Data center cooling using model-predictive control

NEVENA LAZIC, CRAIG BOUTILIER, TYLER LU, EEHERN WONG, BINZ ROY, MOONKYUNG RYU, GREG IMWALLE

Presenter: Judy Lin

Introduction

- Application of reinforcement learning to the task of large-scale commercial data center cooling

- Presentation Structure
  - Background
  - Solution Proposed
  - Experiments
  - Conclusion

Data Centers

- In 2014, data centers in the United States used:
  - 1.8% of power usage
  - 626 billion liters of water

- Data center cooling is well-suited for reinforcement learning
  - Complex, large-scale dynamical system
  - Non-trivial safety constraints
  - Potential for considerable improvements in energy efficiency

Data Center Cooling

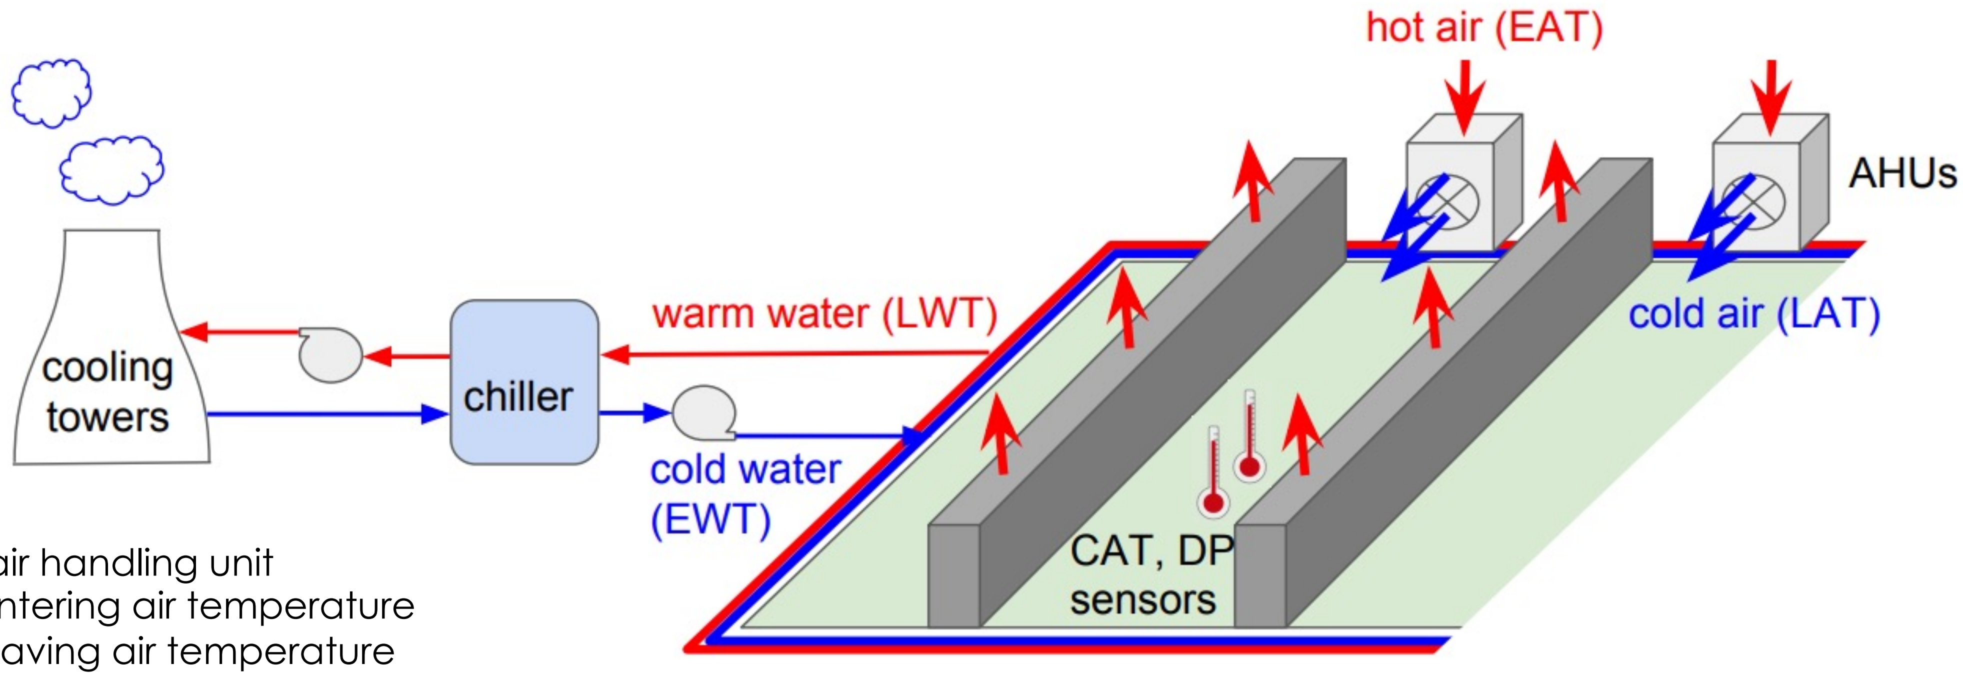

Figure 1

AHU = air handling unit
EAT = entering air temperature
LAT = leaving air temperature
CAT = cold-aisle temperature
DP = differential pressure
EWT = entering water temperature

Controls, States & Disturbances

Controls	• Fan speed • Valve opening
States	• Cold-aisle temperature • Differential pressure • Entering air temperature • Leaving air temperature
Disturbances	• Server power usage • Entering water temperature

Existing Solutions

- Simple and conservative
- Hand-tuned to specific requirements
  - Equipment architectures, layouts, configurations
- Local PID controllers
  - Located on each air handling unit
  - Regulate differential air pressure and leaving air temperature
  - Operates independently which can lead to a suboptimal state

Proposed Solution

- Model-Predictive Control (MPC)
  - System identification phase
    - Safe, random exploration with little, or no, prior knowledge required
  - Control phase
    - Optimizes the cost of a model-predicted trajectory
    - Executes optimized control action at first time step
    - Re-optimize at each time step

- Has been previously applied to regulate building cooling
  - Mostly used historical data or physics-based models
- (Zhou et al., 2012) used model-predictive control to regulate data center cooling
  - Controlled adaptive-vent floor tiles and air-conditioning units

Improvements over Existing Solution

- Directly regulates cold-aisle temperature as variable of interest
  - Models effect of each AHU on a neighbourhood (up to 5 rows)
  - Jointly optimize all controllers

- Simple to deploy
  - Few hours of exploration to identify system dynamics
  - Little or no prior knowledge required

- Operates at less conservative setpoints

MPC: Model Structure

- Linear autoregressive model with exogeneous variables
  - True dynamic are not linear → linear approximation is sufficient
- *T*-Markov model
  - Used cross validation to determine $T = 5$
- Each time step is 30s

$$\mathbf{x}[t] = \sum_{k=1}^{T} A_k \mathbf{x}[t-k] + \sum_{k=1}^{T} B_k \mathbf{u}[t-k] + C\mathbf{d}[t-1] \, . \tag{1}$$

MPC: Model Structure

- Sparsity in parameter matrices
  - Data center layout
  - Variable types
- Linear convolutional structure
  - Shared parameters due to regularity of physical layout

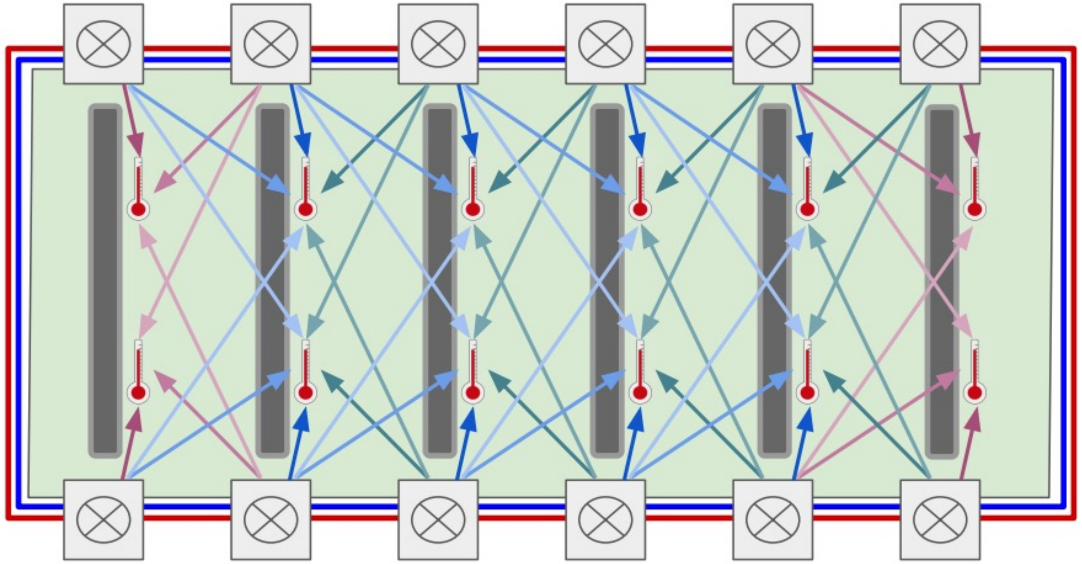

Figure 2

MPC: System Identification

- Initiated with "vacuous" model
- Exploration through simple, range-limited uniform random walk
  - Limits control variables to predetermined safe range
  - Limits maximum absolute change between consecutive time steps
- Model updated using recursive least squares

$$u_i^c[t+1] = \max(u_{\min}^c, \min(u_{\max}^c, u_i^c[t] + v_i^c)), \quad v_i^c \sim \mathrm{Uniform}(-\Delta^c, \Delta^c). \tag{2}$$

MPC: System Identification

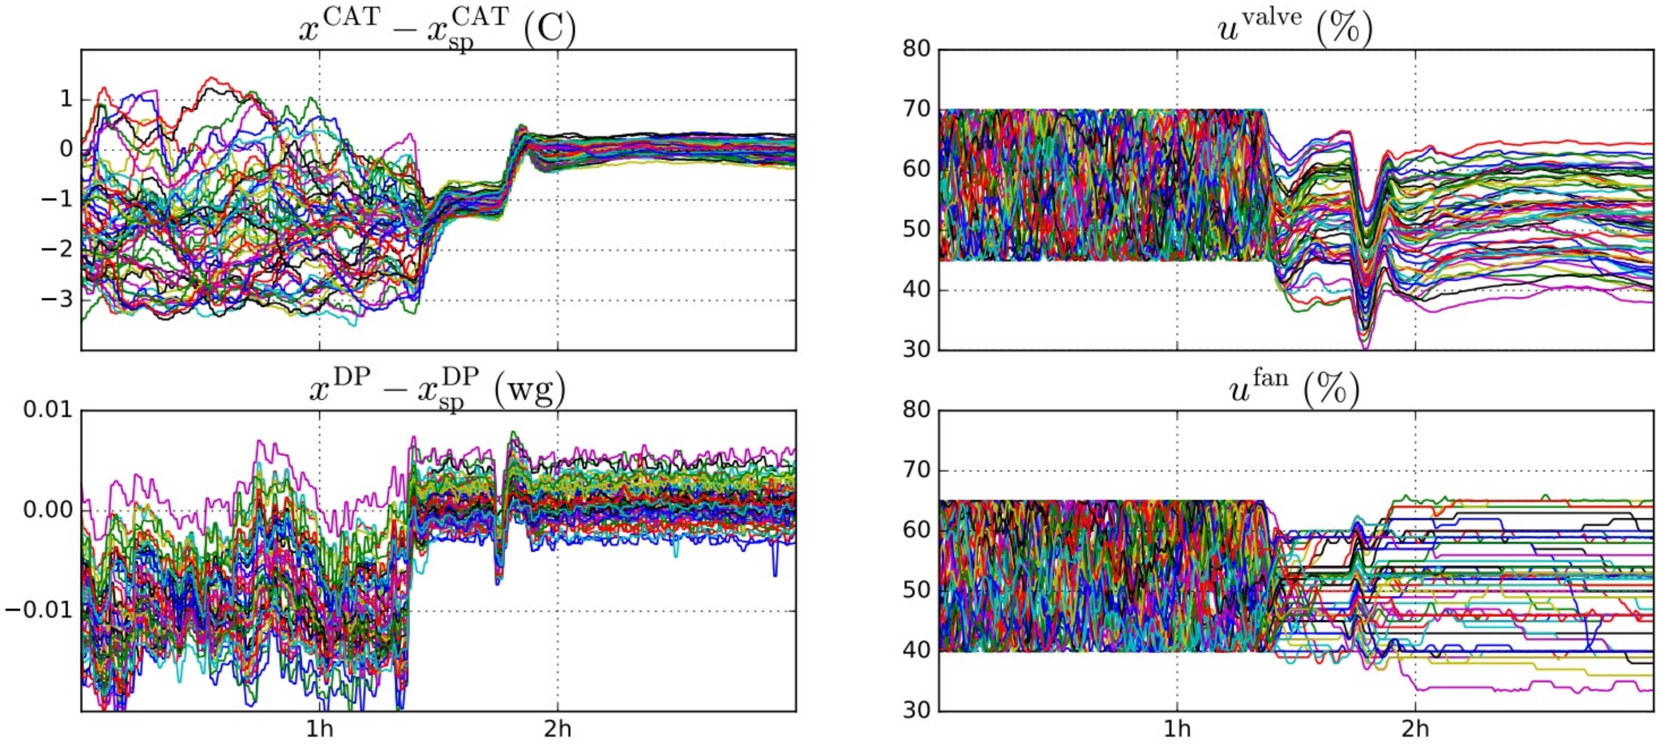

Figure 3

MPC: Control

- Optimize the cost of trajectory with length *L*
  - Only executes optimized control action at first time step
  - Assumes disturbances do not change over time

$$\min_{\mathbf{u}} \sum_{\tau=t}^{t+L} \sum_{i=1}^{M} \left(\sum_{s} q_s (x_i^s[\tau] - x_{\text{sp}}^s)^2 + \sum_{c} r_c (u_i^c[\tau] - u_{\min}^c)^2 \right) \tag{3}$$

$$\text{s.t.} \quad u_i^c \in [u_{\min}^c, u_{\max}^c], \quad |u_i^c[\tau] - u_i^c[\tau-1]| \leq \Delta^c, \quad \mathbf{d}[\tau] = \mathbf{d}[\tau-1] \tag{4}$$

$$\mathbf{x}[\tau] = \sum_{k=1}^{T} A_k \mathbf{x}[\tau-k] + \sum_{k=1}^{T} B_k \mathbf{u}[\tau-k] + C\mathbf{d}[\tau-1] \tag{5}$$

$$t \leq \tau \leq t+L, \; c \in \{\text{fan}, \text{valve}\}, \; \text{s} \in \{\text{DP}, \text{CAT}, \text{LAT}\}. \tag{6}$$

MPC: Control

- In practice, controls are optimized in TensorFlow with the Adam algorithm
  - Simple and fast → converges before 30s time step length

$$u_i^c[\tau] = \max(u_{\min}^c, \min(u_{\max}^c, u_i^c[\tau - 1] + \Delta^c \tanh(z_i^c[\tau]))) \quad (7)$$

- Continue to update model parameters in an online fashion
  - Estimate noise standard deviation as root mean squared error
  - Updates selectively to avoid overwhelming model with steady-state data

Experiments

- Experiments
  - System identification evaluation
  - Comparison to existing local PID controllers

- Challenge
  - Inability to control environmental disturbances
- Solution
  - Filtered to data within 0.25C and 0.004 of target temperature and pressure
  - Stratified data by different ranges of entering water temperature and server load

Experiments: System Identification

- **Model 1:** trained on 3 hours of deliberate exploration data with control following independent random walks limited to a safe range
  - Method in proposed solution

- **Model 2:** trained on a week of historical data generated by local PID controllers
  - Trained on 56 times more data than others

- **Model 3:** trained on 3 hours of data with controls recommend by a certainty-equivalent controller limited to a safe range
  - No exploratory actions

Experiments: System Identification

Figure 4

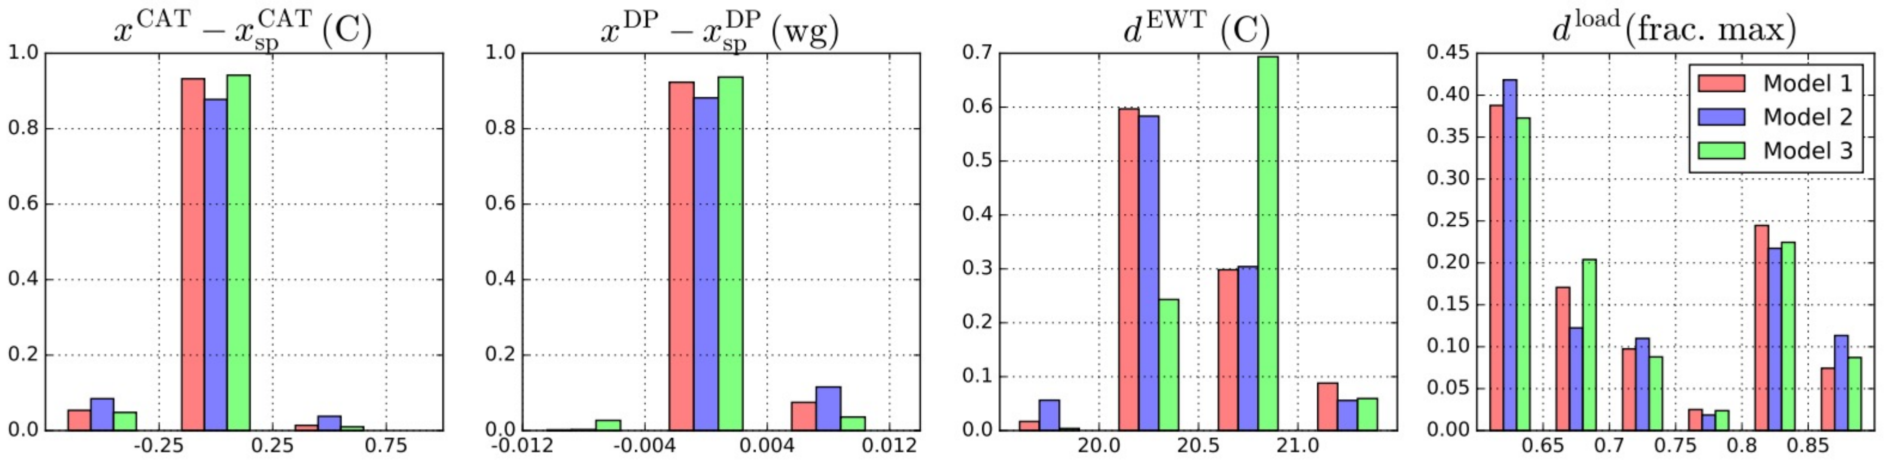

Table 2

Entering water temperature (C)	Server load (fraction of max)	Model 1 cost (% data)	Model 2 cost (% data)	Model 3 cost (% data)
≤ 20.5	≤ 0.7	**84.3** (31%)	94.4 (29.9%)	99.6 (13.7%)
> 20.5	≤ 0.7	**85.8** (17.6 %)	93.8 (14.1 %)	112.7 (36.0 %)
≤ 20.5	> 0.7	**142.4** (21.9 %)	149.4 (20.4 %)	178.2 (8.3 %)
> 20.5	> 0.7	**144.6** (15.3 %)	148.9(12.8 %)	182 .1 (29.9 %)
any	any	**110.2** (85.8%)	117.9 (77.2%)	140.4 (87.9%)

Experiments: Local PID Controllers

Figure 5

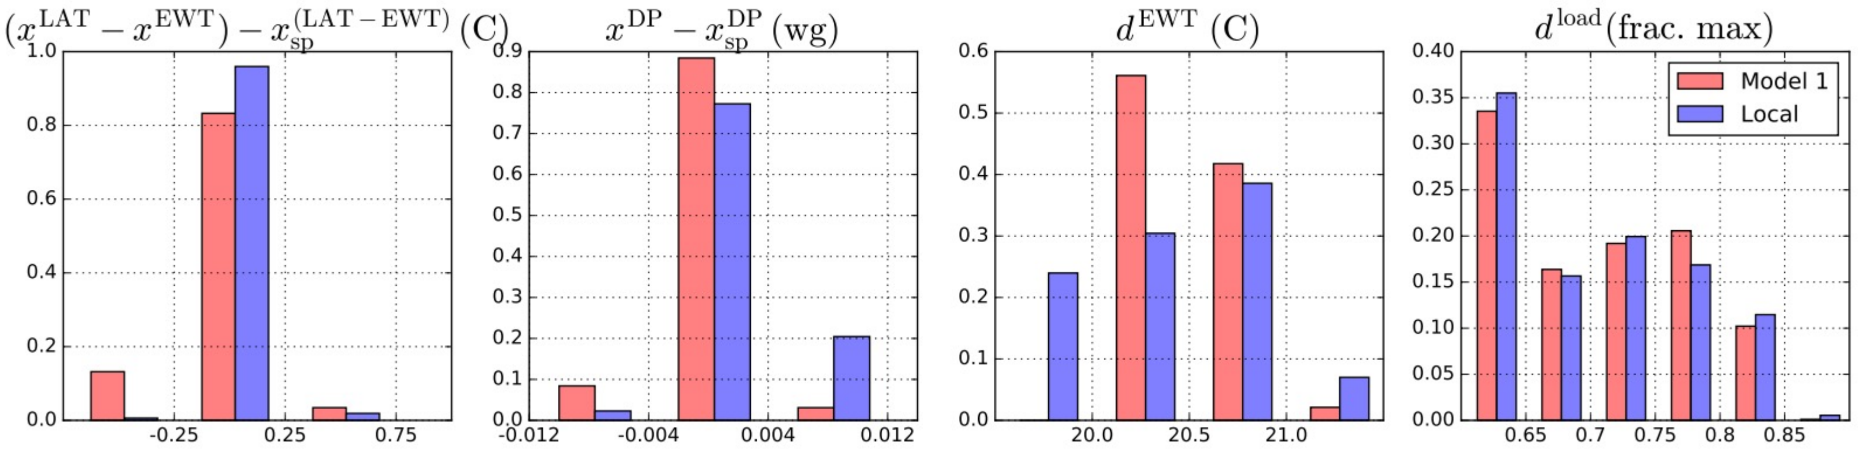

Table 3

Entering water temperature (C)	Server load (frac. max)	Local controllers cost (% data)	Model 1 cost (% data)
≤ 20.5	≤ 0.7	**95.3** (19.8 %)	106.4 (22.6 %)
> 20.5	≤ 0.7	107.9 (13.8 %)	**104.9** (15.0 %)
≤ 20.5	> 0.7	170.3 (20.1%)	**130.5** (18.8 %)
> 20.5	> 0.7	187.8 (20.4 %)	**128.7** (18.0 %)
any	any	142.2 (74.4%)	**116.7** (74.1%)

Conclusion

- Contributions
  - Simple linear model identified from only a few hours of exploration is sufficient for effective regulation of temperatures and airflows in a large-scale commercial data center
  - Proposed solution is more cost effective than use of local PID controllers and controllers based on non-exploratory data
- Possible extensions
  - Use of a higher-capacity model
  - Learn a mixture of linear models
  - Improve overall data center efficiency

References

- Lazic, Nevena, Craig Boutilier, Tyler Lu, Eehern Wong, Binz Roy, Moonkyung Ryu, and Greg Imwalle. "Data center cooling using model-predictive control." *Advances in Neural Information Processing Systems* 31, (2018): 3814-3823.

- Zhou, Rongliang, Cullen Bash, Zhikui Wang, Alan McReynolds, Thomas Christian, and Tahir Cader. "Data center cooling efficiency improvement through localized and optimized cooling resources delivery." *American Society of Mechanical Engineers*, (2012): 1789–1796.