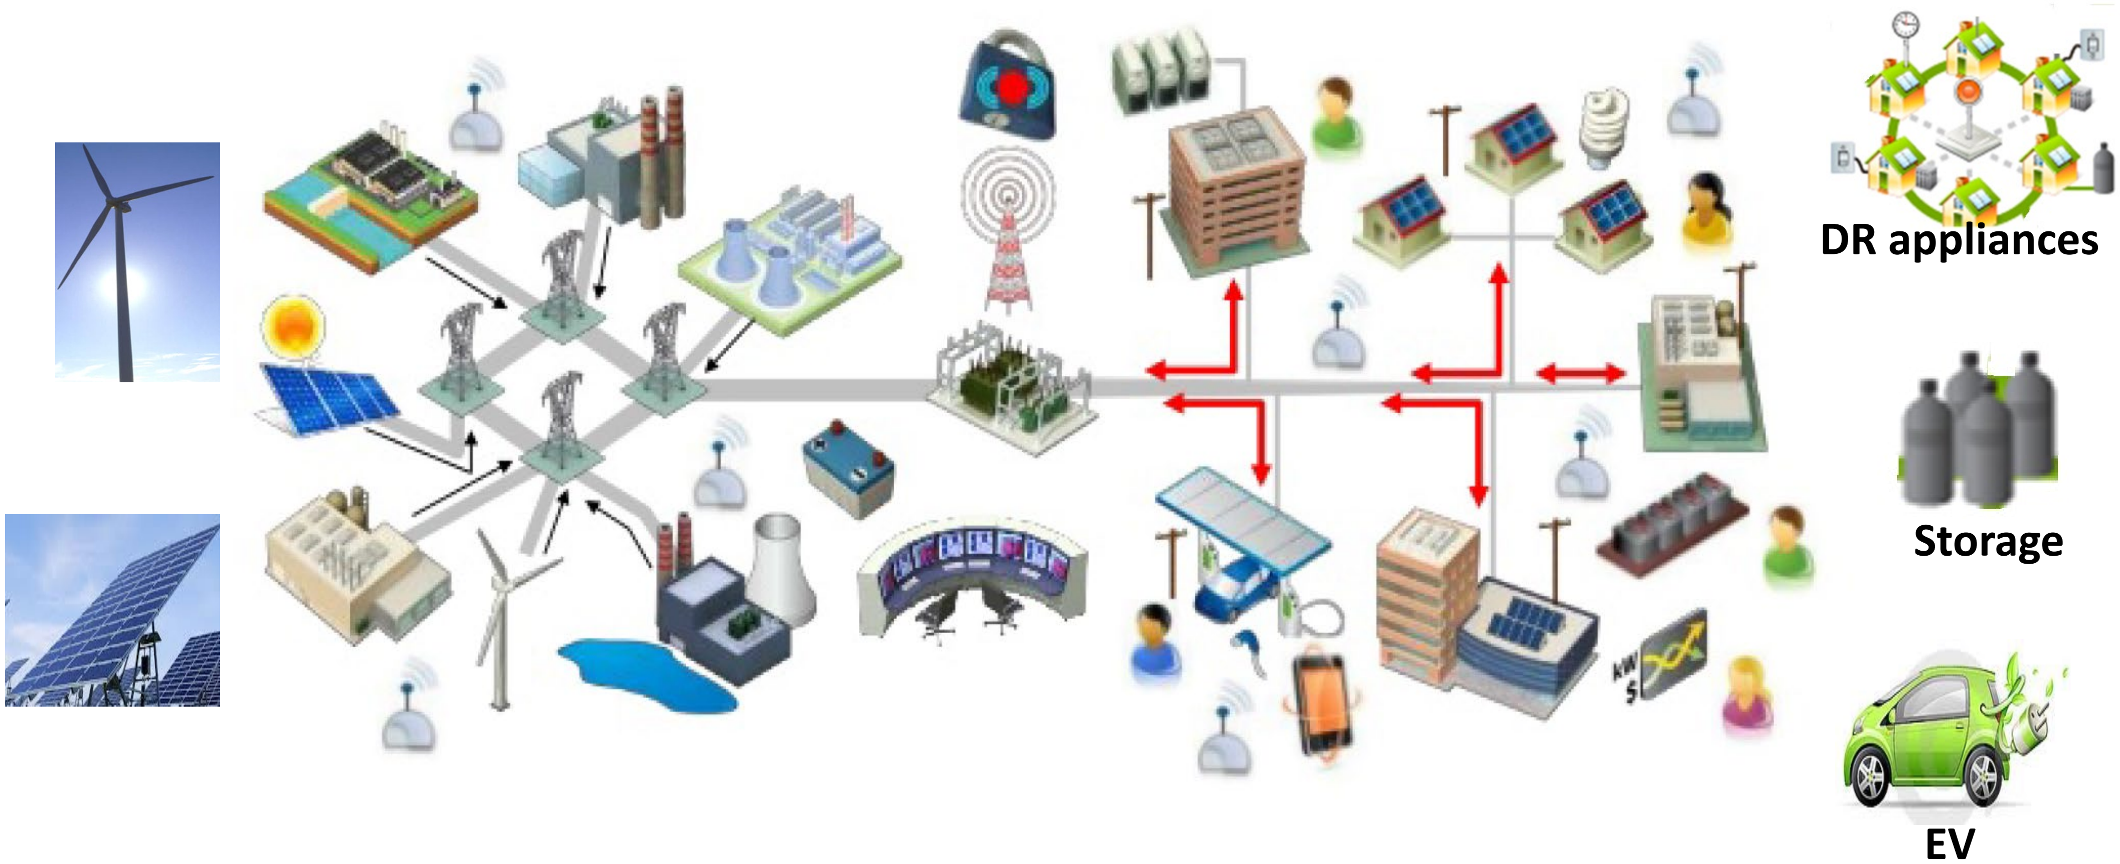

Supply = Demand

Increasing uncertainties in the grid

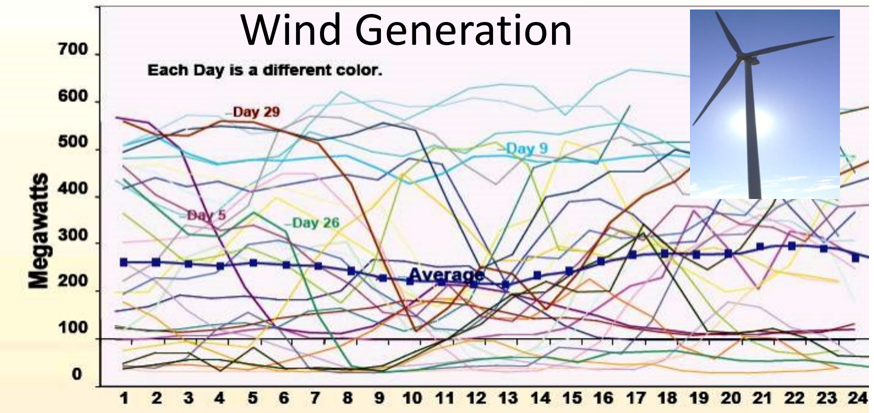

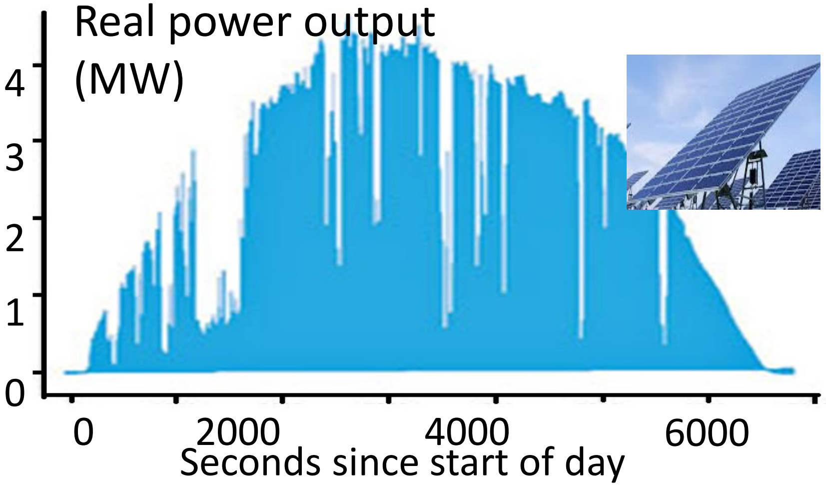

Random, Intermittent

DR appliances

Storage

EV

Responsive ← **Unresponsive**

Match supply

Increasing uncertainties in the grid

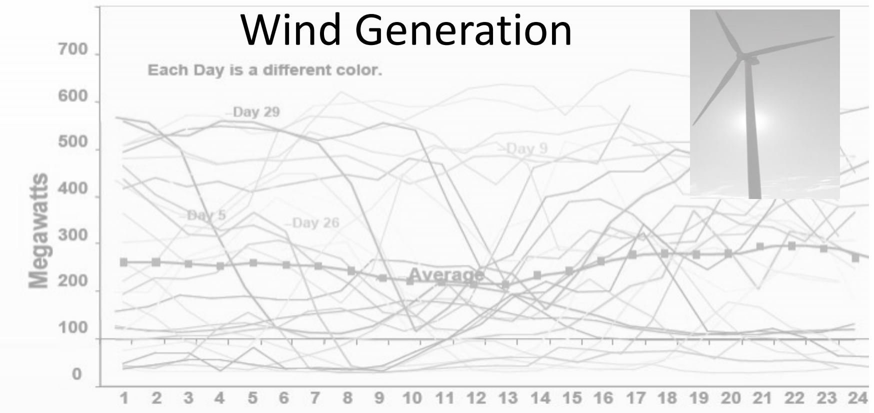

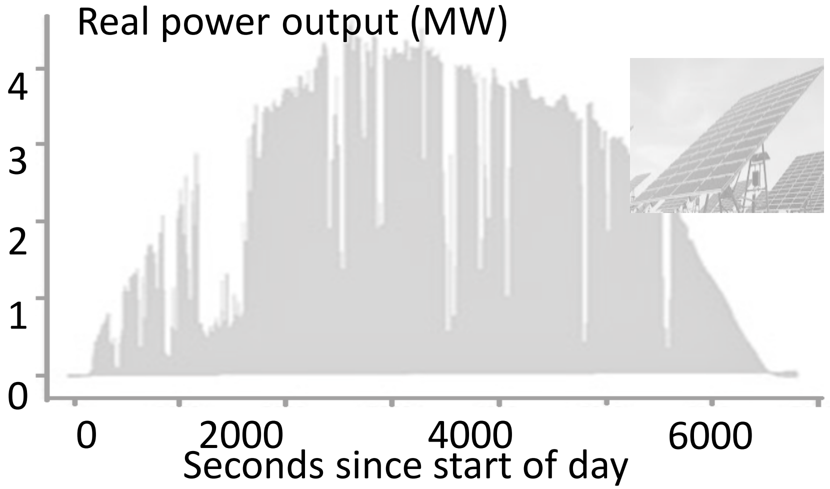

Random, Intermittent

??

Preferred room temperature

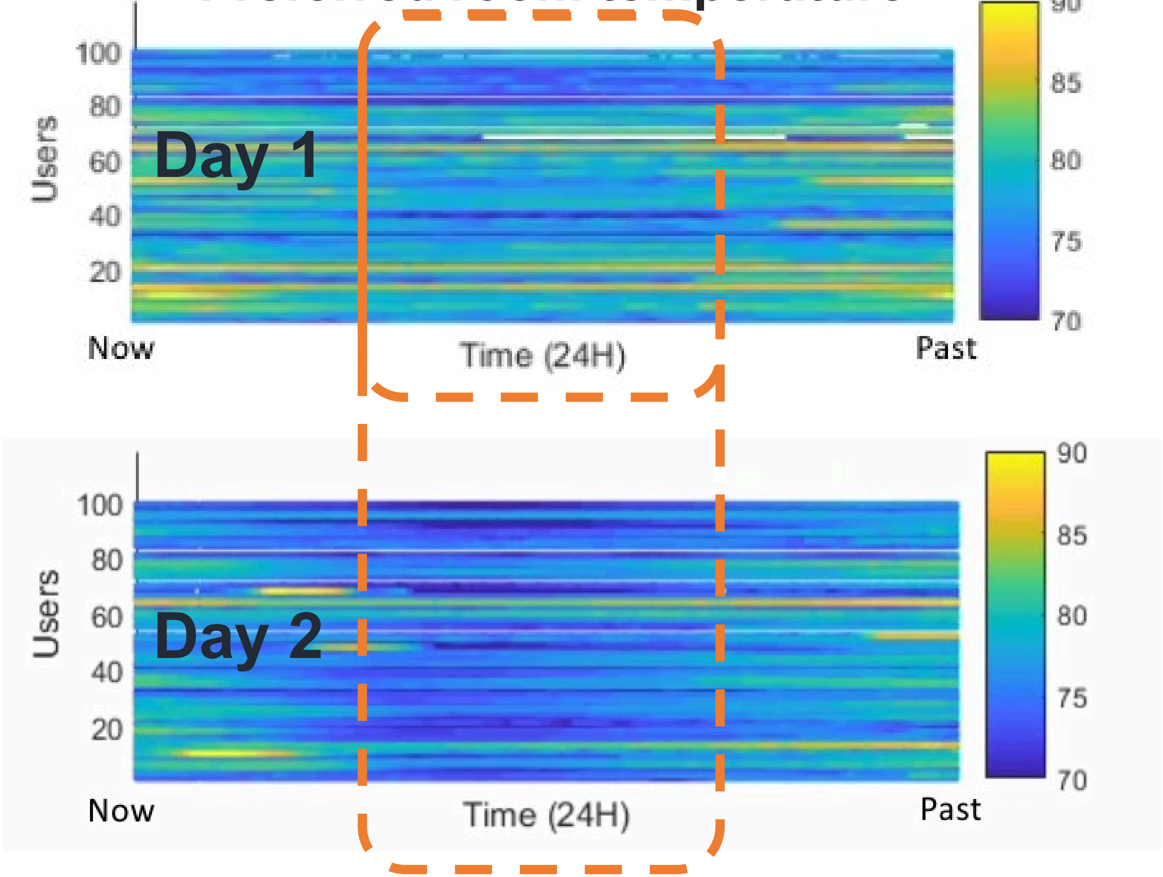

Different and **random** preferences:
Unknown to the utility/DR company

Source: ThinkEco pilot study 2016 in Connecticut

Online learning and decision making

Online decision-making under a time-varying environment

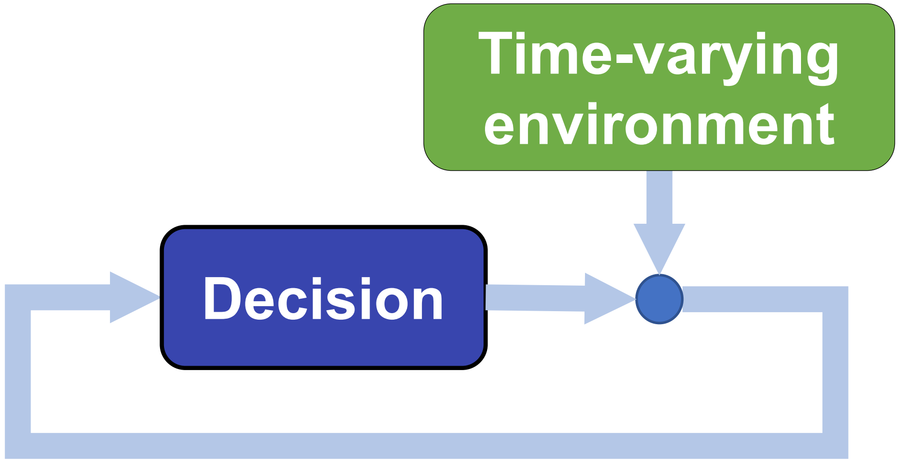

Online learning and control to learn & handle unknown uncertainties

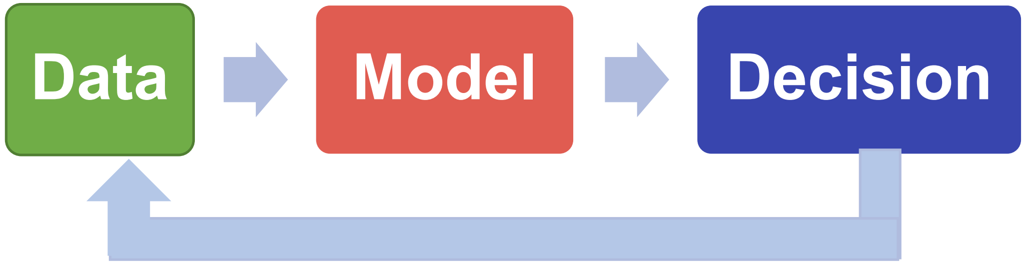

Online learning and decision making

Online decision-making under a time-varying environment

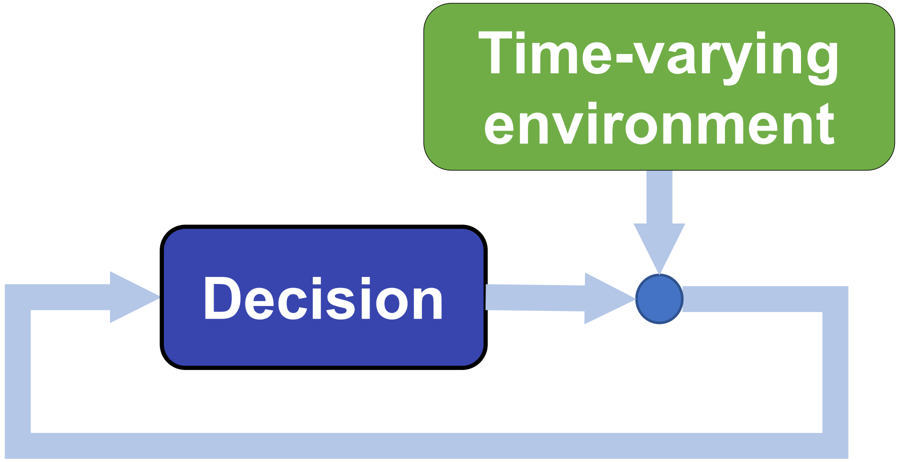

Online learning and control to learn & handle unknown uncertainties

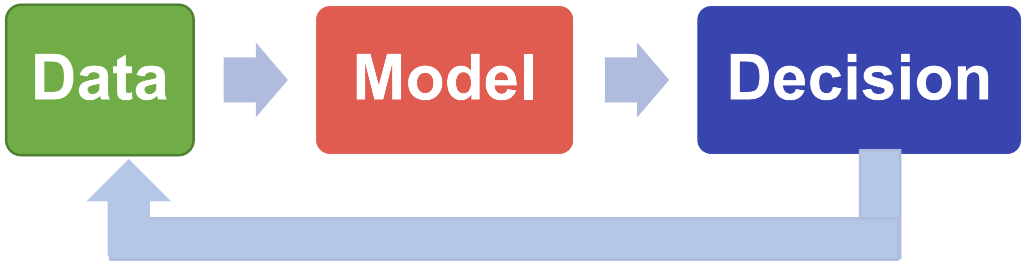

Today's talk

Online control (online optimization with switching cost) using Prediction

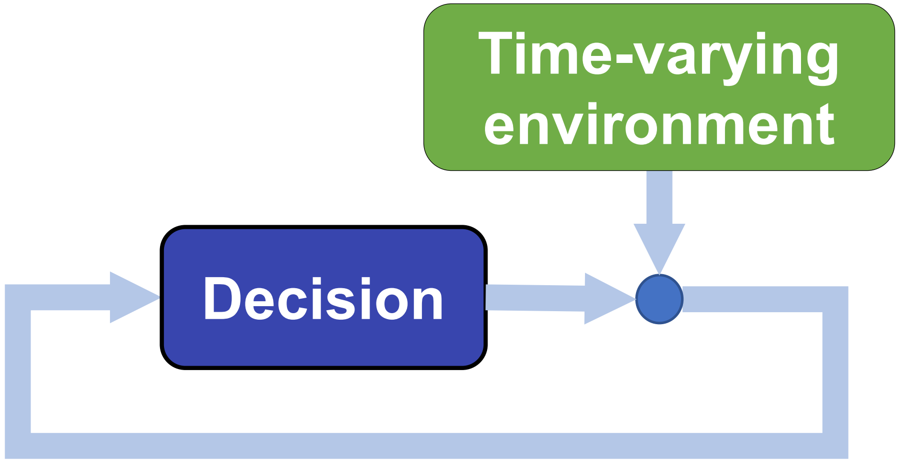

Online learning and control to learn & handle unknown uncertainties