Competitive Algorithms for Online Control

Yisong Yue

This Talk

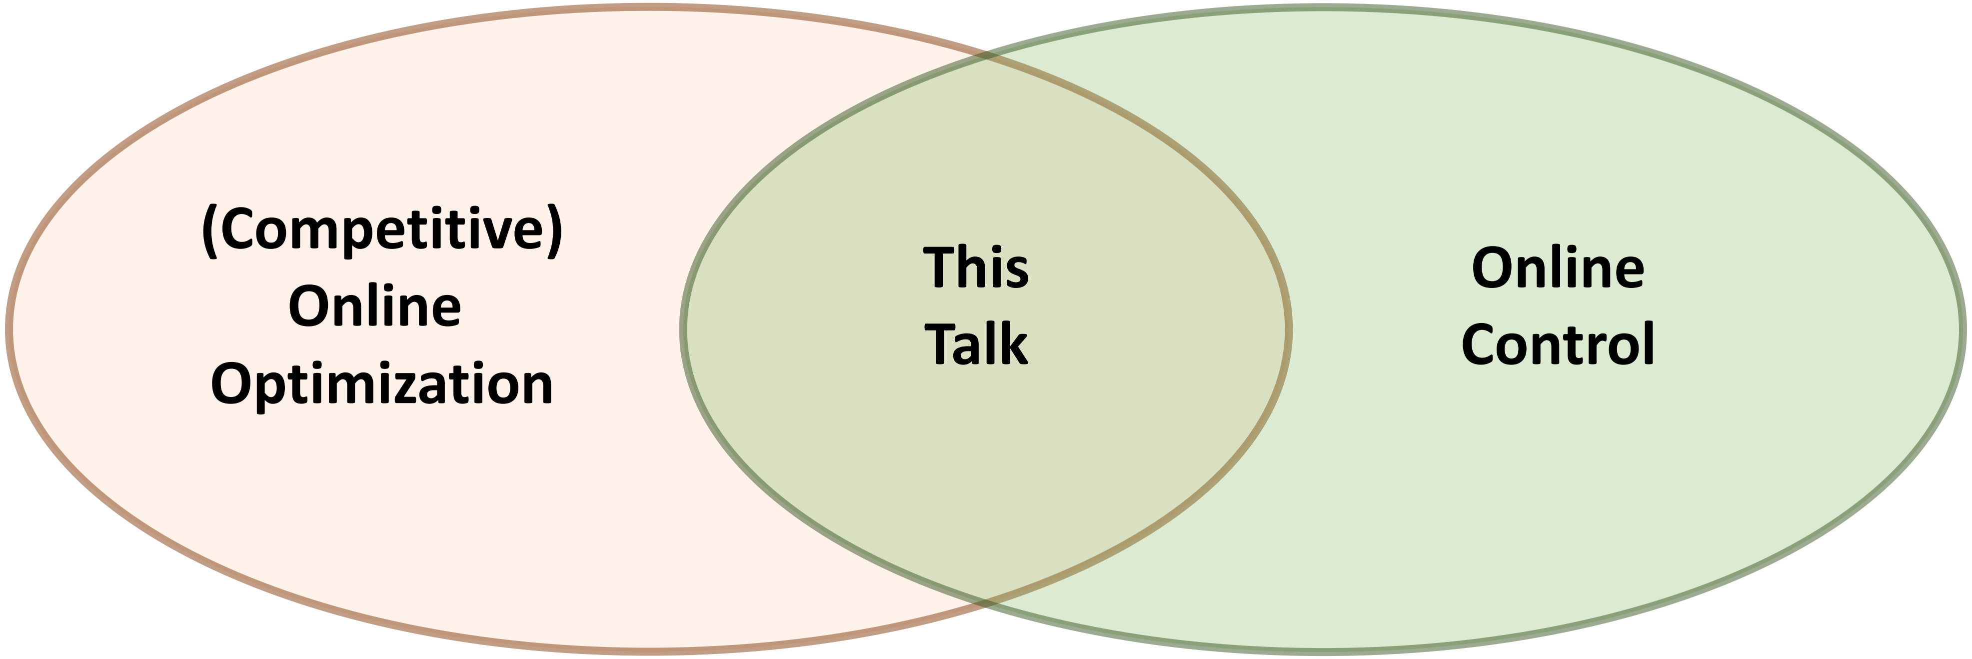

- Bulk of work in collaboration with Adam Wierman's group.

Adam Wierman

Online Optimization

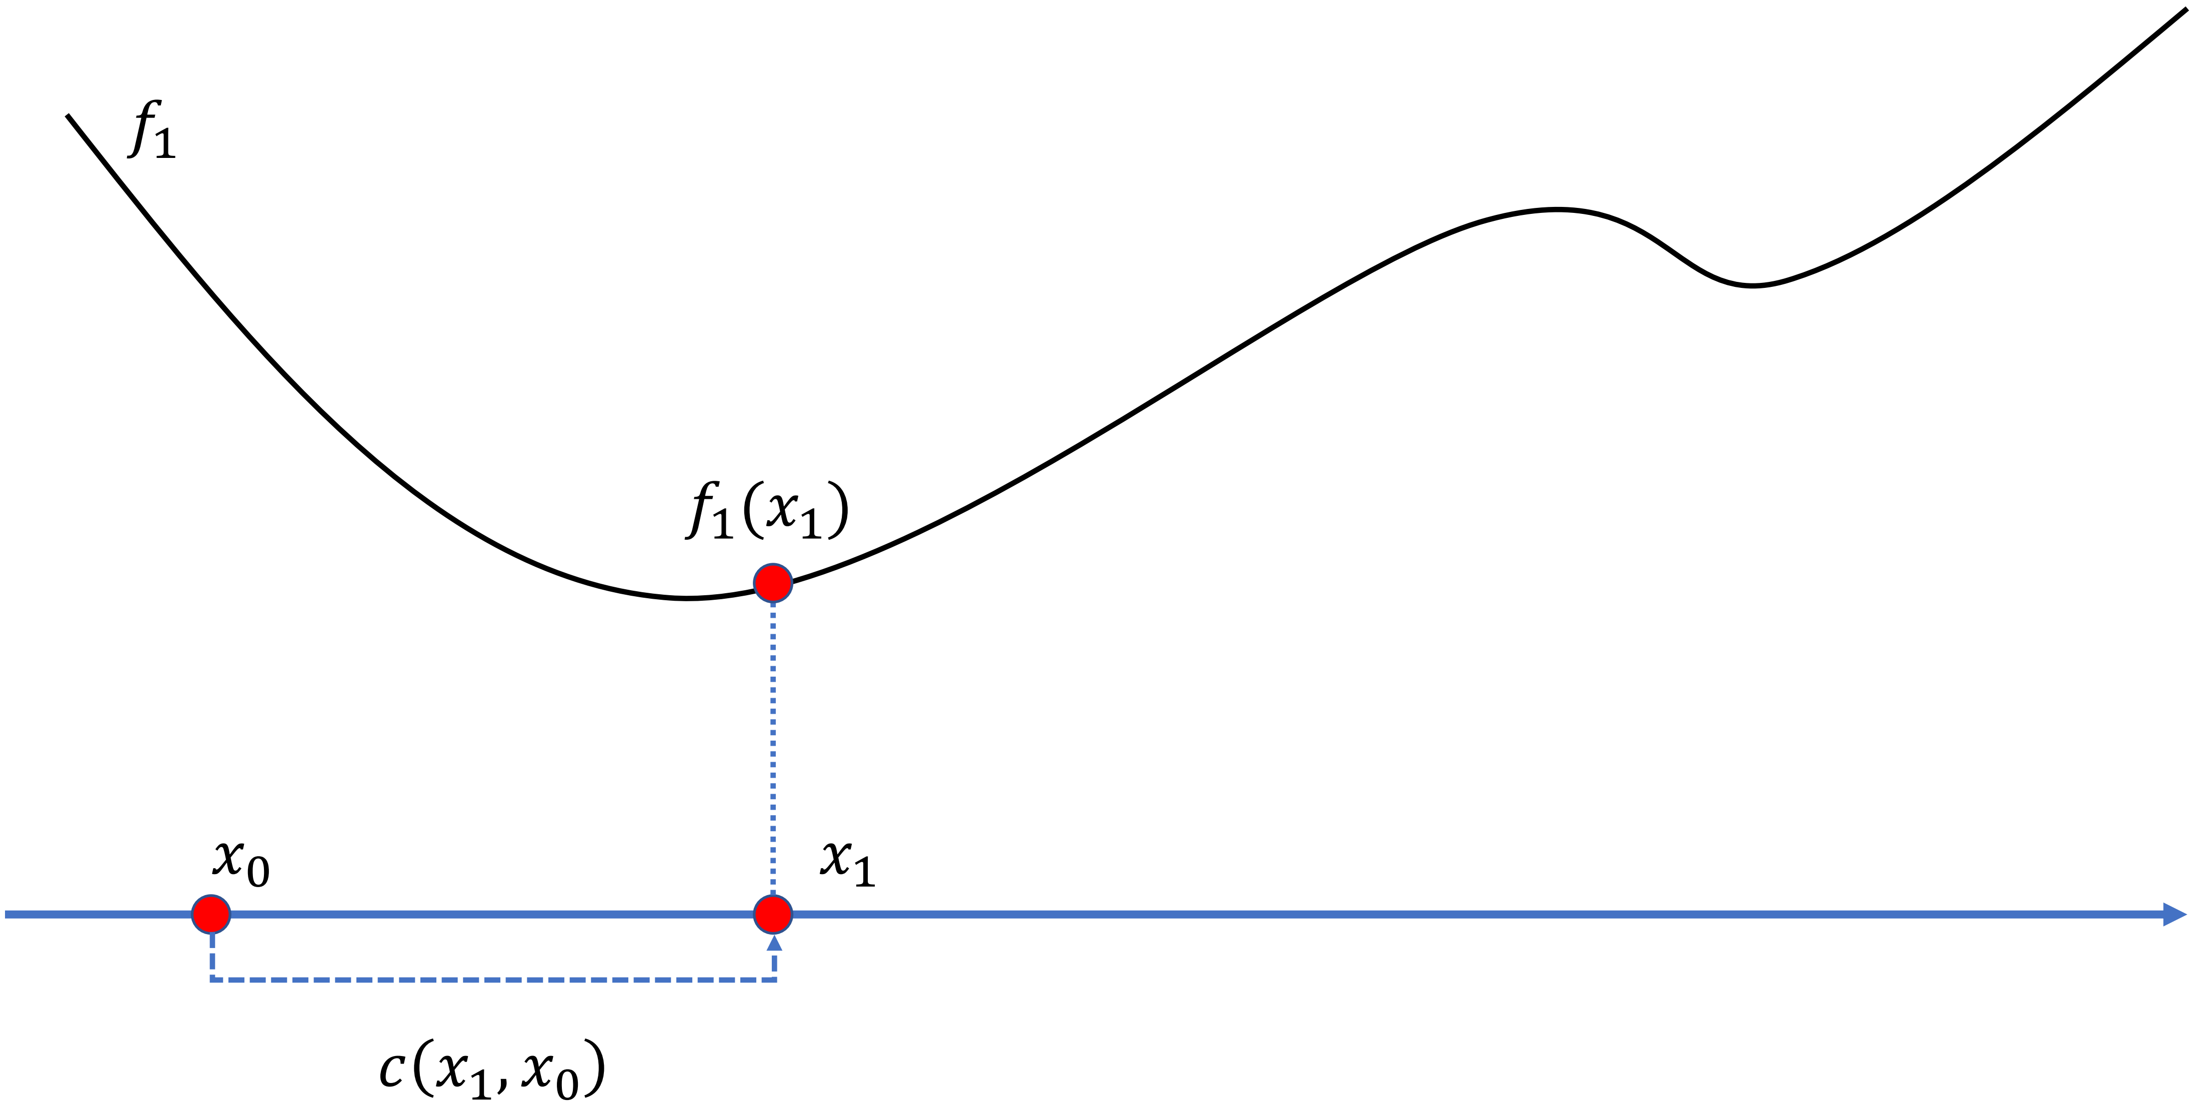

Material from Adam Wierman

Online Optimization

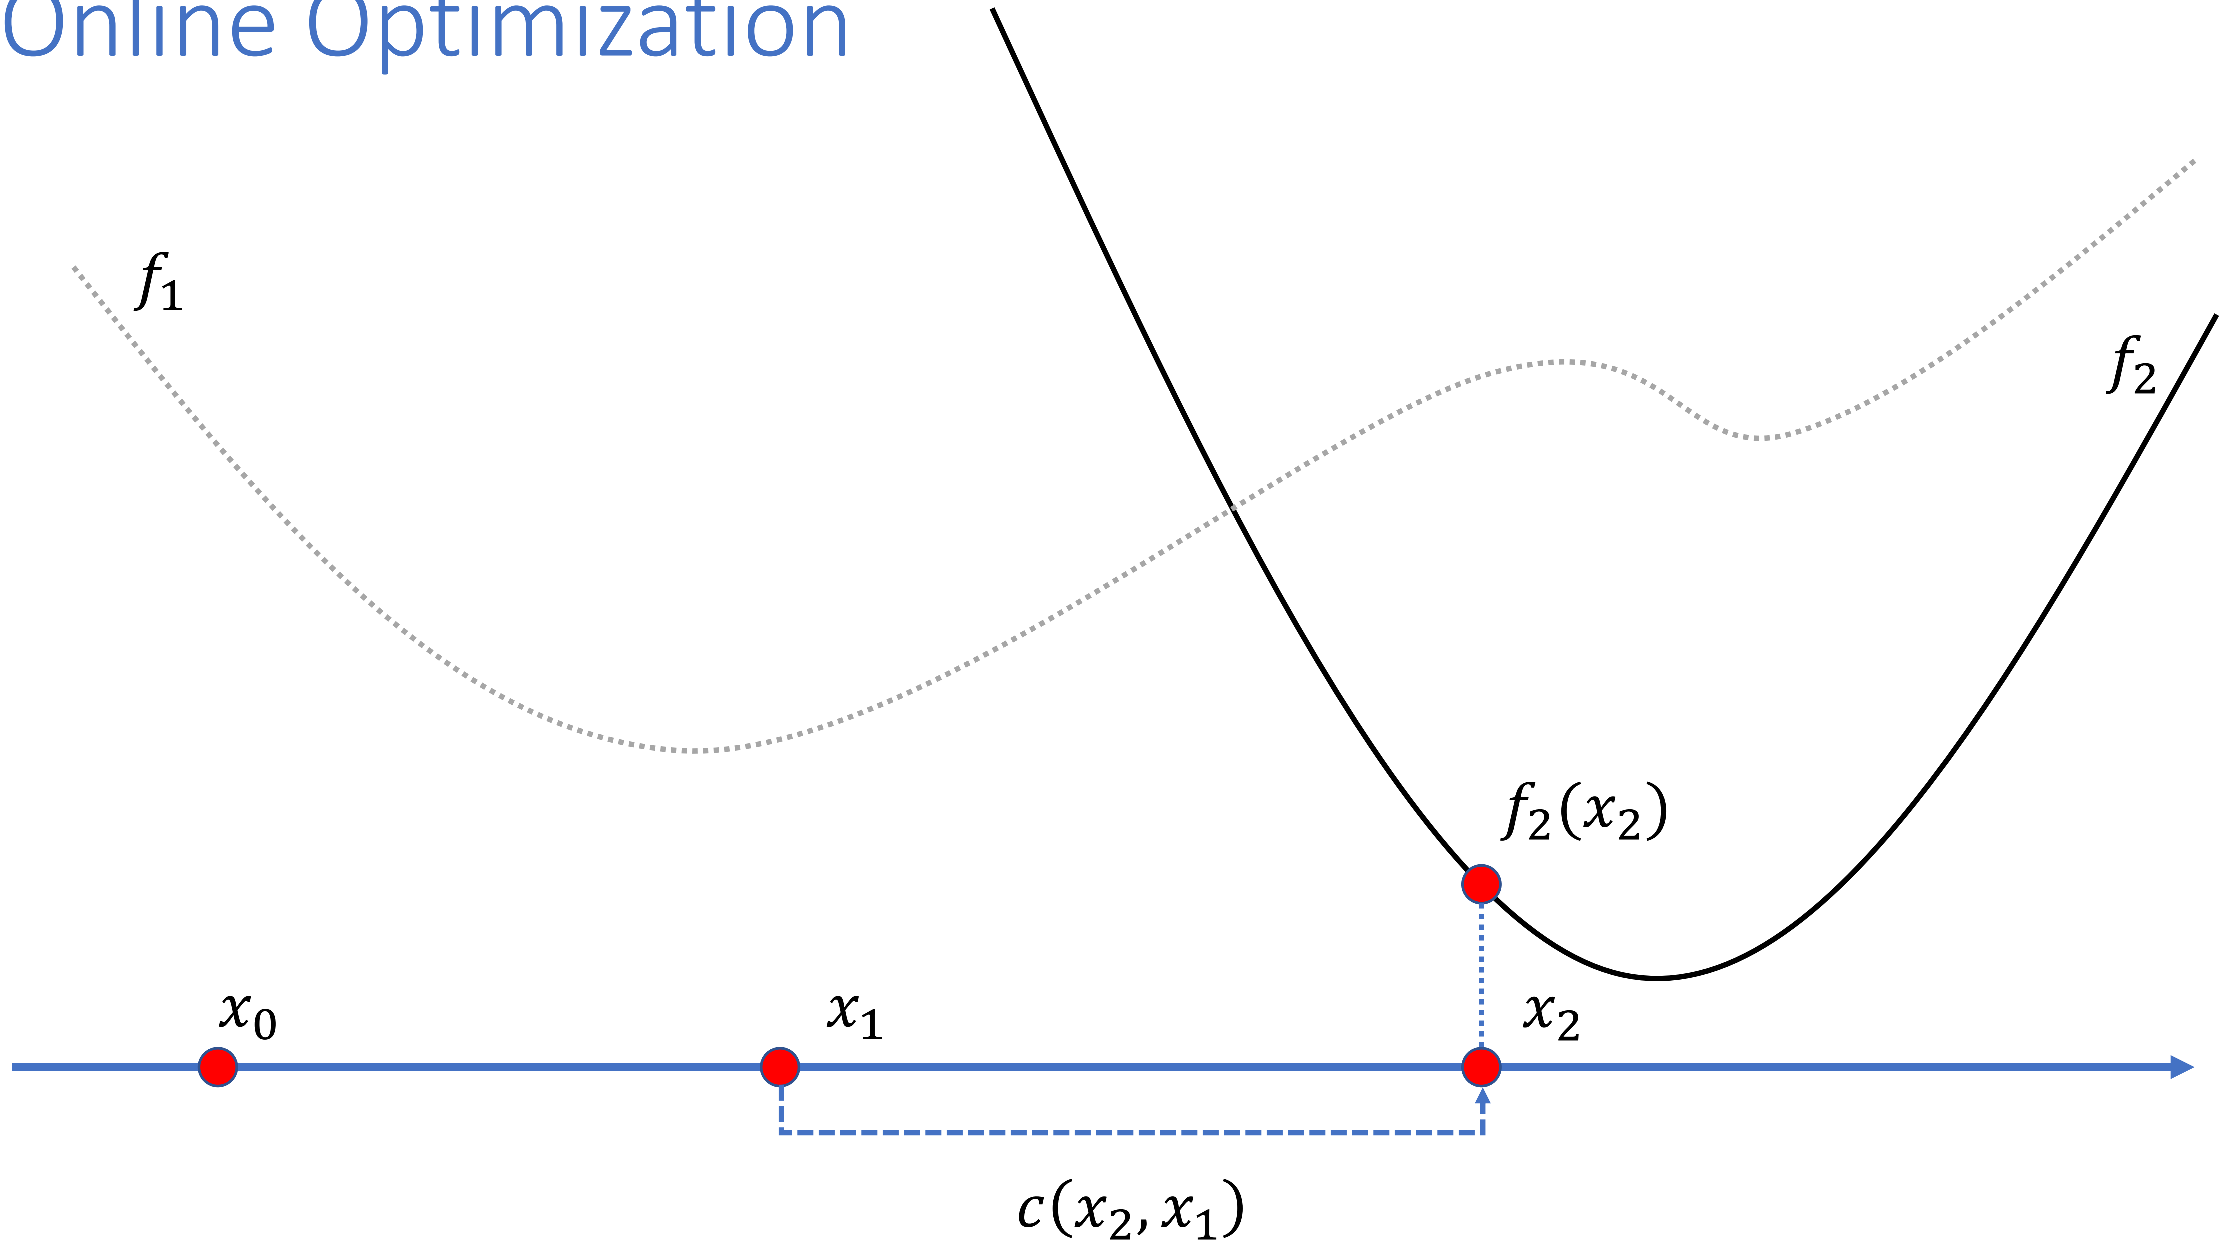

Material from Yiheng Lin

Online Optimization

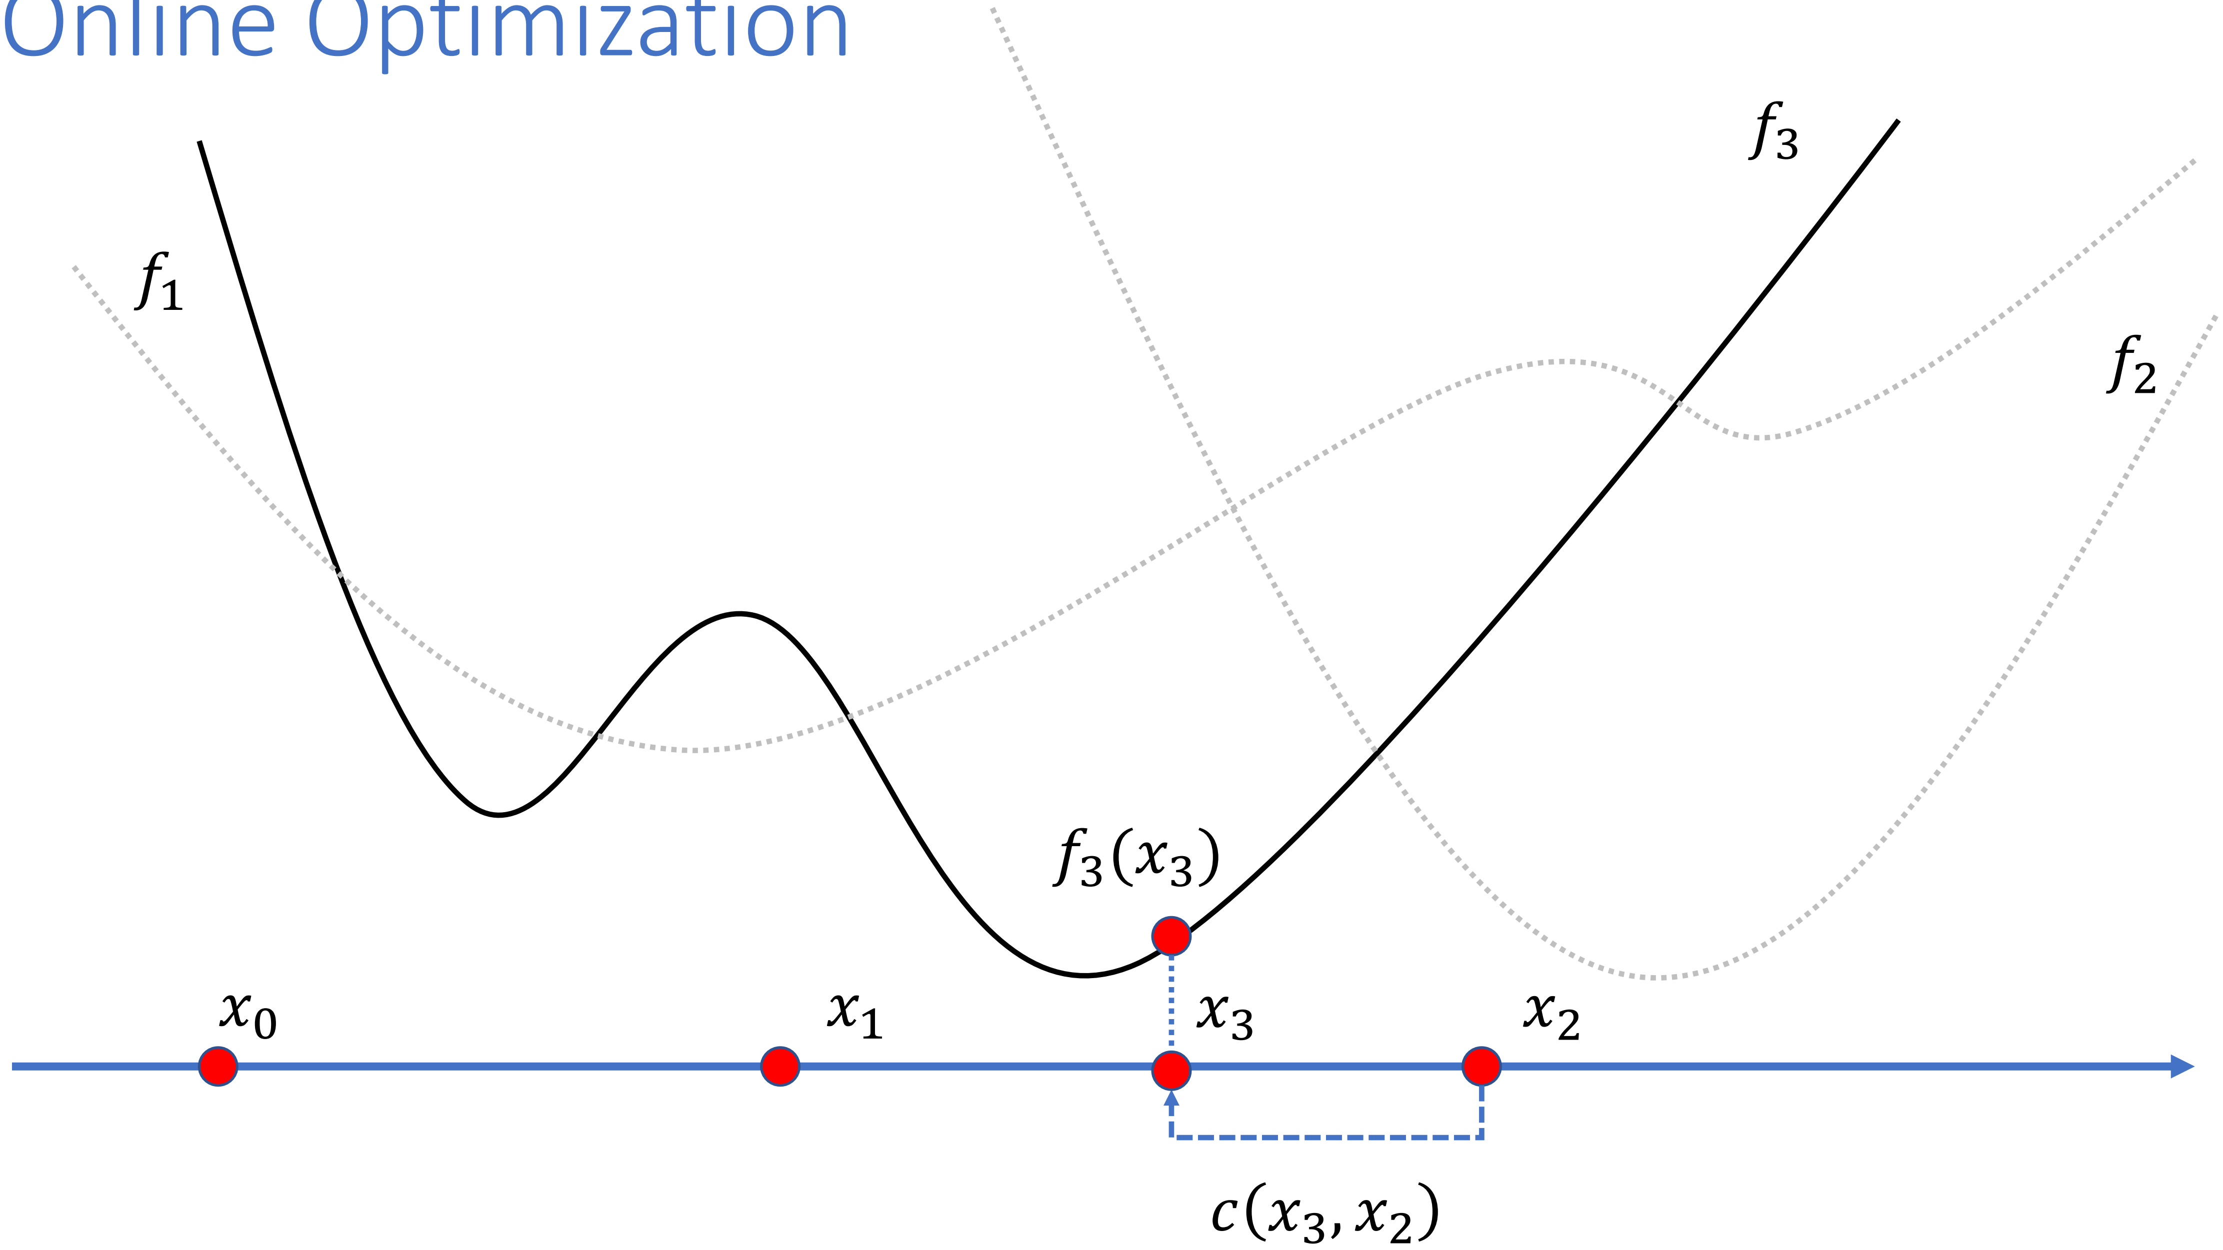

Material from Yiheng Lin

Primer: Smoothed Online Convex Optimization (SOCO)

https://arxiv.org/abs/1803.10366

$$\min_{x_t \in F_t} \sum_{t=1}^{T} f_t(x_t) + c(x_t, x_{t-1})$$

(observe f_t at time t)

(convex) hitting cost

movement (switching) cost

Goal: Design algorithms to minimize the total cost.

Performance metric: competitive ratio: $sup_{\{f_t\}} \frac{cost(ALG)}{cost(OPT)}$

Material from Yiheng Lin

Comments on Competitive Ratio

$$sup_{\{f_t\}} \frac{cost(ALG)}{cost(OPT)}$$

- Stronger criteria than regret
- Offline OPT gets to see the future
- Goal is typically to achieve constant competitive ratio
  - (note that CCR typically leads to linear regret)

Example 1: Sustainable Data Centers

Can a data center run (almost) entirely on renewable sources?

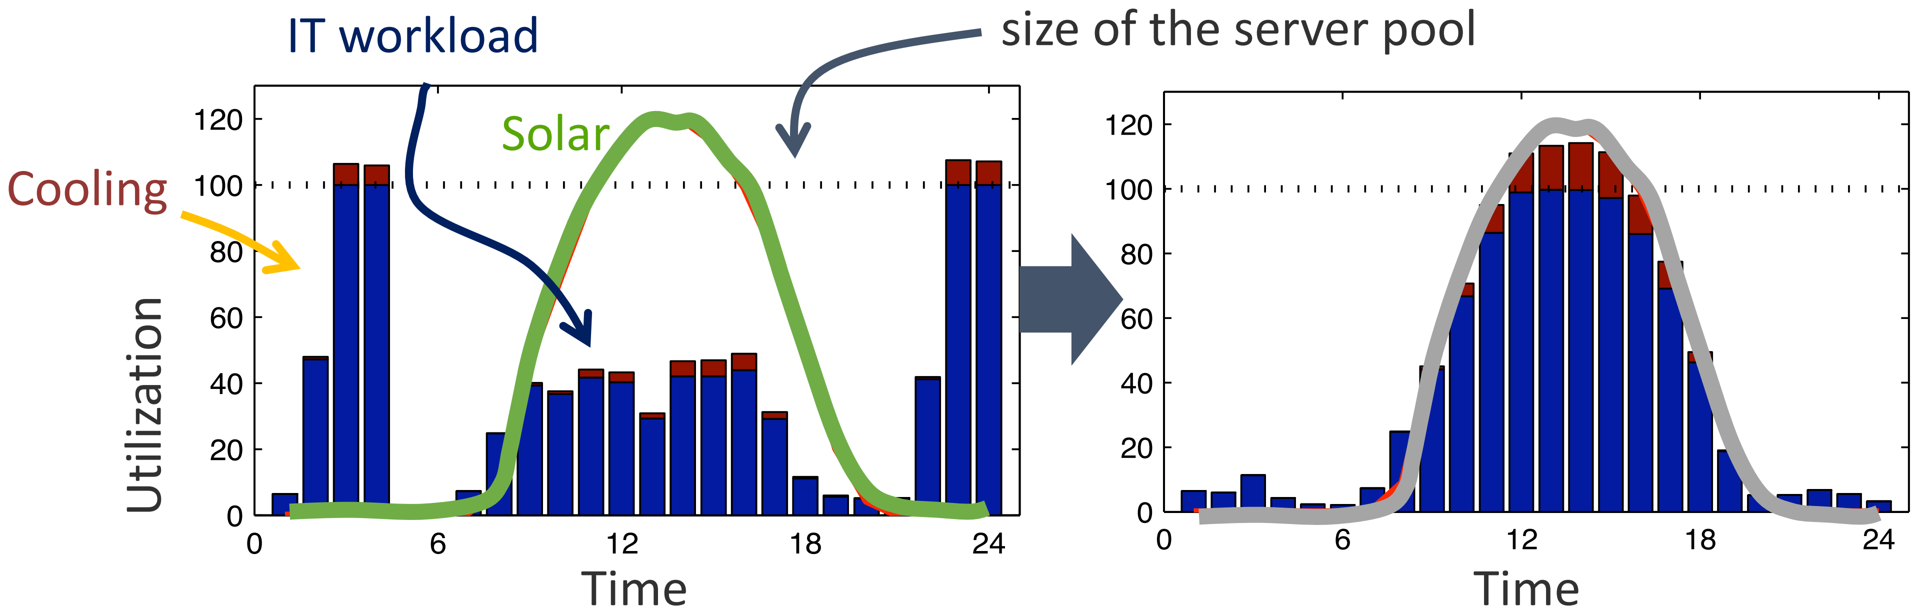

Requires dynamic rightsizing of capacity and smart deferral of workloads

Material from Adam Wierman

Example 1: Sustainable Data Centers

Can a data center run (al… …ces?

Cool

Utilizati…

R… …of workloads

24

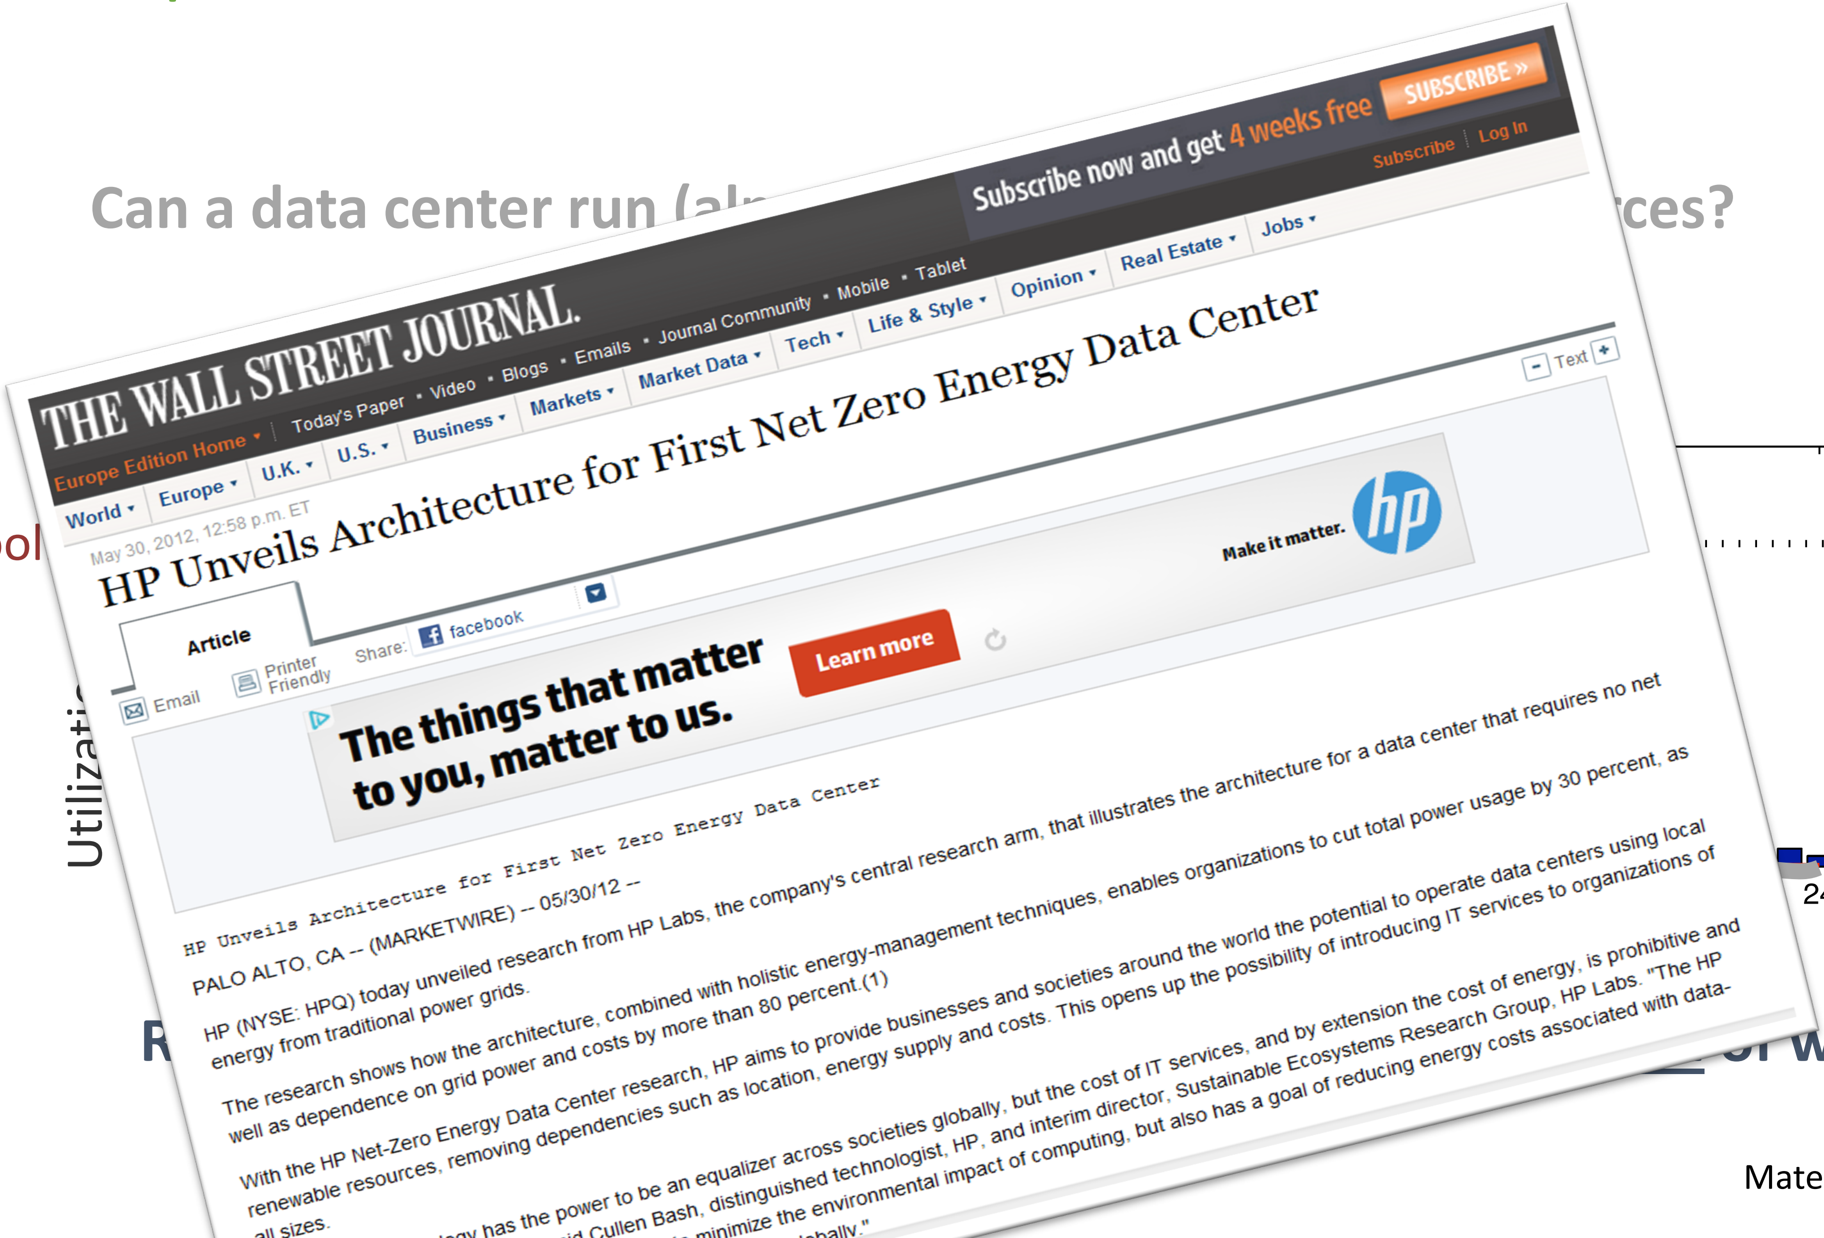

THE WALL STREET JOURNAL.

Europe Edition Home · Today's Paper · Video · Blogs · Emails · Journal Community · Mobile · Tablet

Subscribe now and get 4 weeks free SUBSCRIBE »

Subscribe | Log In

World · Europe · U.K. · U.S. · Business · Markets · Market Data · Tech · Life & Style · Opinion · Real Estate · Jobs

May 30, 2012, 12:58 p.m. ET

HP Unveils Architecture for First Net Zero Energy Data Center

Article · Email · Printer Friendly · Share: facebook

Text

HP Unveils Architecture for First Net Zero Energy Data Center

PALO ALTO, CA -- (MARKETWIRE) -- 05/30/12 --

HP (NYSE: HPQ) today unveiled research from HP Labs, the company's central research arm, that illustrates the architecture for a data center that requires no net energy from traditional power grids.

The research shows how the architecture, combined with holistic energy-management techniques, enables organizations to cut total power usage by 30 percent, as well as dependence on grid power and costs by more than 80 percent.(1)

With the HP Net-Zero Energy Data Center research, HP aims to provide businesses and societies around the world the potential to operate data centers using local renewable resources, removing dependencies such as location, energy supply and costs. This opens up the possibility of introducing IT services to organizations of all sizes.

…logy has the power to be an equalizer across societies globally, but the cost of IT services, and by extension the cost of energy, is prohibitive and …id Cullen Bash, distinguished technologist, HP, and interim director, Sustainable Ecosystems Research Group, HP Labs. "The HP …minimize the environmental impact of computing, but also has a goal of reducing energy costs associated with data-…lobally."

Material from Adam Wierman

Example 2: Tracking Control under Predictable Disturbances

Connection to Convex Body Chasing

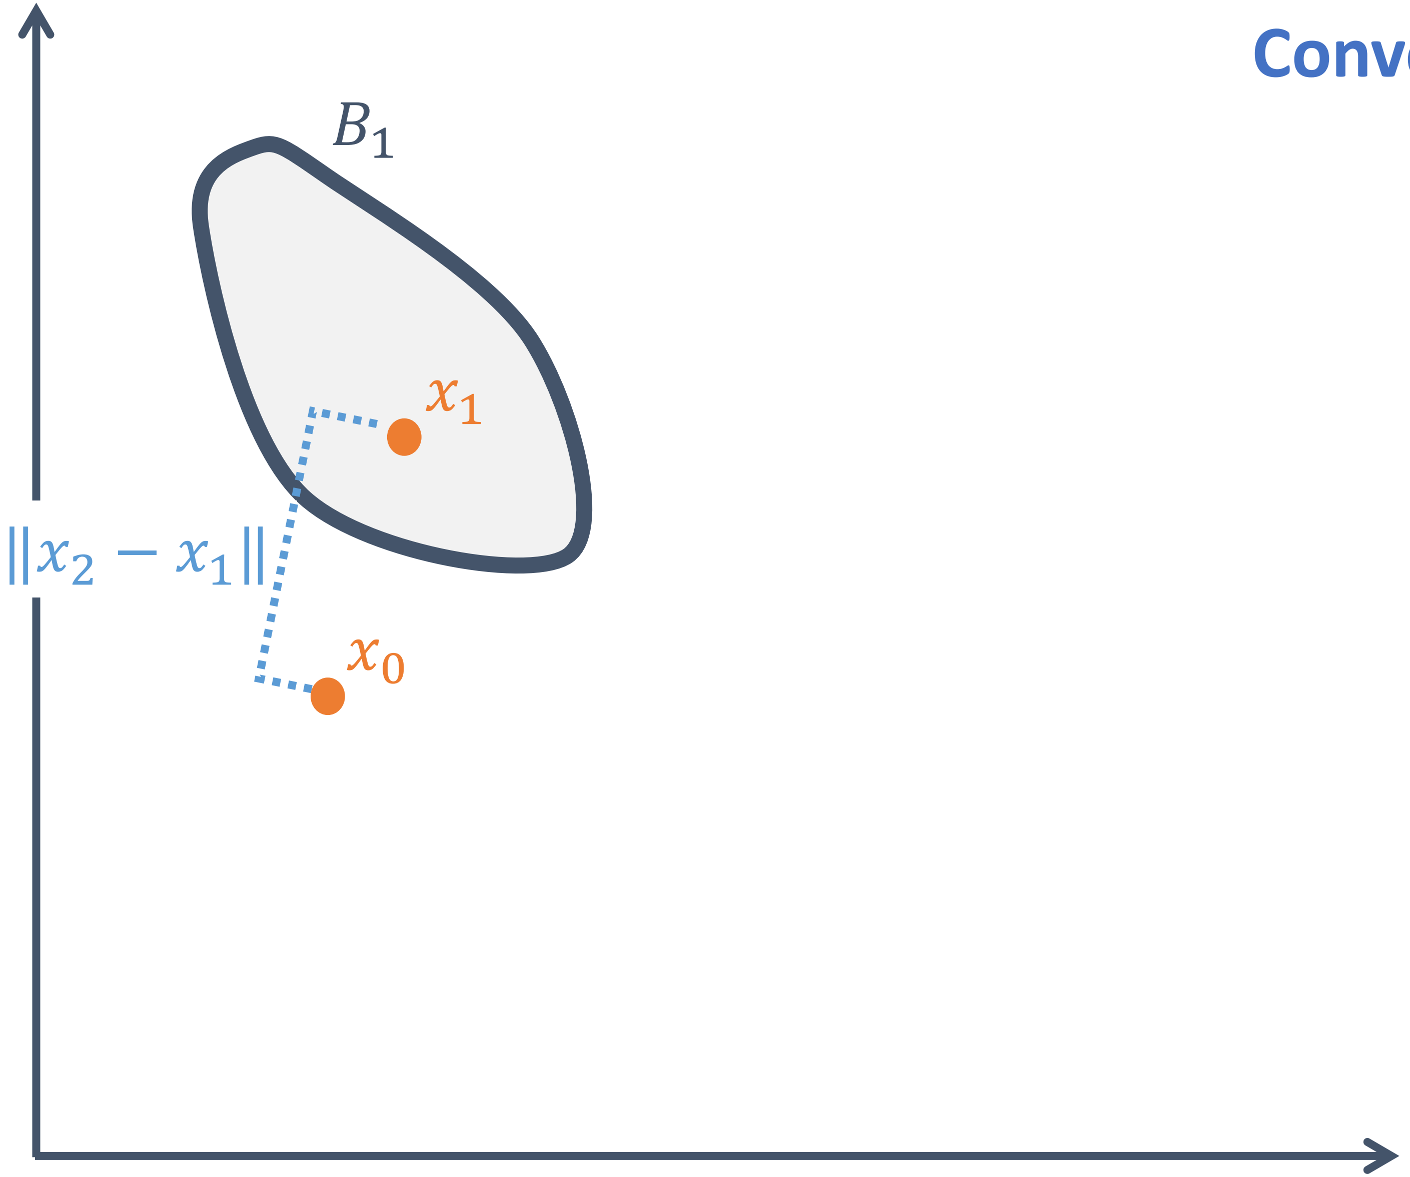

Material from Adam Wierman

Connection to Convex Body Chasing

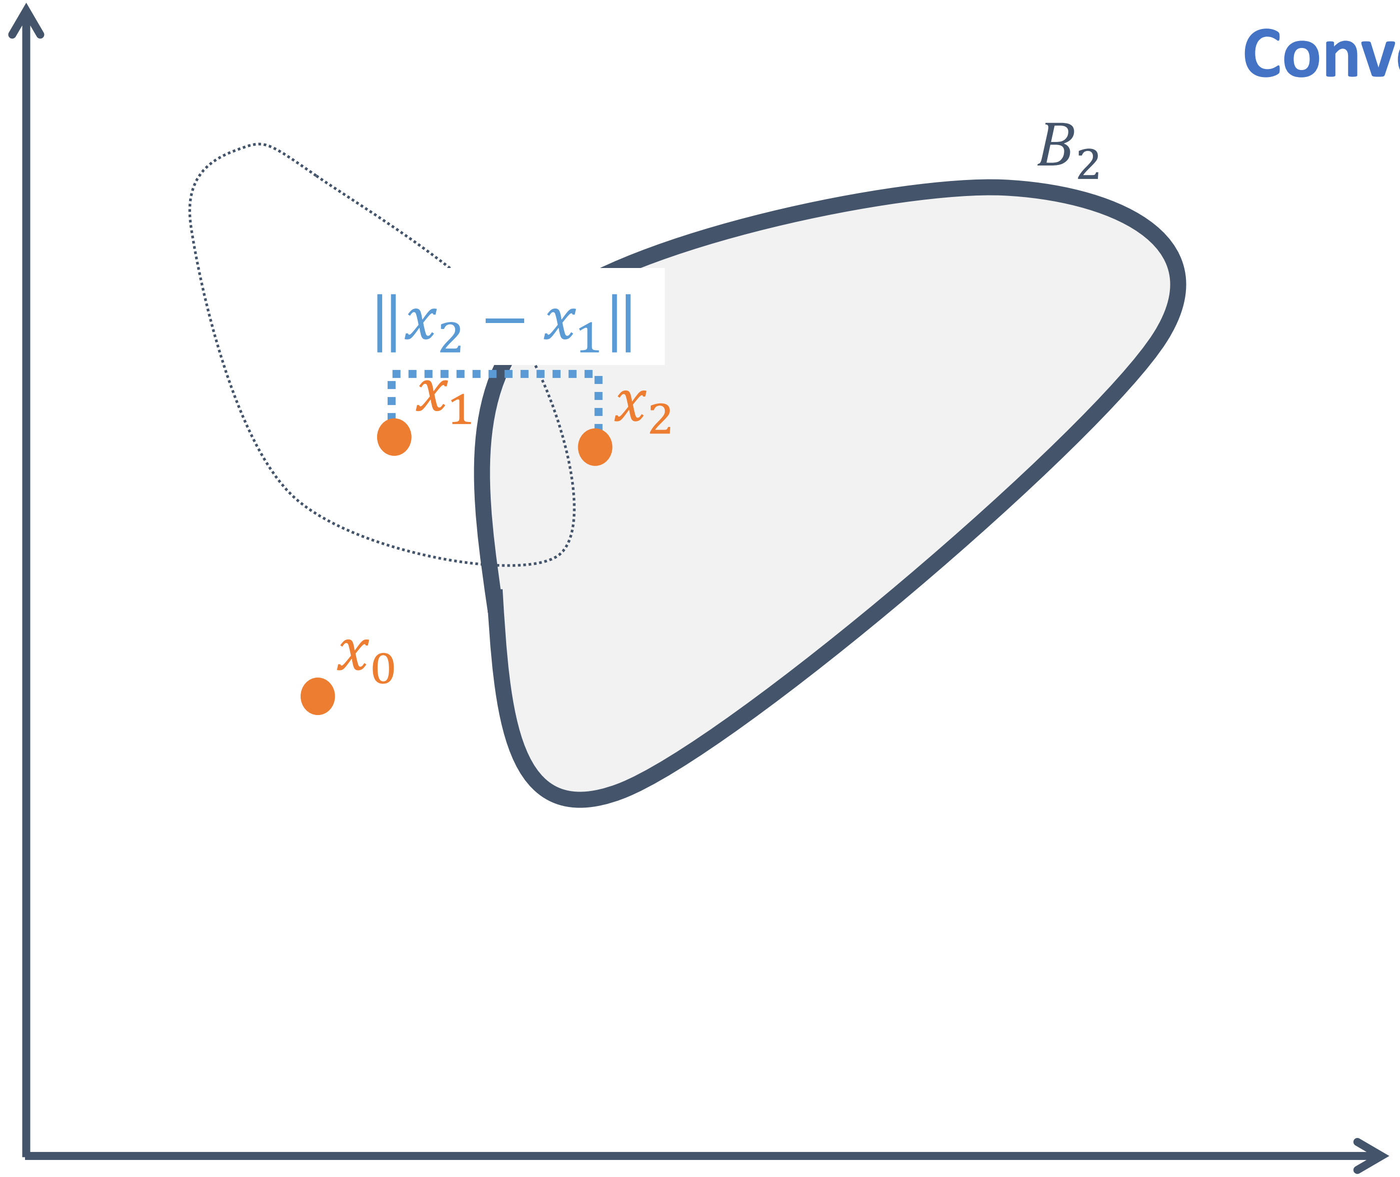

Material from Adam Wierman

Connection to Convex Body Chasing

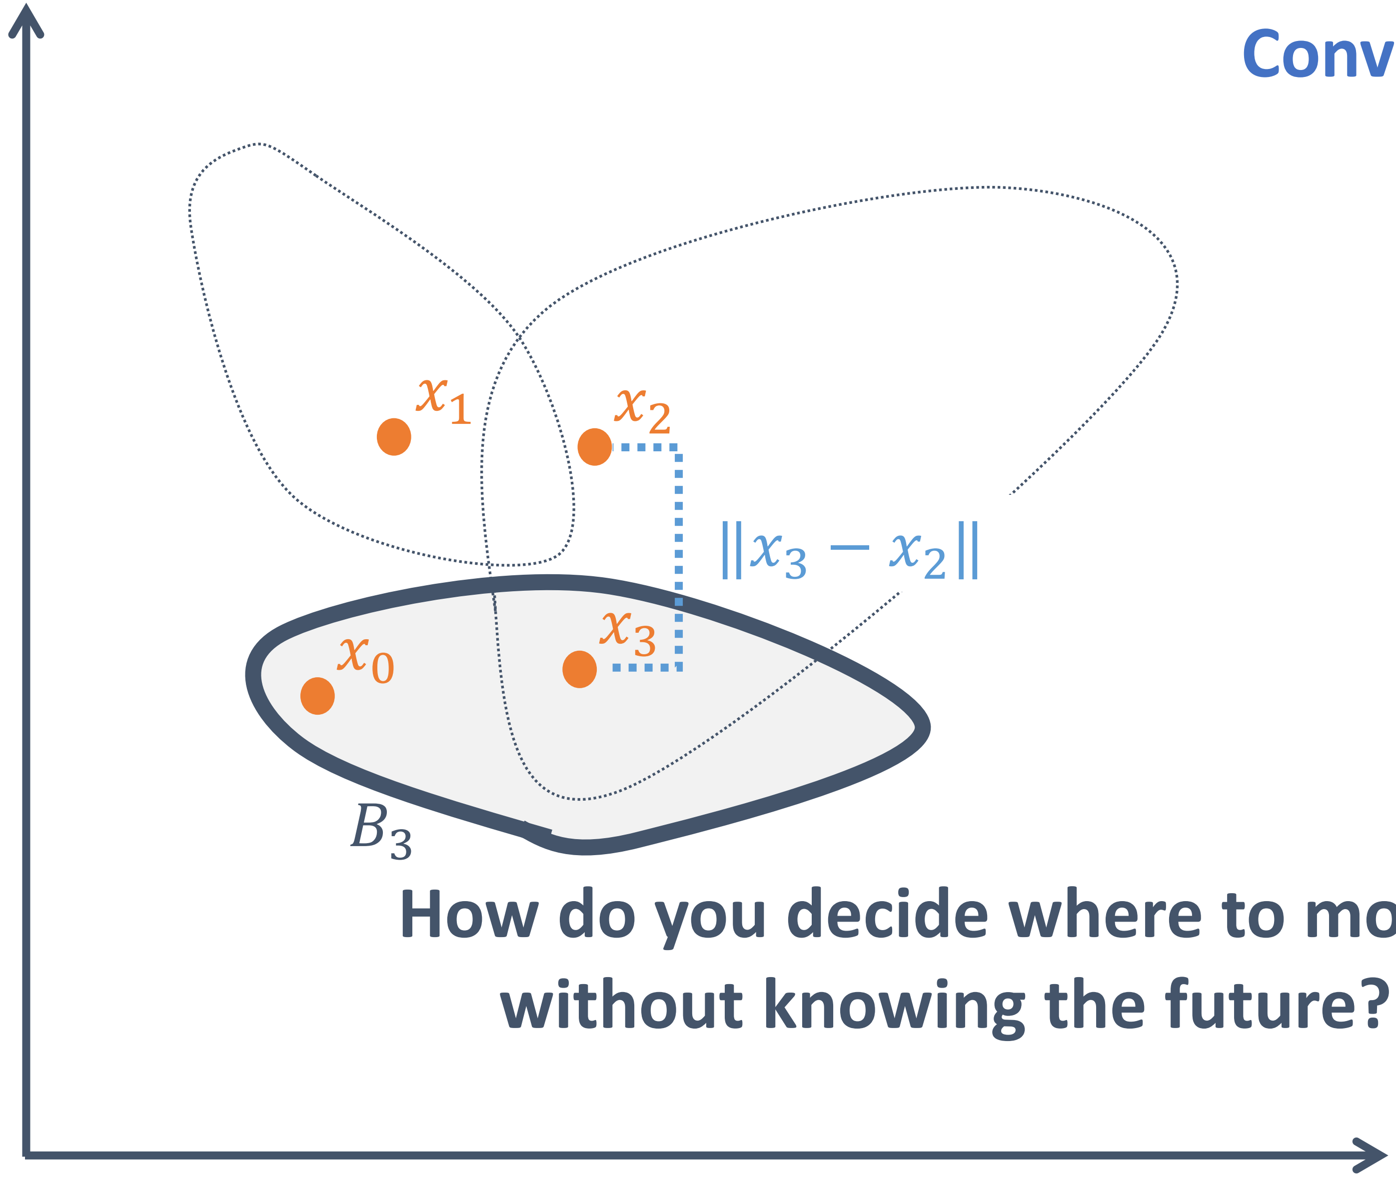

How do you decide where to move without knowing the future?

Material from Adam Wierman

Connection to Convex Body Chasing

Clearly, A SOCO algorithm can solve CBC…

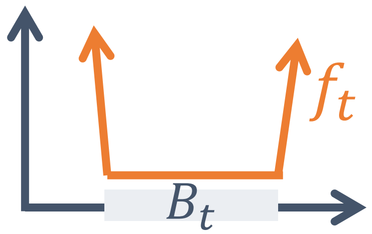

…a CBC algorithm can also solve a SOCO instance!

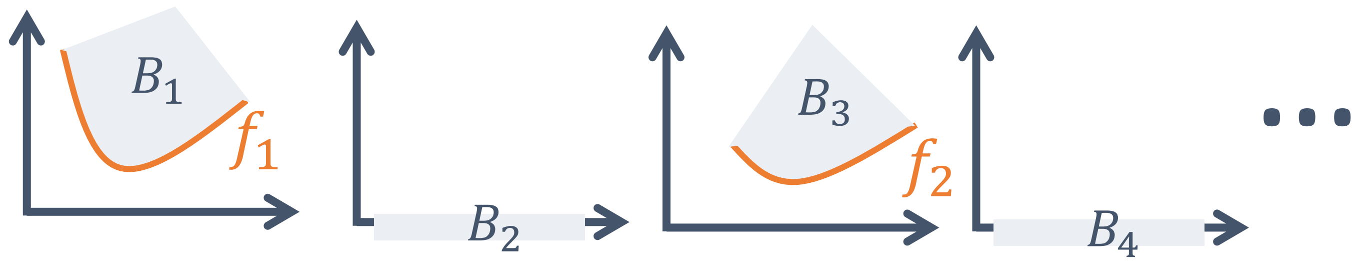

https://arxiv.org/abs/1811.00999 [Bubeck et al, SODA 2018]

Material from Adam Wierman

Online Control

- (Discrete Time) Dynamical System: $x_{t+1} = g(x_t, u_t, w_t)$

Linear Time Invariant: $g(x_t, u_t, w_t) = Ax_t + Bu_t + w_t$

- Additive Control Objective: $\sum_t Cost_t(x_t, u_t)$

Quadratic: $Cost_t(x_t, u_t) = x_t^T Q x_t + u_t^T R u_t$

- LQR: LTI System w/ Additive Quadratic Control Cost

Online Control (cont.)

- Restrict to LQR-like settings:

$$\min \sum_t x_t^T Q x_t + u_t^T R u_t$$
$$\text{s.t. } x_{t+1} = A x_t + B u_t + w_t$$

- Problem Settings:
  - Do we know A & B a priori or must we learn?
  - What assumptions can we make on A & B?

  } Known A & B, stabilizable

  - Is x_t fully or partially observed?

  } Fully observed

  - Is w_t observed before or after committing to u_t?
  - What assumptions can we make on w_t?

  } $w_{t:t+\ell-1}$ predictable & bounded

  - How do we measure performance?

  } Competitive Ratio

Recall: Tracking Control under Predictable Disturbances

Competitive Control (for LQR-like problems)

State **Control**

$$\min \sum_t x_t^T Q x_t + u_t^T R u_t$$

Disturbance

$$\text{s.t. } x_{t+1} = A x_t + B u_t + w_t$$

- At each time t:
  - Observe $x_t, \widehat{w}_t, \ldots, \widehat{w}_{t+\ell-1}$
  - Choose u_t
  - Repeat

(imperfect) knowledge of $w_t, \ldots, w_{t+\ell-1}$
"Fixed-Horizon Control" or "Model-Predictive Control"
Focus on 1-step prediction: $\ell = 1$

- Minimize Competitive Ratio: $sup_{\{x_0, w_1, \ldots, w_T\}} \dfrac{cost(ALG)}{cost(OPT)}$

Comments on Best Linear Controller & Static Regret

State Control

$$\min \sum_t x_t^T Q x_t + u_t^T R u_t$$

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

Disturbance

- Common to measure regret v.s. best static linear controller:
  - $\boldsymbol{u_t = K(x_t^* - x_t)}$
- However, the best static linear controller may have arbitrarily large competitive ratio vs offline optimal.

https://arxiv.org/abs/2003.00189 [Foster & Simchowitz]
https://arxiv.org/abs/2002.02574 [Goel & Hassibi]

Reduction to Online Convex Optimization w/ Structured Memory

(extends reduction in https://arxiv.org/abs/1810.10132 [Goel & Wierman 2019])

- Special case of LQR: Input-Disturbed Squared Regulator (IDSR):
  - (finalizing results on general LQR setting)

$$\min \sum_t \frac{q_t}{2} \|x_t\|^2 + \frac{1}{2} \|u_t\|^2$$

$$\text{s.t. } x_{t+1} = Ax_t + B(u_t + w_t)$$

- Many robotic systems are well described by IDSRs
  - E.g., https://arxiv.org/abs/1811.08027

Yiheng Lin

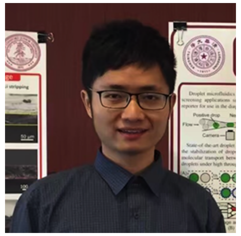

Guanya Shi

Online Optimization with Memory and Competitive Control, Guanya Shi, Yiheng Lin, et al., NeurIPS 2020

Reduction to Online Convex Optimization w/ Structured Memory

(extends reduction in https://arxiv.org/abs/1810.10132 [Goel & Wierman 2019])

- OCO w/ Structured Memory

New OCO setting!

$$\min \sum_t \frac{q_t}{2}\|x_t\|^2 + \frac{1}{2}\|u_t\|^2$$

s.t. $x_{t+1} = Ax_t + B(u_t + w_t)$

$$\min \sum_t f_t(y_t) + \underbrace{\frac{1}{2}\left\| y_t - \sum_{i=1}^{p} C_i y_{t-i} \right\|^2}_{c_t(y_t,\, y_{t-p:t-1})}$$

- y_t is a transformed representation of state x_t
    - p depends on dynamics (single, double integrator, etc.)
- Choose next state indirectly via control action u_t
- Knowing w_t defines OCO hitting & switching costs

Yiheng Lin

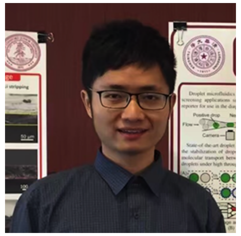

Guanya Shi

Online Optimization with Memory and Competitive Control, Guanya Shi, Yiheng Lin, et al., NeurIPS 2020

Roadmap to Optimistic ROBD

- Online Balanced Descent (OBD)

- Regularized Online Balanced Descent (ROBD)

- Optimistic ROBD

Online Balanced Descent (OBD)

https://arxiv.org/abs/1803.10366 [Chen, Goel, Wierman, 2018]

Choose level set to balance hitting and switching costs.

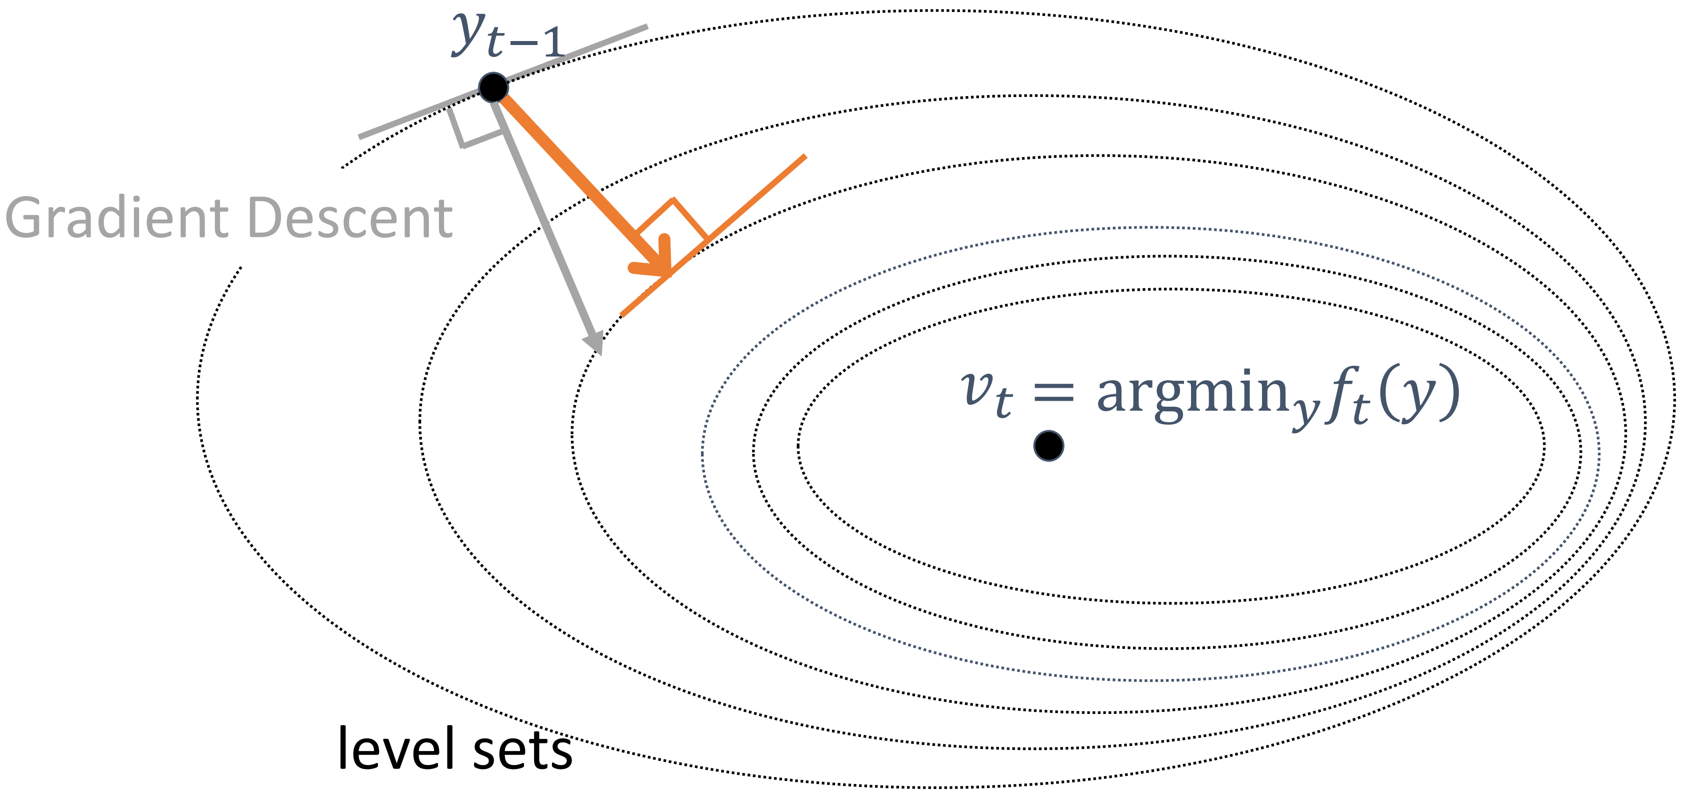

Material from Adam Wierman

Greedy/Regularized Online Balanced Descent

https://arxiv.org/abs/1905.12776 [Goel, Lin, Sun, Wierman, 2019]

Equivalent to regularization: $\arg\min_{y} f_t(y) + \lambda_1 c(y, y_{t-p:t-1}) + \frac{\lambda_2}{2}\|y - v_t\|^2$

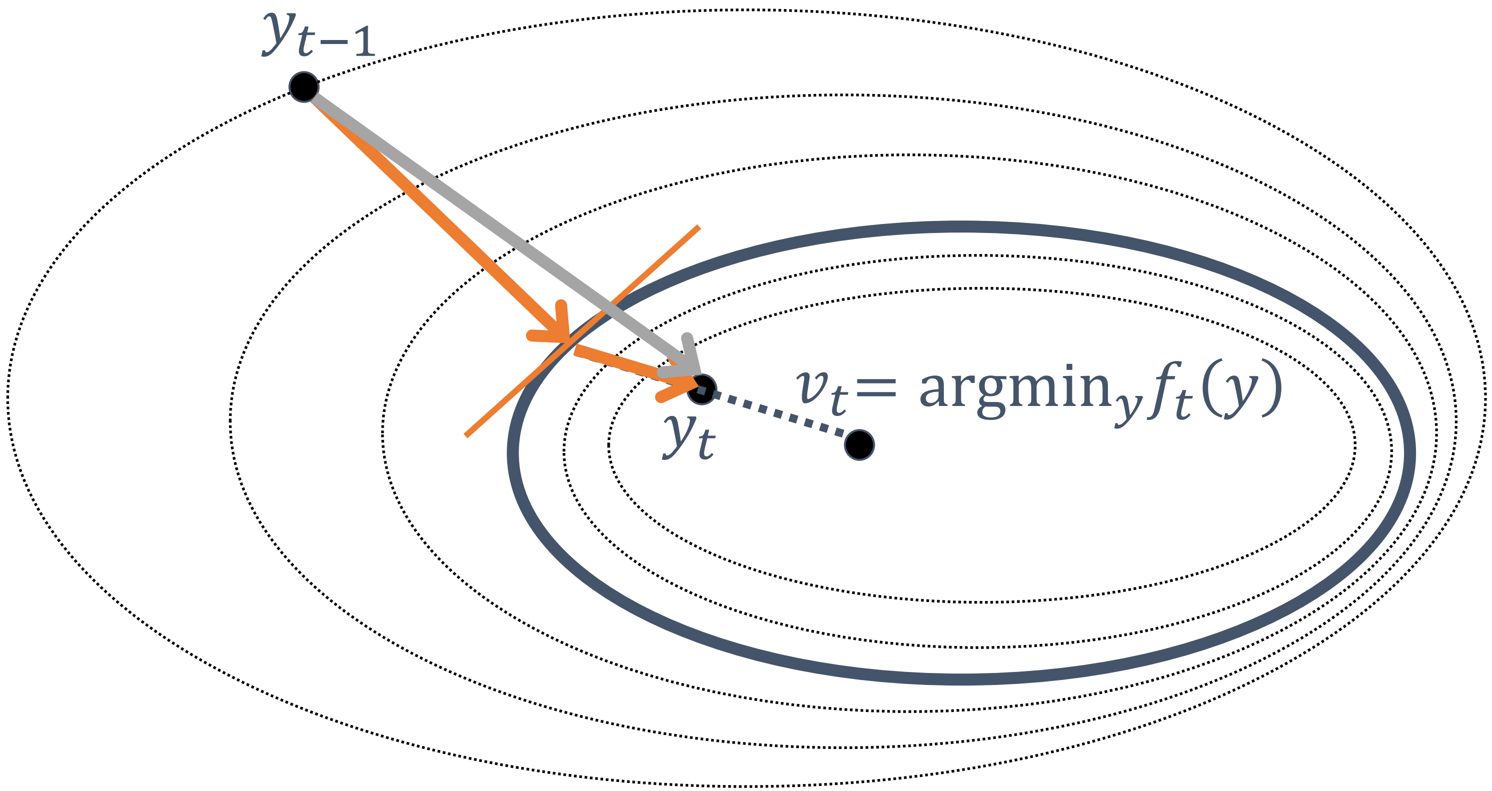

Material from Adam Wierman

Optimistic ROBD


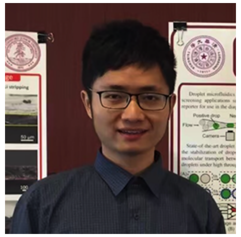

Yiheng Lin Guanya Shi

$w_t \in \Omega_t$ (w_t approximately known) $\rightarrow$ f_t approximately known

Choose w_t optimistically to minimize total cost

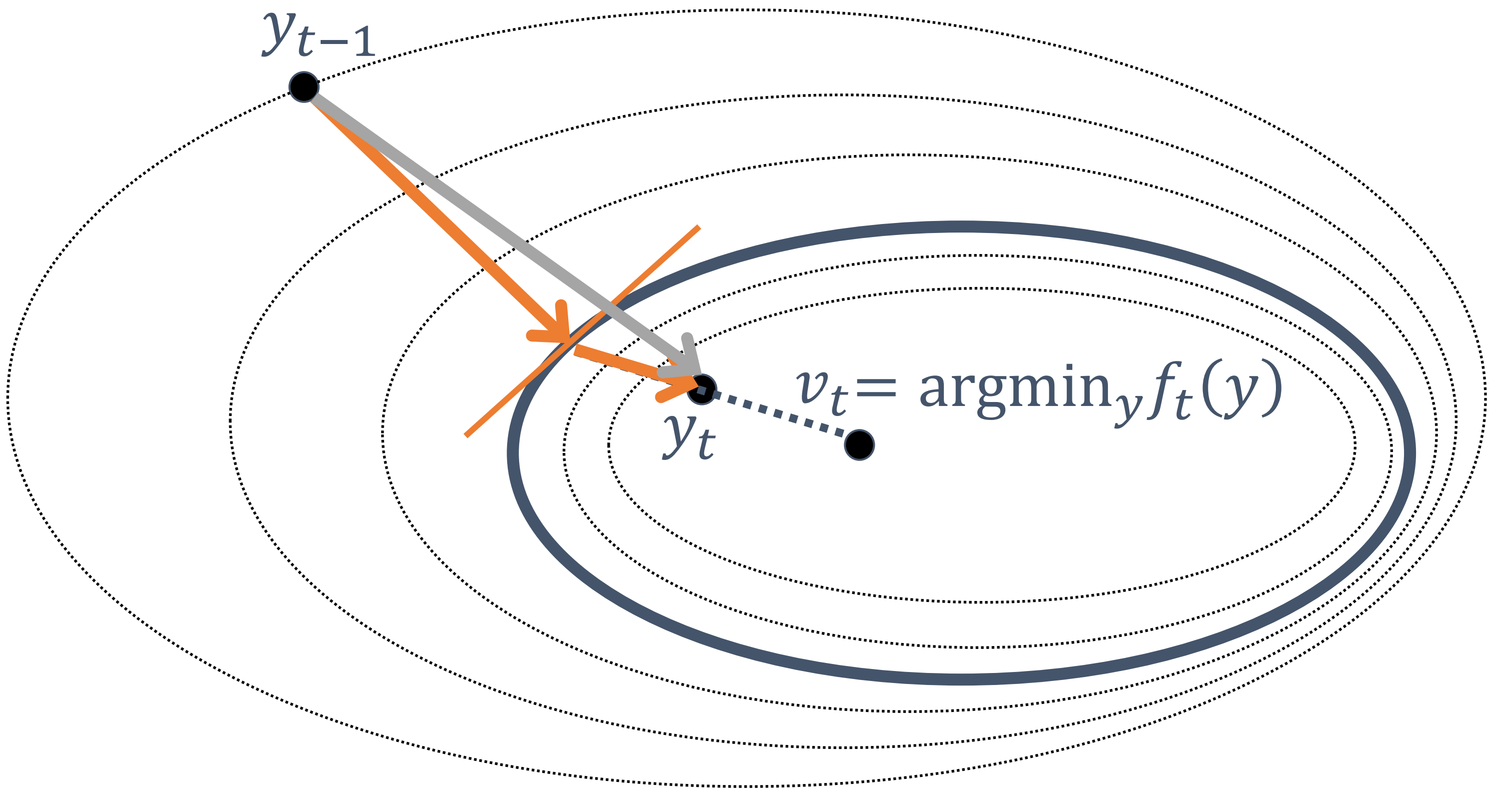

Online Optimization with Memory and Competitive Control, Guanya Shi, Yiheng Lin, et al., NeurIPS 2020

Material from Adam Wierman

Optimistic ROBD

Yiheng Lin

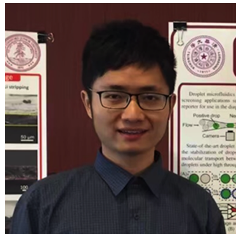

Guanya Shi

$w_t \in \Omega_t$ (w_t approximately known) $\rightarrow$ f_t approximately known

Choose w_t optimistically to minimize total cost

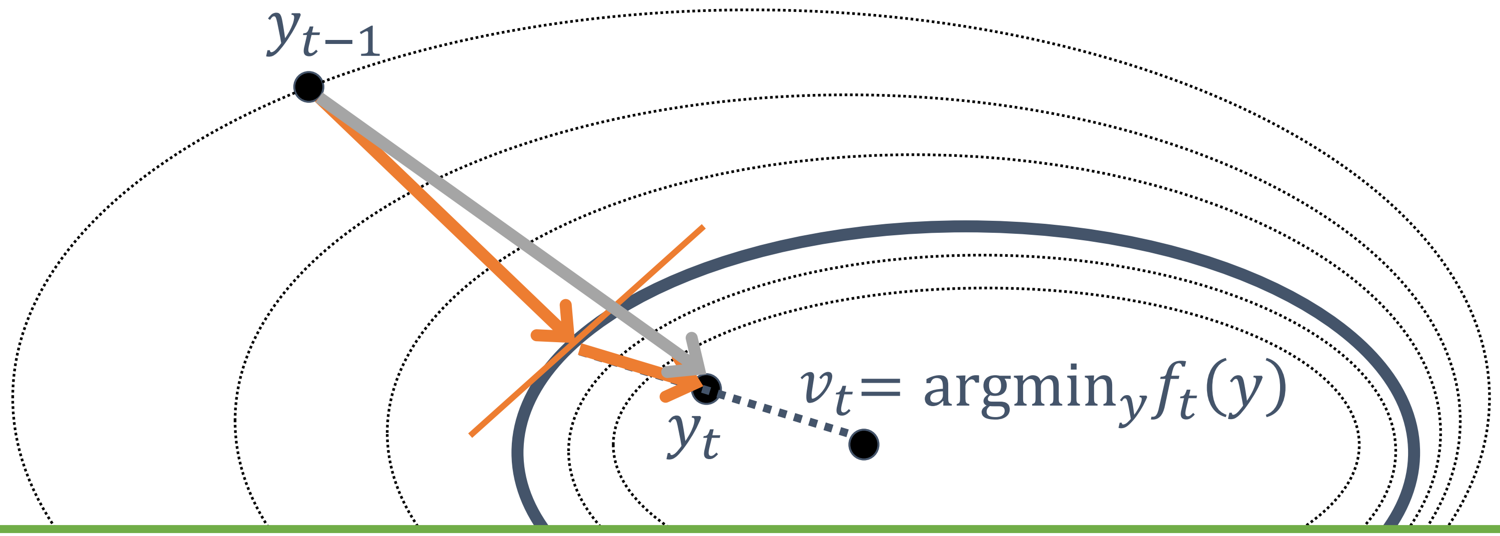

(Simplified) CR: $$O\left(\left(q_{max} + 4\alpha^2\right) \max\left\{\frac{1}{\lambda}, \frac{\lambda + q_{min}}{(1-\alpha^2)\lambda + q_{min}}\right\}\right)$$

Online Optimization with Memory and Competitive Control, Guanya Shi, Yiheng Lin, et al., NeurIPS 2020

Material from Adam Wierman

Summary

- Competitive ratio is a strong benchmark (practically relevant)
- New algorithms for competitive control
  - Online Optimization w/ Structured Memory
- Also working on forecasting & delay
  - **Competitive Control with Delayed Imperfect Information** https://arxiv.org/abs/2010.11637
  - **The Power of Predictions in Online Control** https://arxiv.org/abs/2006.07569
- Also working on nonlinear dynamical systems
  - **Online Learning of Nonlinear Control with Guaranteed Success: An Oracle-Based Robust Control-Powered Approach** *(preprint coming soon)*

Yiheng Lin

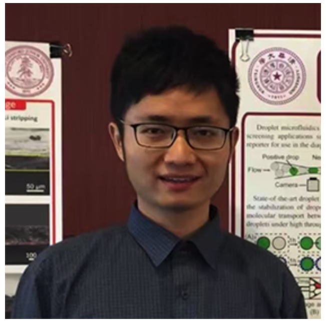

Guanya Shi

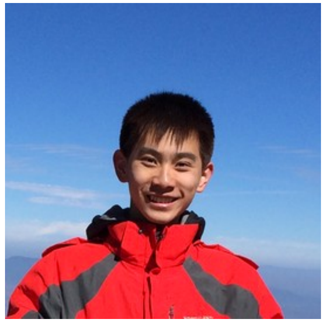

Chenkai Yu

Weici Pan

Dimitar Ho

Hoang Le

Adam Wierman

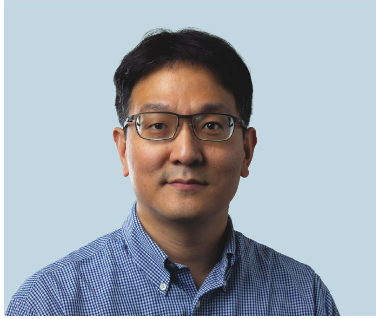

Soon-Jo Chung

John Doyle

Online Optimization with Memory and Competitive Control, Guanya Shi, Yiheng Lin, et al., NeurIPS 2020

On the Power of Predictions in Online Control, Chenkai Yu, Guanya Shi, et al., NeurIPS 2020

Competitive Control with Delayed Imperfect Information, Chenkai Yu, Guanya Shi, et al., arXiv

Online Learning of Nonlinear Control with Guaranteed Success: An Oracle-Based Robust Control-Powered Approach, Dimitar Ho, Hoang Le, et al., *(preprint available soon)*