Green Scheduling for Energy-Efficient Building Controls

Prof. Rahul Mangharam
Director, Real-Time & Embedded Systems Lab
Dept. Electrical & Systems Engineering
Dept. Computer & Information Science
University of Pennsylvania
rahulm@seas.upenn.edu

Penn Engineering

mLAB
Real-Time & Embedded

Green Scheduling for Energy-Efficient Building Controls

With Truong X. Nghiem, Madhur Behl and Willy Bernal

Published in: Green Computing’11, CDC’11, ACC’12,
CDC’12, BuildSys’12, RTSS’12, ACC’13, ICCPS’14

The Peak Power Minimization Problem

$$\text{supply} = \text{demand}$$

111 million viewers watch the SuperBowl

During a commercial, millions of refrigerators and microwaves trigger simultaneously, causing massive spikes in the energy demand

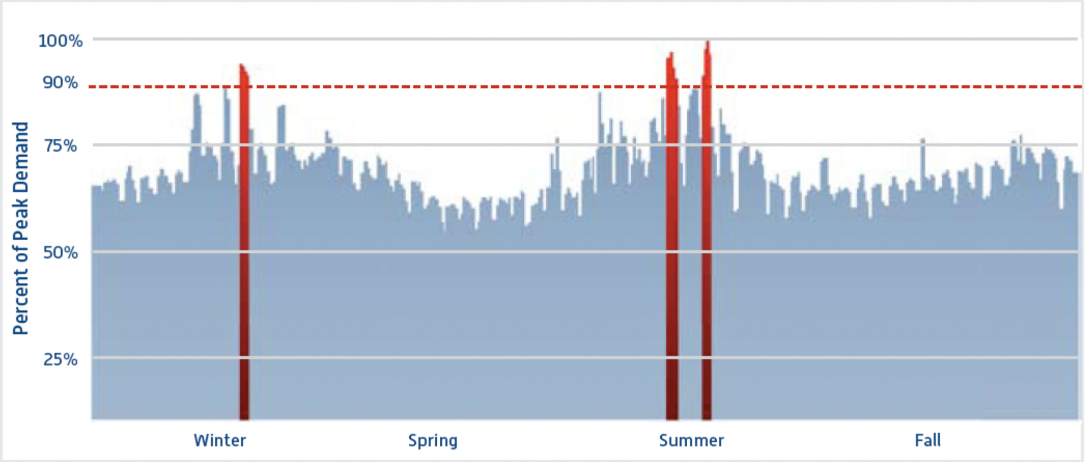

Human behavior and environmental conditions are responsible for high temporal correlation of energy demand.

Penn's Power Consumption

Estimated Annual Electricity Bill for UPenn is **$28,330,027.75**

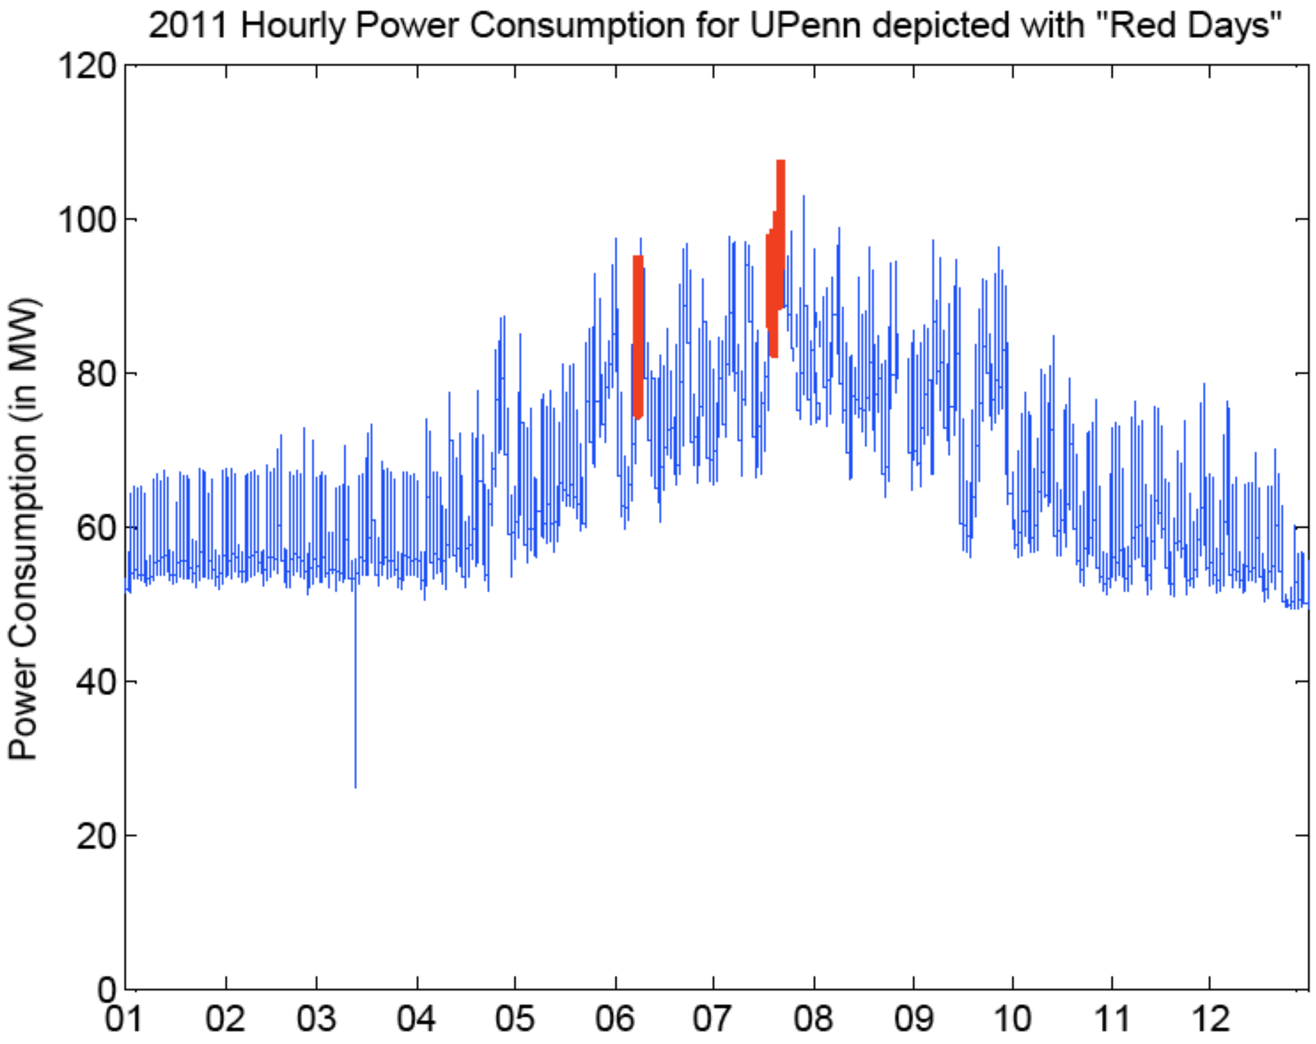

Red Days 5 days electric bill → $1.5 Million

Demand across the grid (and hence the price of electricity) is very high

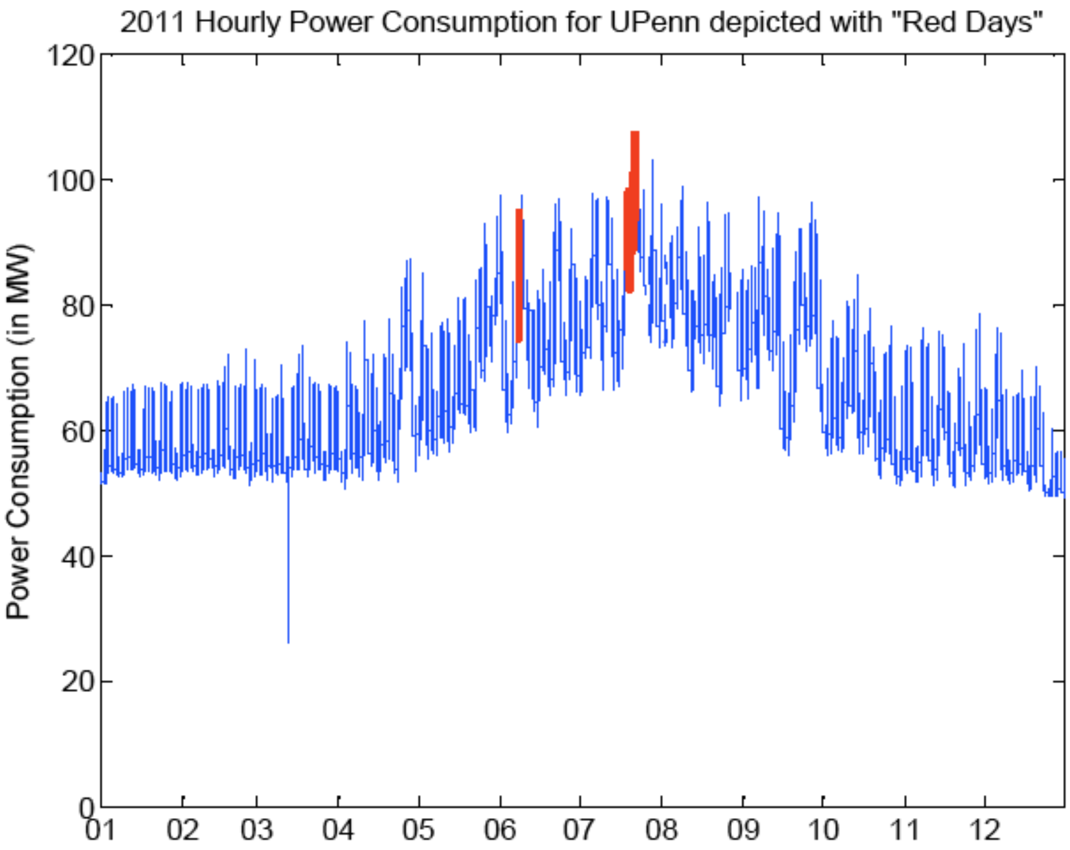

Red days alone cost
Penn **$1,470,228.05**
(**5.1%** of the total bill)

Penn's Electricity Demand in 2011

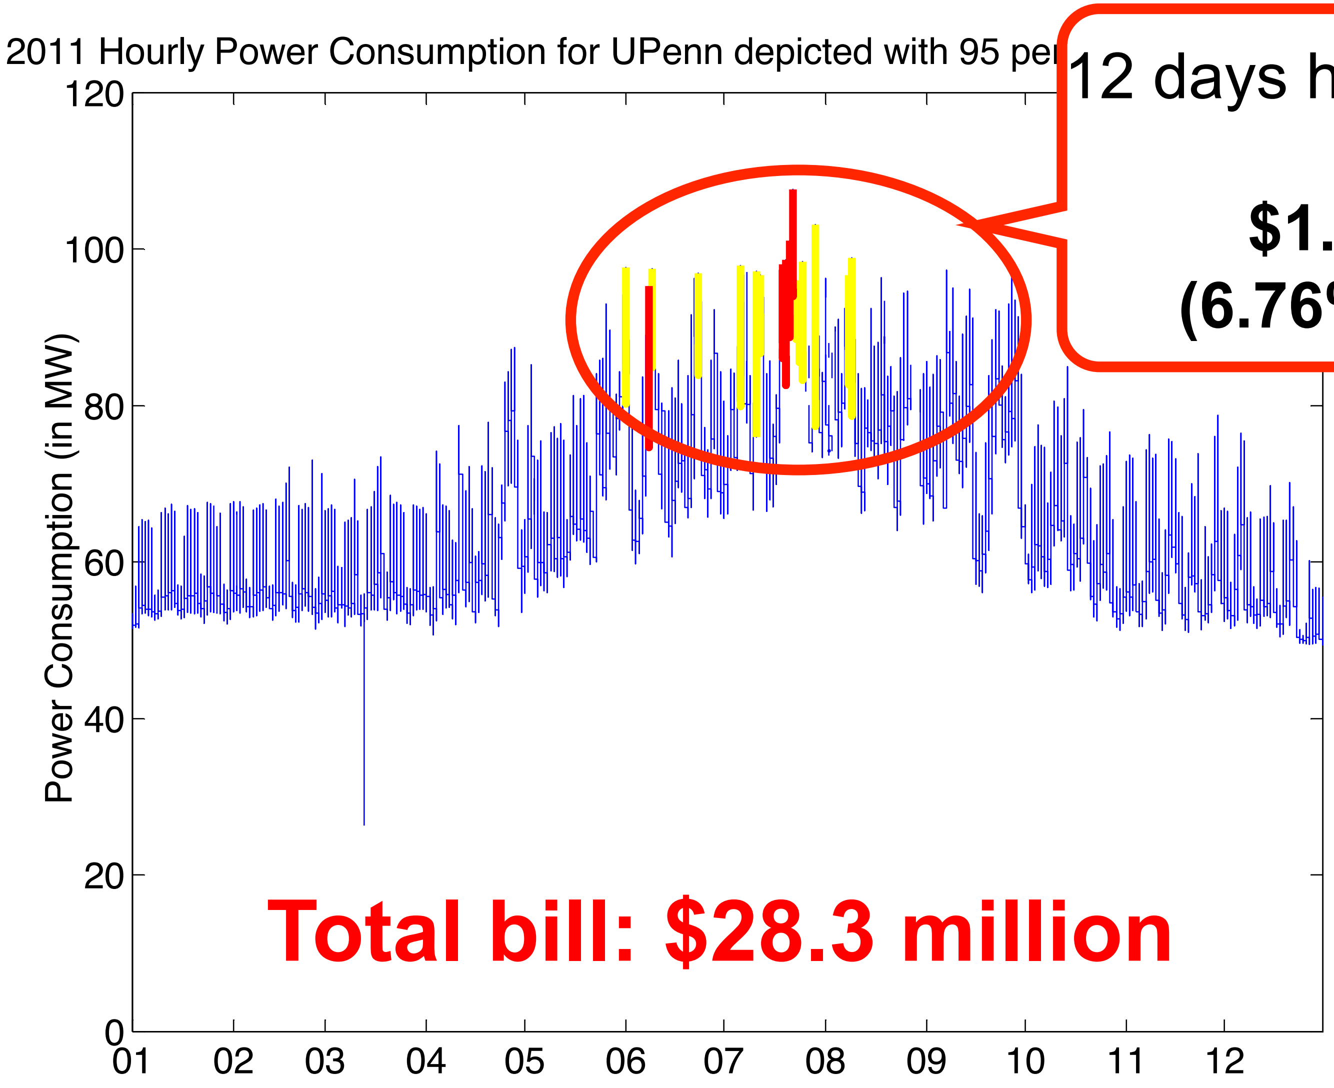

12 days highest demand cost
$1.9 million
(6.76% total bill)

Total bill: $28.3 million

Peak Electricity Demand is Expensive!

2011 Hourly Real Time Mkt prices based on Locational Marginal Pricing (LMP) for PECO by PJM

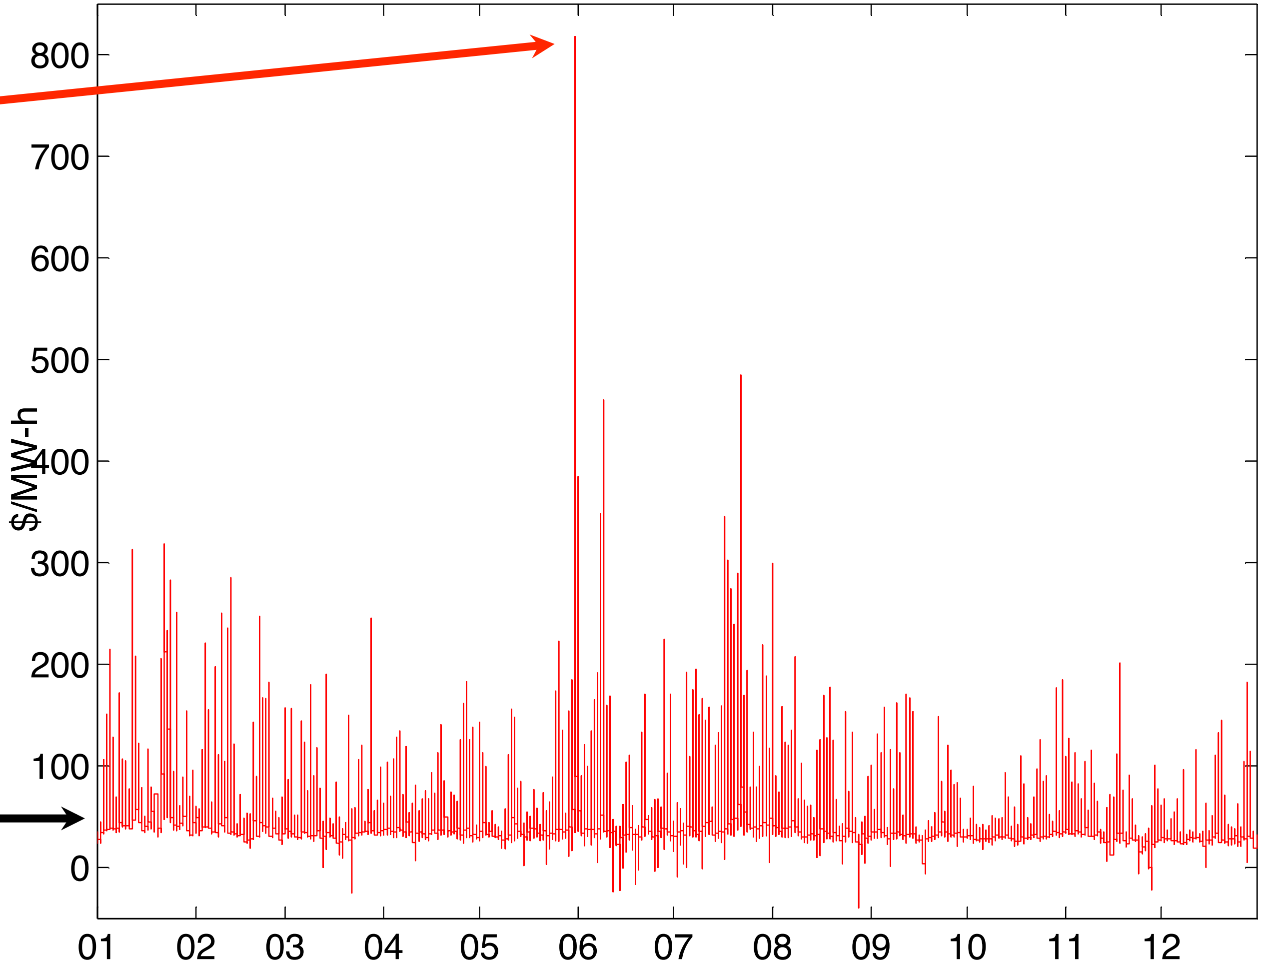

Peak $817/MWh

27 x

Normally $30/MWh

Peak Electricity Demand is Expensive!

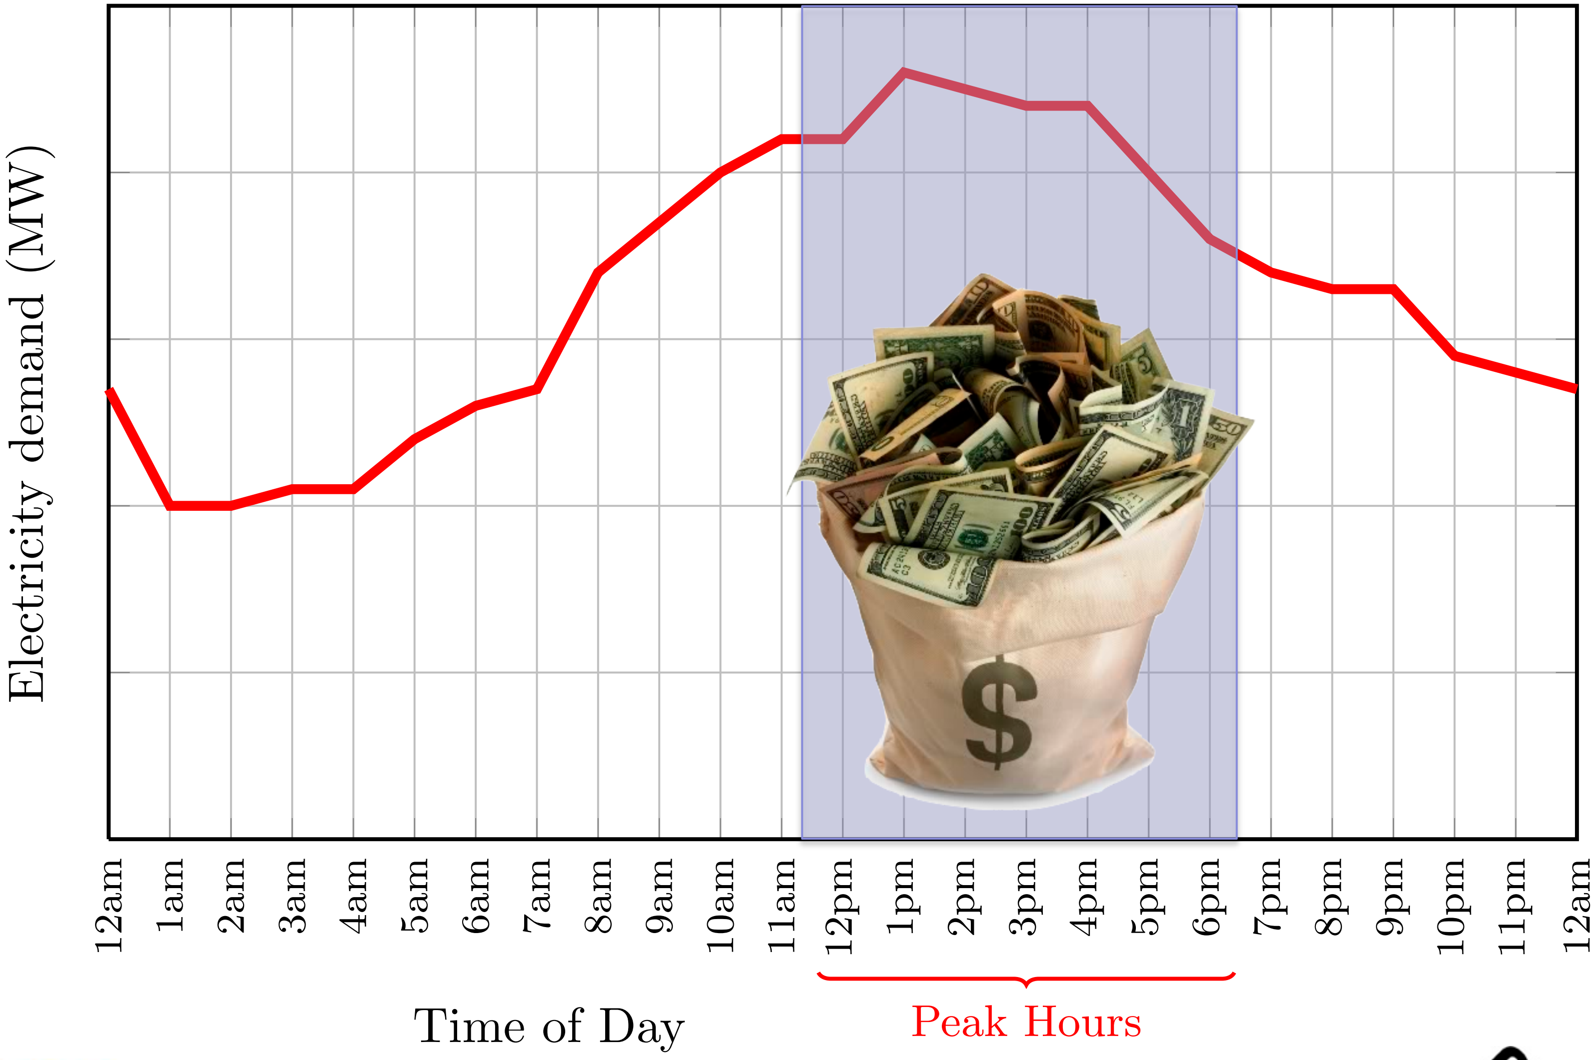

Peak Demand Reduction

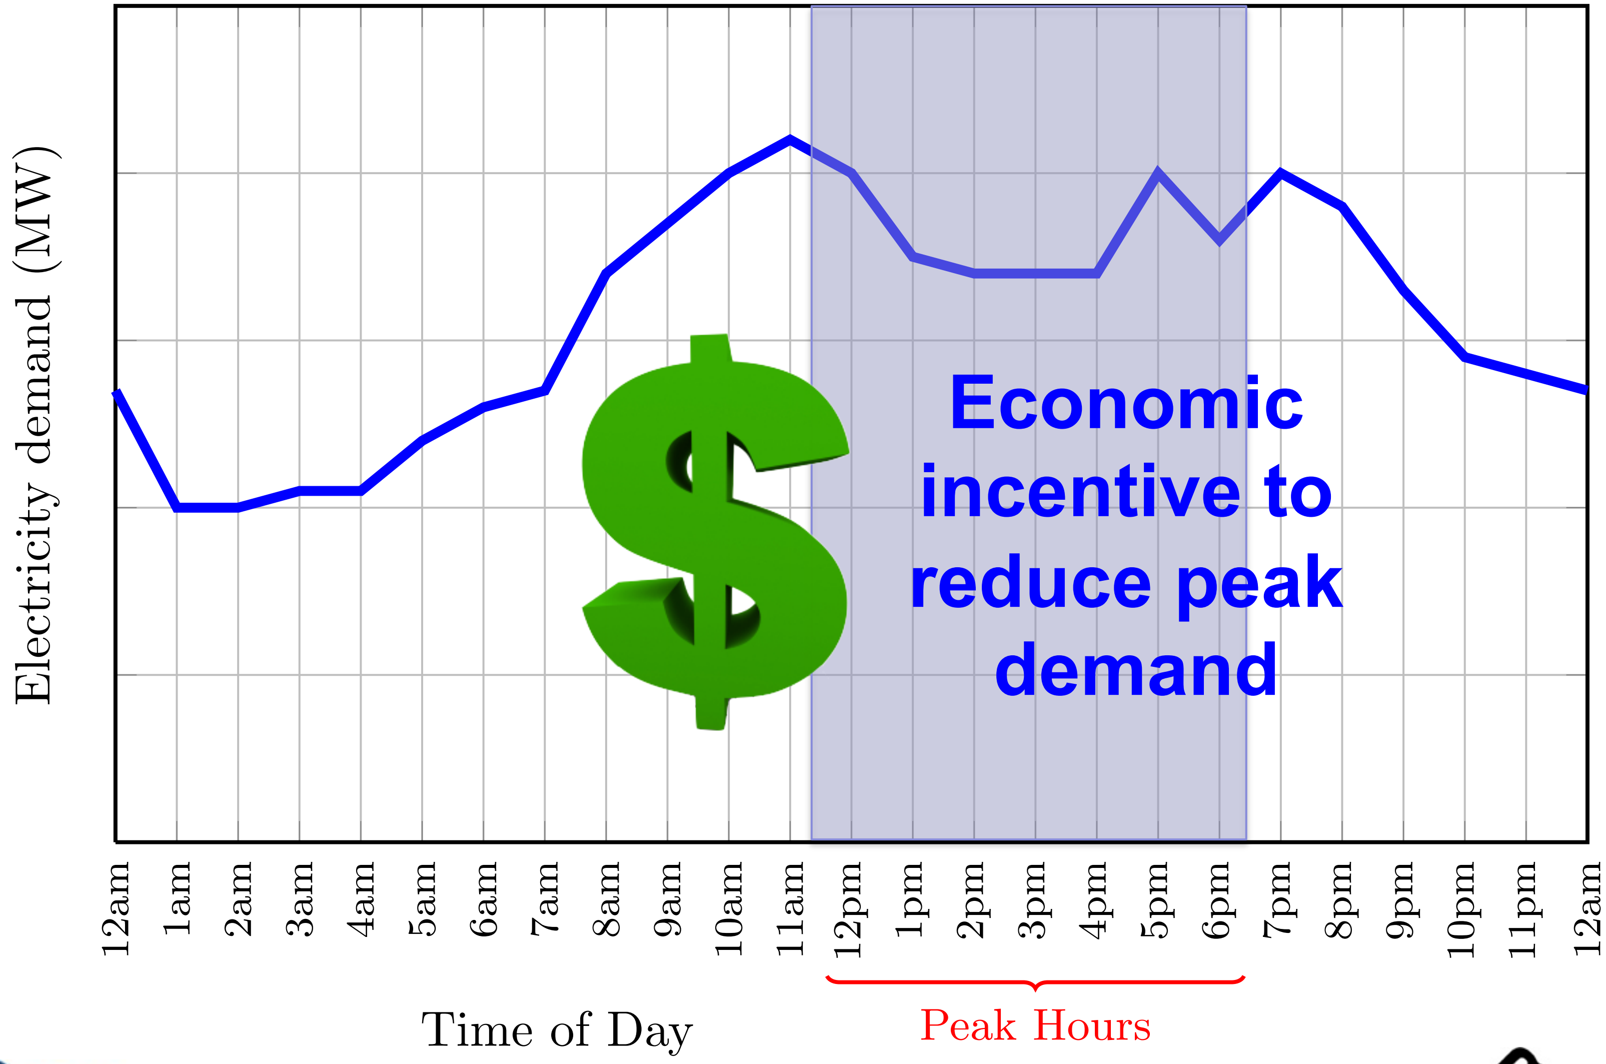

Economic incentive to reduce peak demand

How to Reduce Peak Demand?

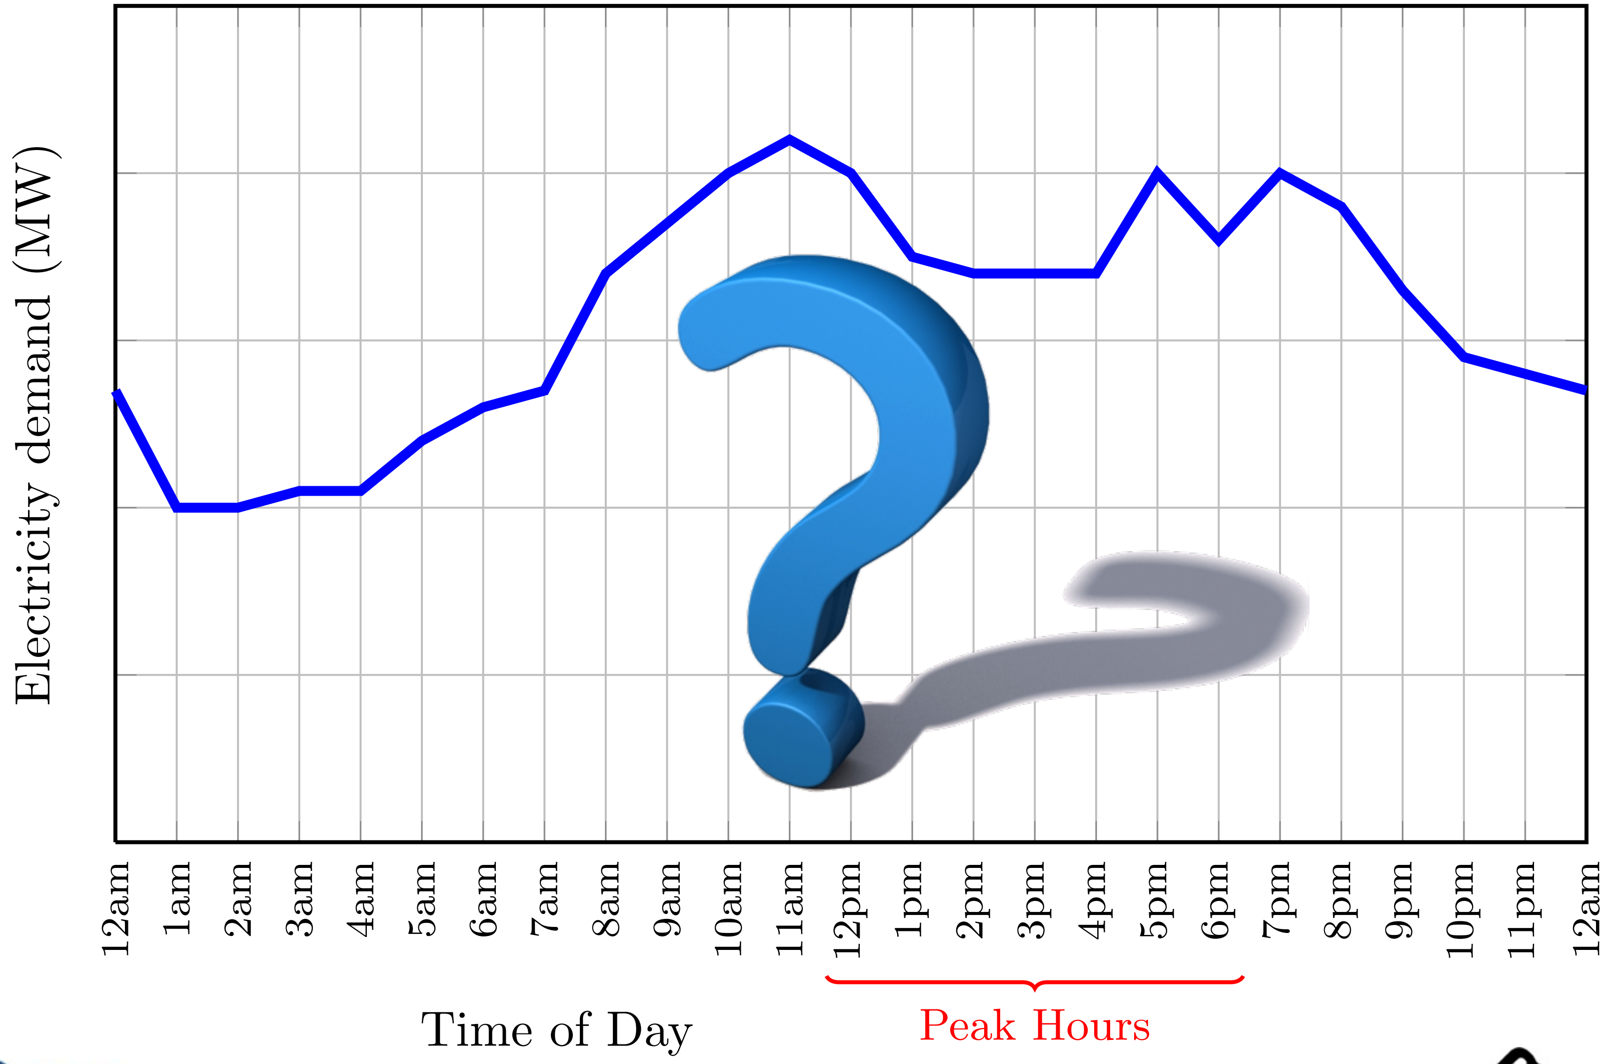

Peak Demand Reduction Approaches

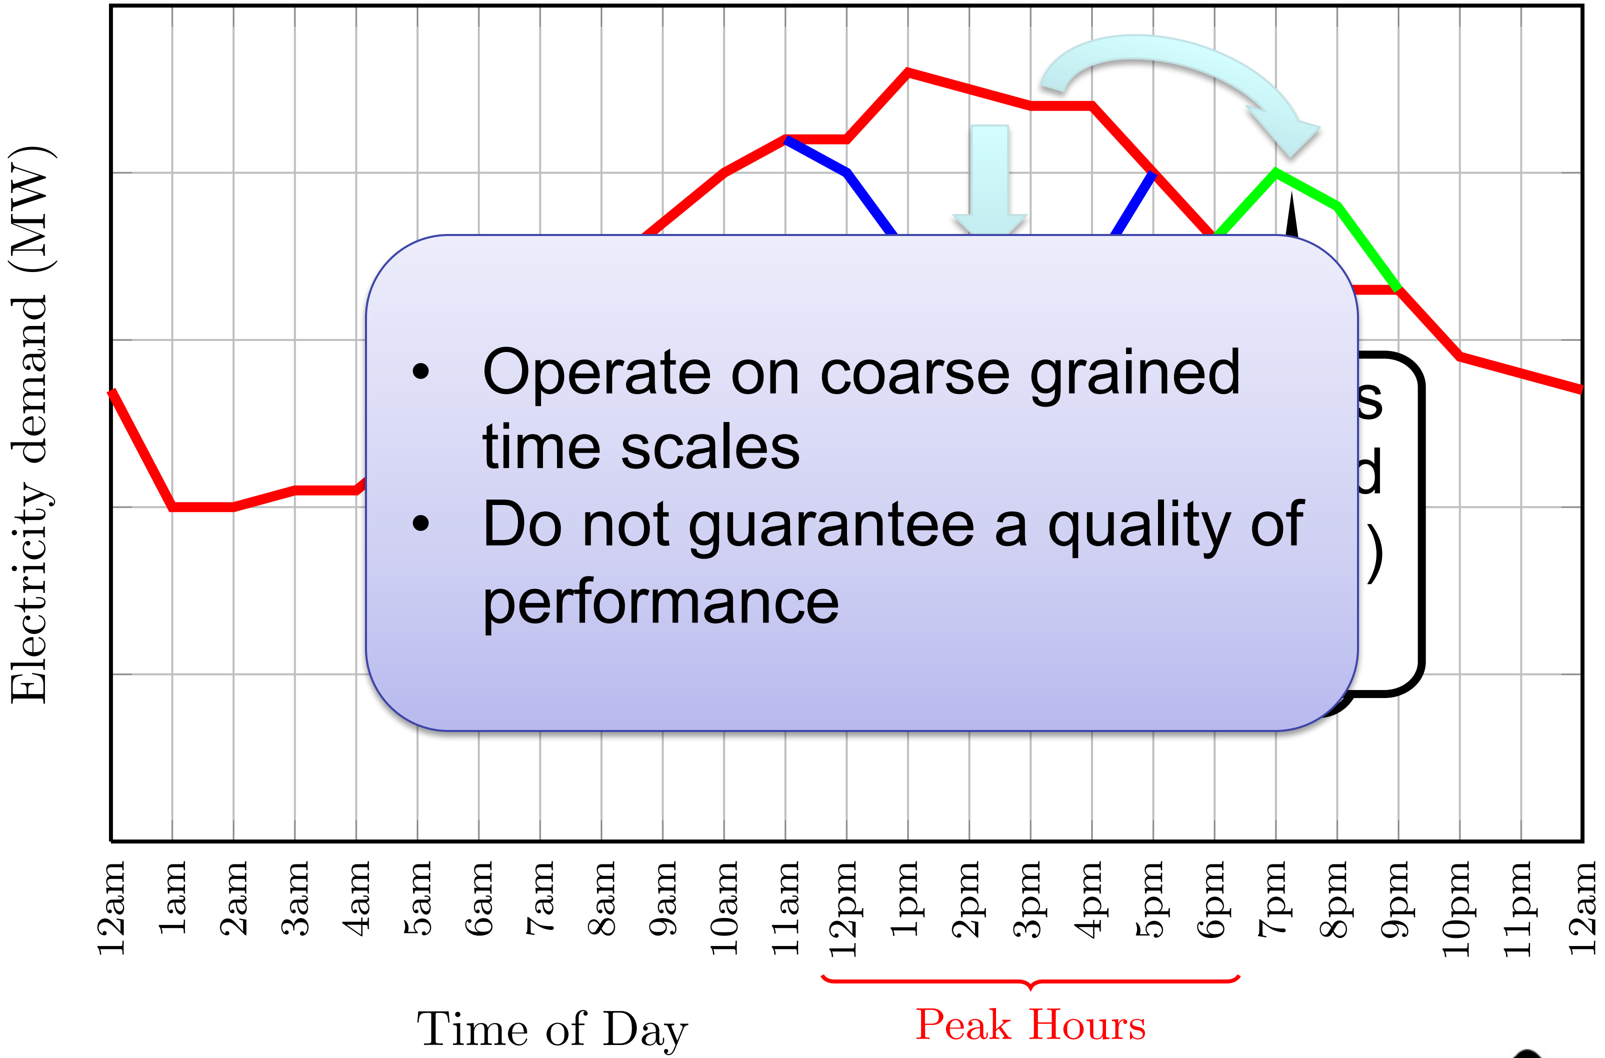

- Operate on coarse grained time scales
- Do not guarantee a quality of performance

Un-coordinated Control Systems

- HVAC (Heating,Ventilation and Air Conditioning) systems, chiller systems and lighting systems operate independently of each other
    - frequently result in temporally correlated energy demand surges (peaks)

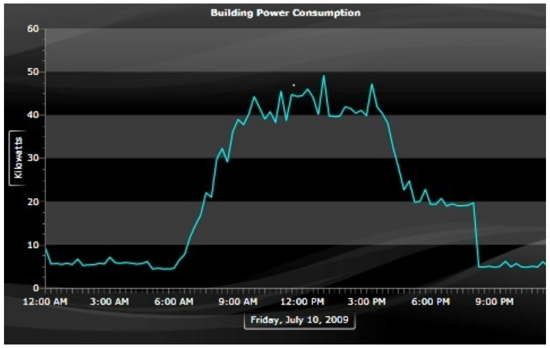

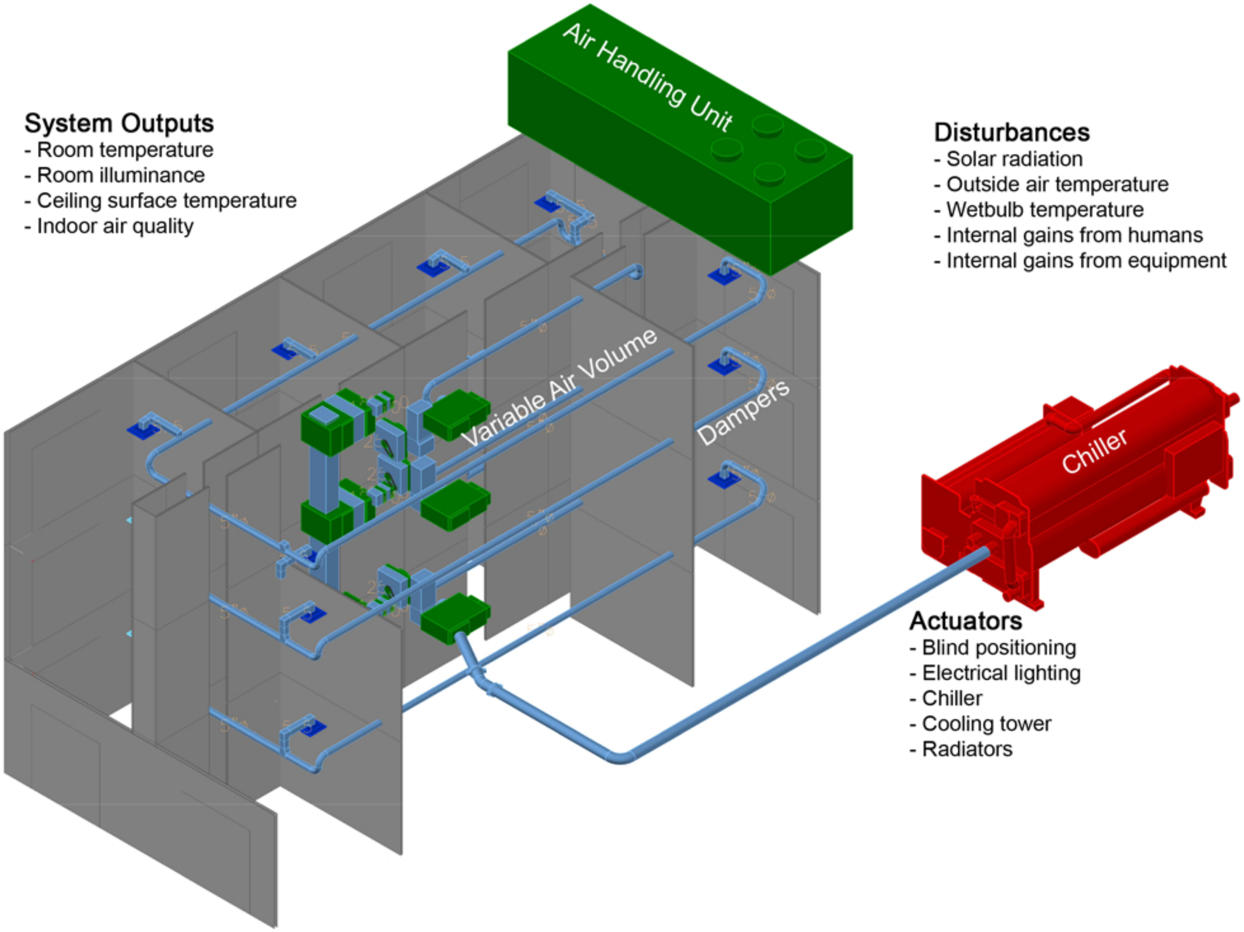

System Outputs
- Room temperature
- Room illuminance
- Ceiling surface temperature
- Indoor air quality

Disturbances
- Solar radiation
- Outside air temperature
- Wetbulb temperature
- Internal gains from humans
- Internal gains from equipment

Actuators
- Blind positioning
- Electrical lighting
- Chiller
- Cooling tower
- Radiators

GREEN SCHEDULING APPROACH

Green Scheduling (GS) Approach

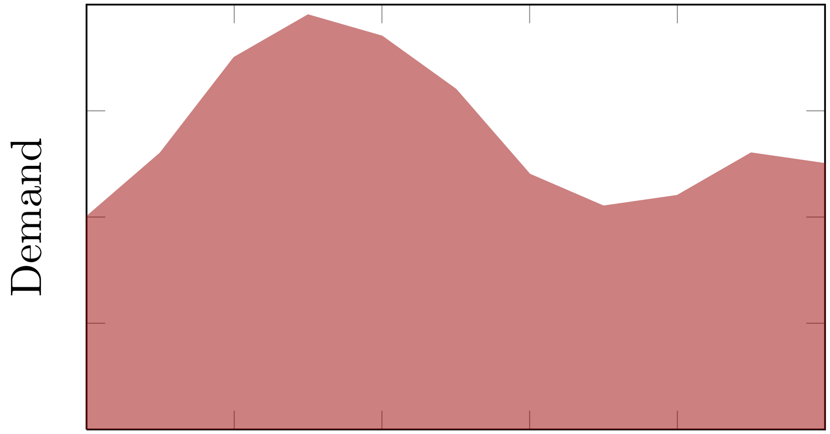

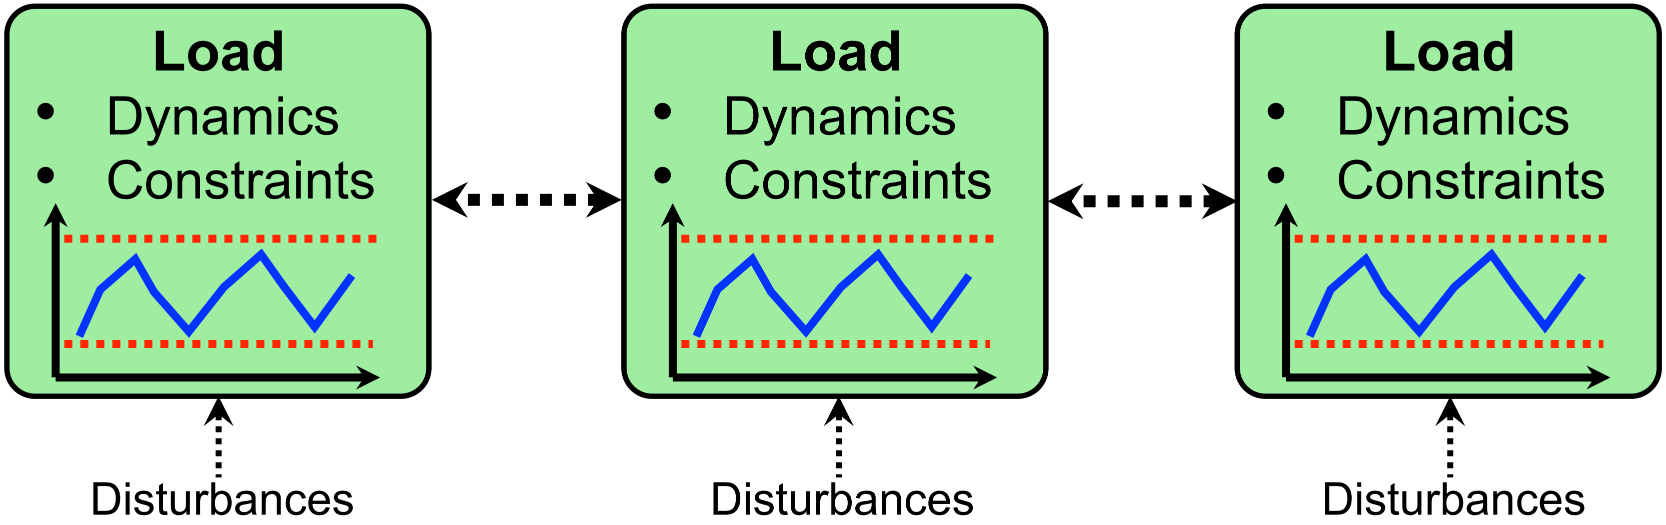

Load
- Dynamics
- Constraints

Disturbances

Load
- Dynamics
- Constraints

Disturbances

Load
- Dynamics
- Constraints

Disturbances

Green Scheduling (GS) Approach

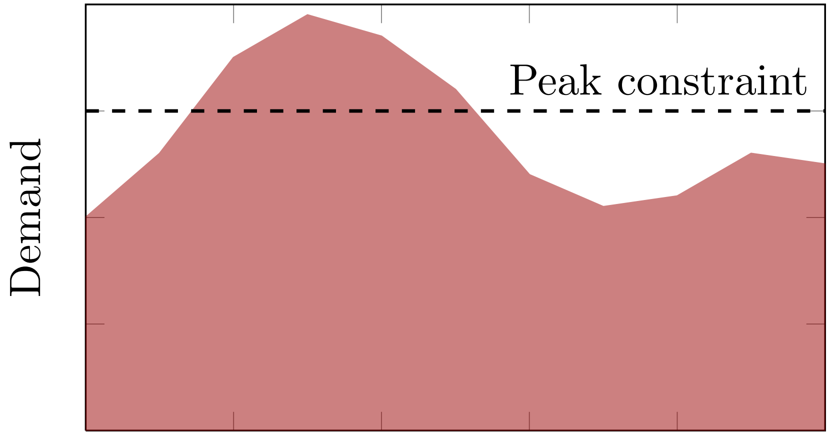

Set a peak constraint

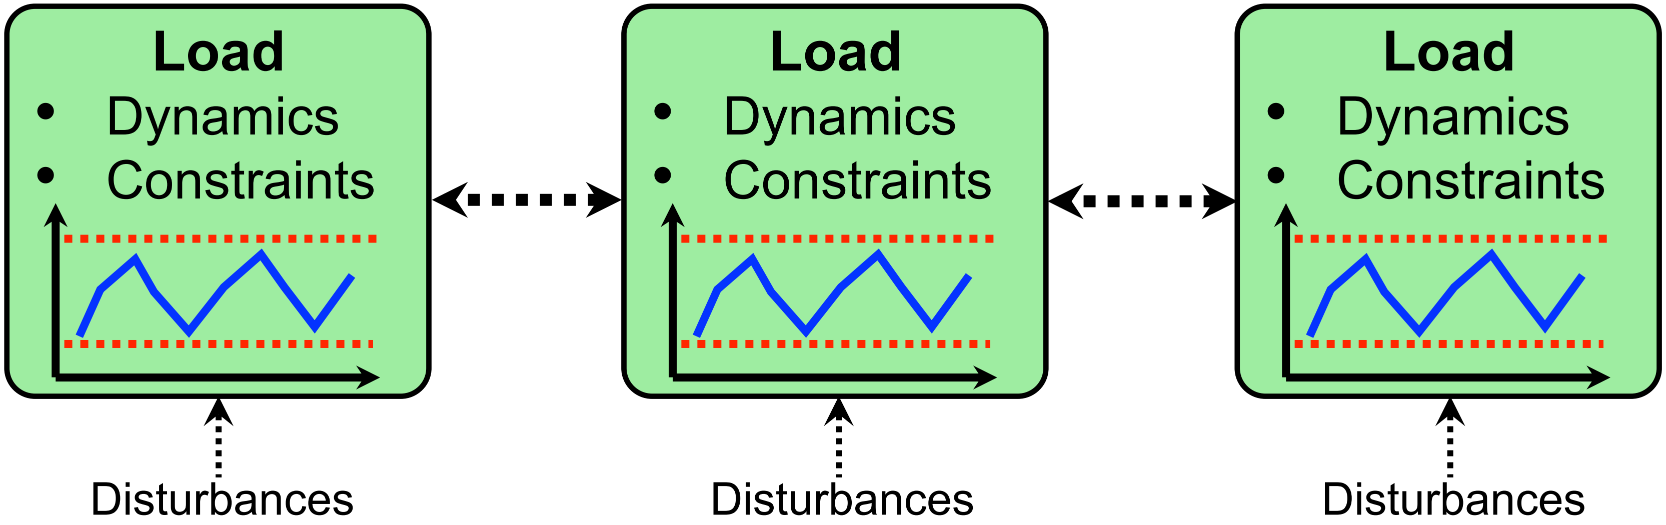
Load
- Dynamics
- Constraints

Disturbances

Load
- Dynamics
- Constraints

Disturbances

Load
- Dynamics
- Constraints

Disturbances

Green Scheduling (GS) Approach

Coordinate dynamical loads

1. *Under peak envelope*
2. *Satisfy safety constraints*

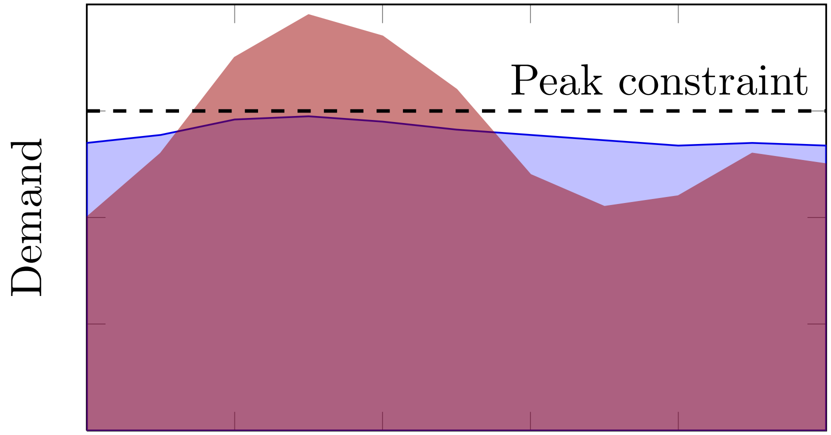

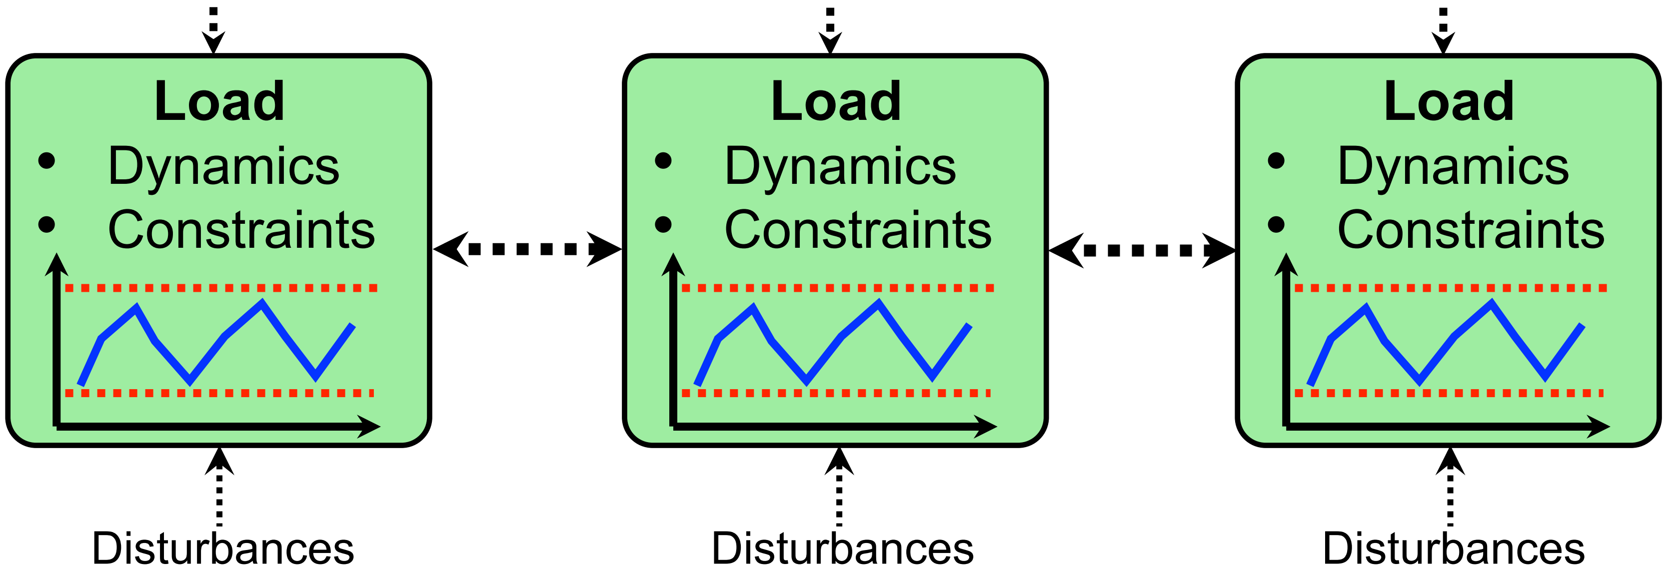

Load
- Dynamics
- Constraints

Disturbances

Load
- Dynamics
- Constraints

Disturbances

Load
- Dynamics
- Constraints

Disturbances

From Control to Scheduling

- Control loops are abstracted as tasks
- Extract temporal parameters across multiple control loops
- Compute a global schedule, reduce peak power by de-correlating systems

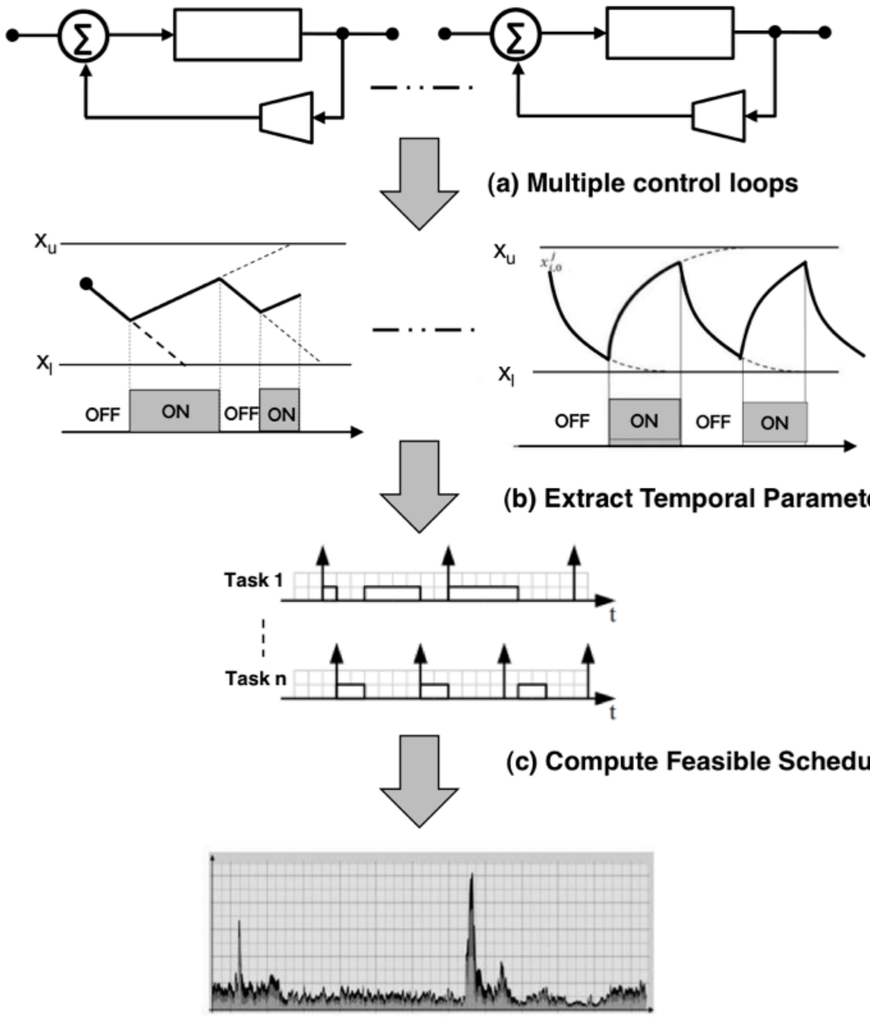

(a) Multiple control loops

(b) Extract Temporal Parameters

(c) Compute Feasible Schedule

(d) Execute Schedule

GS: Analysis & Synthesis

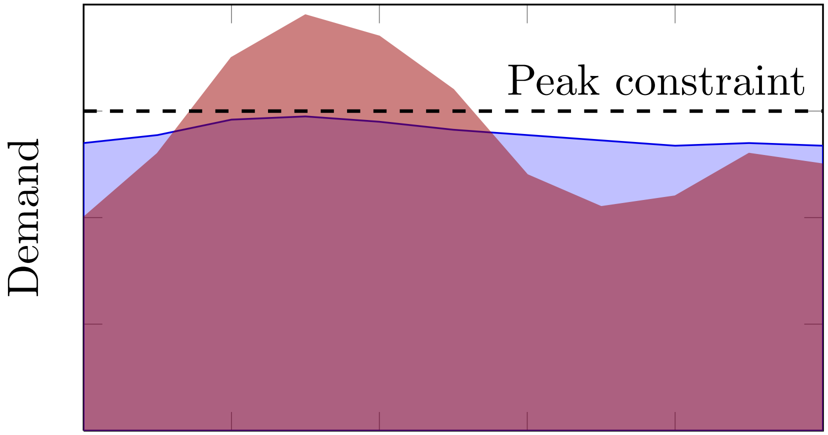

Coordinate dynamical loads

1. *Under peak envelope*
2. *Satisfy safety constraints*

Schedulability analysis

Is a peak constraint feasible? (how to choose a peak constraint?)

How to schedule the loads safely under a feasible peak constraint?

Schedule/Control Synthesis

GS Control/Scheduling Structure

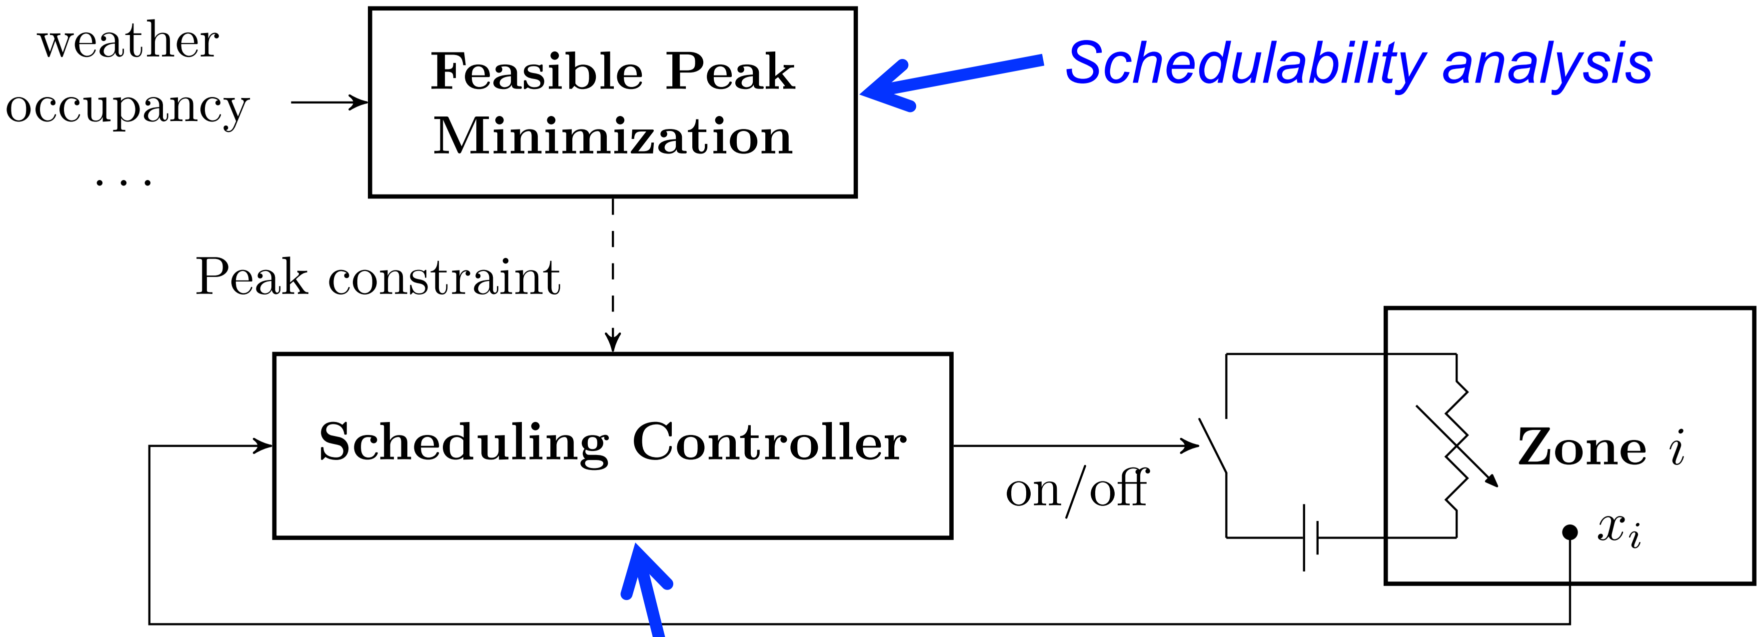

Schedulability analysis

Scheduling synthesis

GS Control/Scheduling Structure

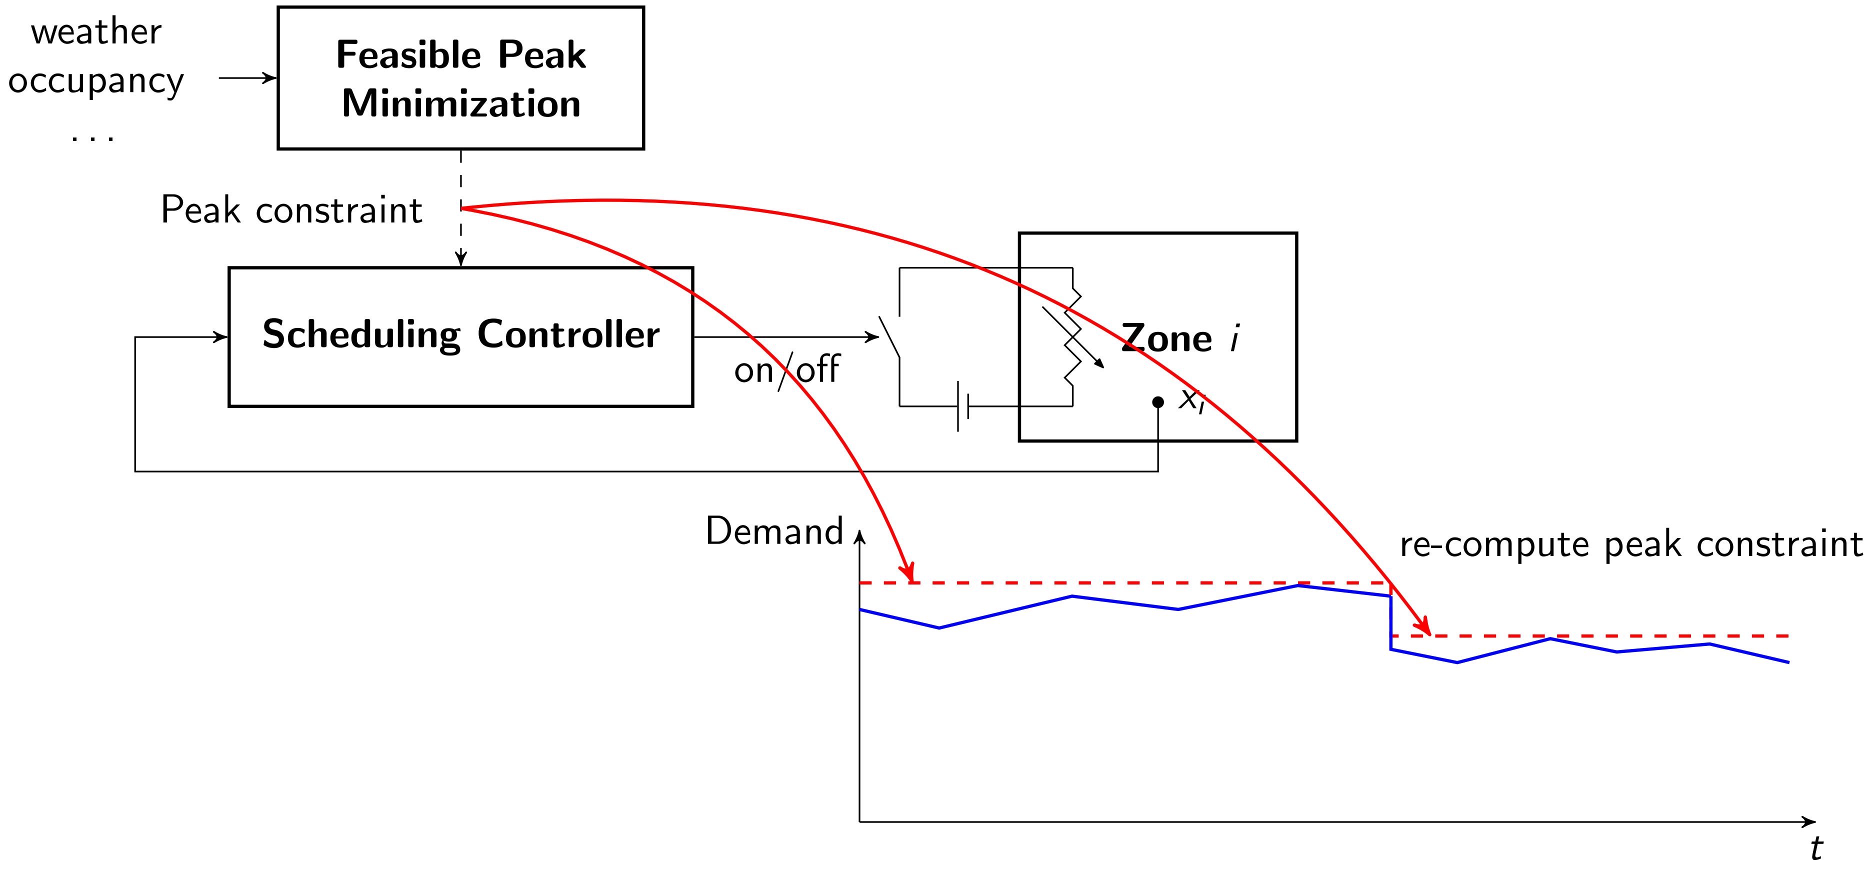

From Control to Scheduling

In Energy Control Systems

- *Execution time* is dependent on
    - **Plant Dynamics**: dimensions of the room, ingress and egress airflow
    - **Environmental Conditions**: outside weather, human occupancy, air quality
    - **Initial State**
- Tasks have **elastic execution times** where a task may have to perform more work, the longer its response time.
- Aim: keep the state of a system within a deadband

Both are resource contrained problems

- Here the resource is electricity/energy as opposed to a processor
- PSU (power supply unit) scheduling instead of CPU scheduling.

Task Model

- Tasks satisfy this differential equation: $x(t+1) = dyn(m(t), x(t))$
- Two operational modes, $M = \{\text{ON}, \text{OFF}\}$

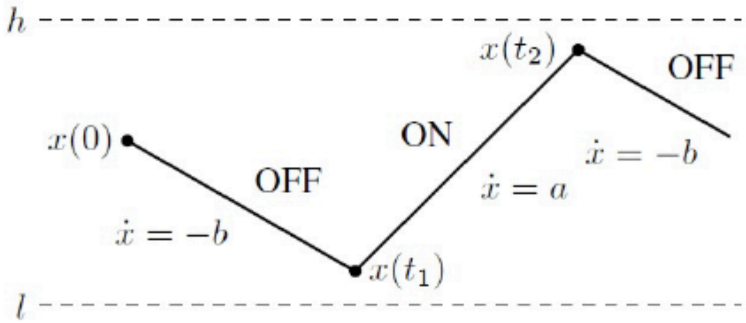

- *Linear tasks* have dynamics defined as:

$$\dot{x}(t) = dyn(m(t), x(t)) = \begin{cases} a & \text{if } m(t) = \text{ON} \\ -b & \text{if } m(t) = \text{OFF} \end{cases}$$

- A task set $\mathcal{T}$ is *schedulable* by a policy π if π can schedule the tasks in $\mathcal{T}$ so that they are all safe and that all system-wide constraints are met.

De-correlation Constraint

At most one task can be in the ON mode at any time

First Order Task Model

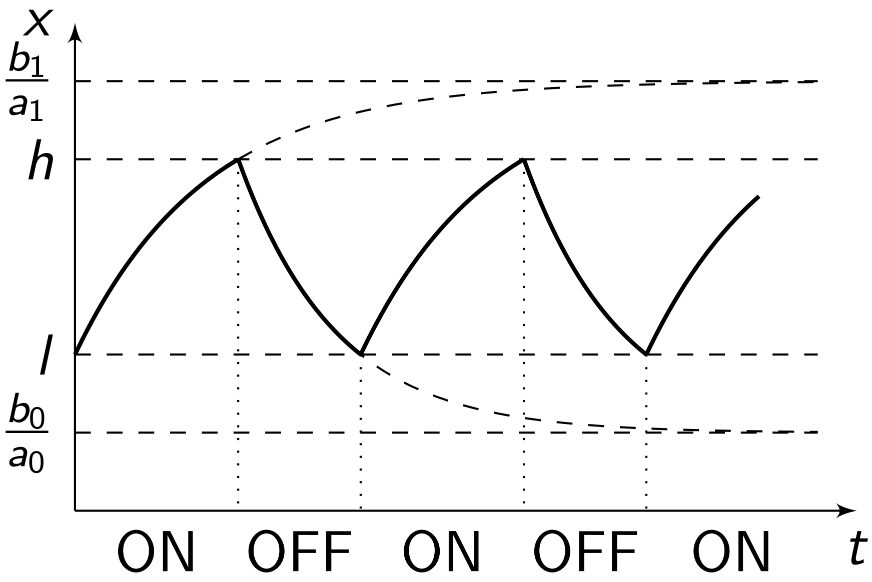

Simplified heat balance equation to model each zone

$$C_i \frac{\mathrm{d}x_i}{\mathrm{d}t} = K_i \left(T_a - x_i \right) + Q_i \qquad (1)$$

Zone temperature is governed by the following affine differential equation:

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = -\frac{K_i}{C_i} x_i + \left(\frac{K_i}{C_i} T_a + \frac{Q_i}{C_i} \right) = -a_i x_i + b_i \qquad (2)$$

Geometric Interpretation

Intuitive and simple framework for scheduling a system of linear tasks.

- Two linear tasks T_1 and T_2, normalized so that their bounds are both $[0, 1]$
- Define a 2-dimensional state vector $x = [x_1, x_2]^T \in \mathbb{R}^2$

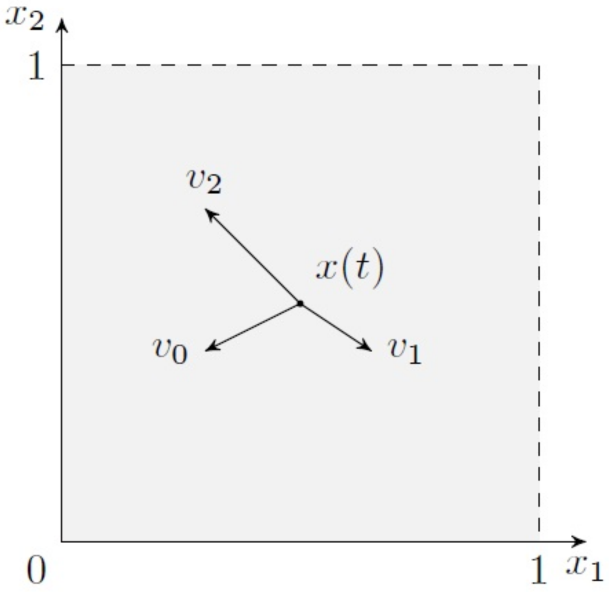

There are three scheduling modes:

- **Mode 0**: T_1 and T_2 are OFF (vector v_0)
- **Mode 1**: T_1 is ON and T_2 is OFF (v_1)
- **Mode 2**: T_1 is OFF and T_2 is ON (v_2)

Scheduling Policy

Keeps $x(t)$ within bounds (invariant set) using mode vectors v_0, v_1 and v_2

Scheduling Policy as Hybrid Automaton

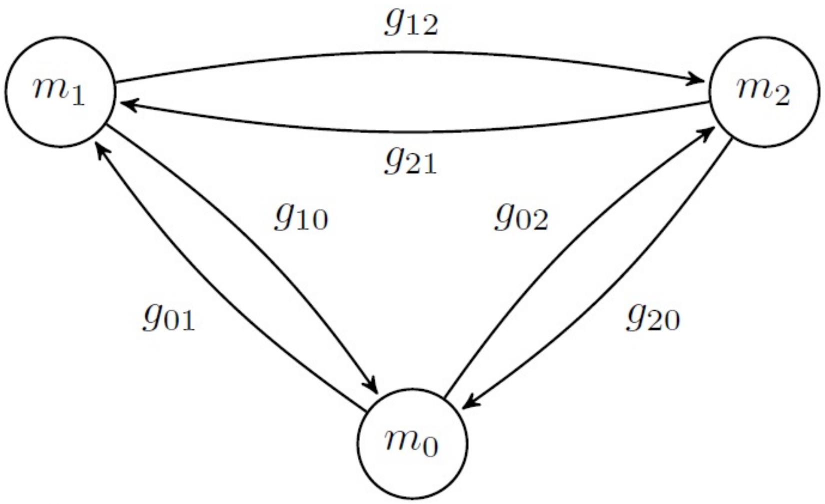

- States $\overset{correspondto}{\rightarrow}$ scheduling modes
- Edges $\overset{correspondto}{\rightarrow}$ switching between modes
- g_{ij} is the guard associated for each edge, for the transition from mode i to mode j.
- Scheduling policy π for the task set is simply a set of guards $\{g_{ij}\}$

Lazy Scheduling Policy: Hybrid Automaton

Lazy Policy: All tasks stay in their current modes as long as they are safe

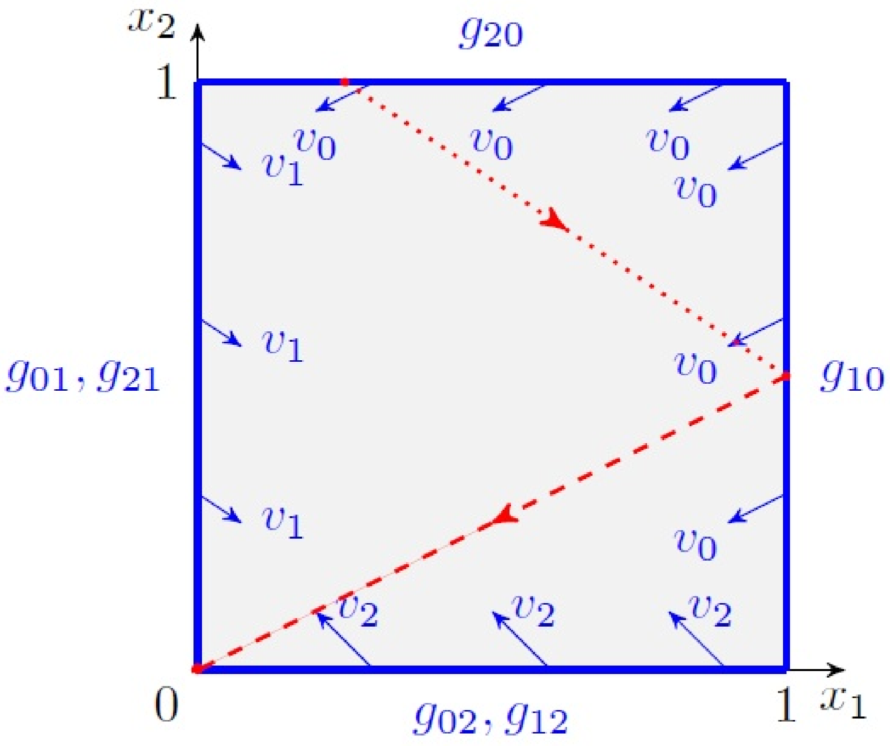

- g_{01} and g_{21} are both $(x_1 \leq l_1)$;
- g_{02} and g_{12} are both $(x_2 \leq l_2)$;
- g_{10} is $(x_1 \geq h_1 \wedge x_2 > l_2)$;
- g_{20} is $(x_2 \geq h_2 \wedge x_1 > l_1)$.

The Lazy Scheduling Algorithm

Lazy, because switching decisions are made only if a task is approaching its thresholds

- **Critical zone**
- **Alarming zone**
- **OFF zone**

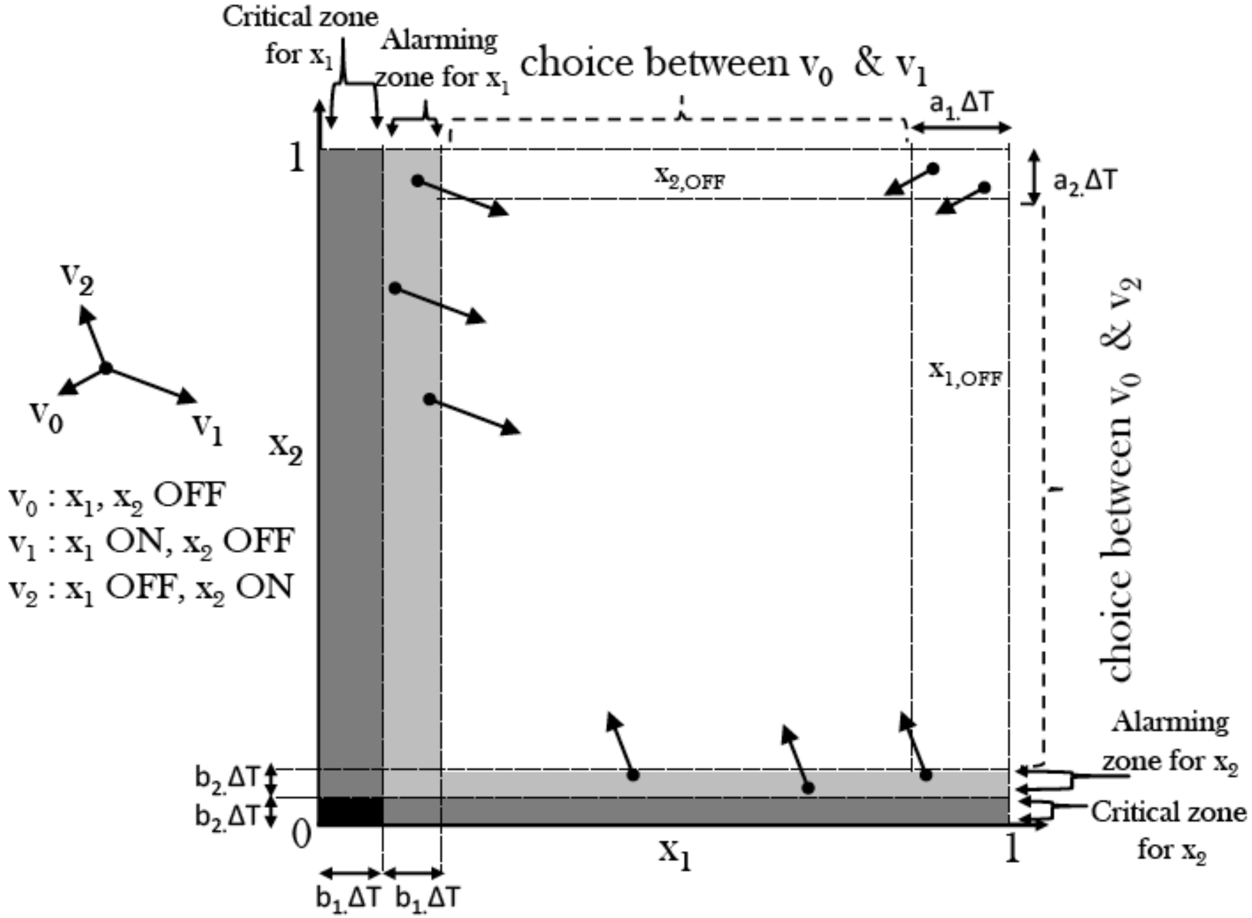

Simulation: Two-Task System

- Feasibility constraint: Keep temperature centered around mean 70°F
- Heating system operates with a power of 12000 BTU/h or 3.517 kW
- Cooling occurs through heat loss and does not consume any extra power
- Each time step of the algorithm is of 15 minute duration

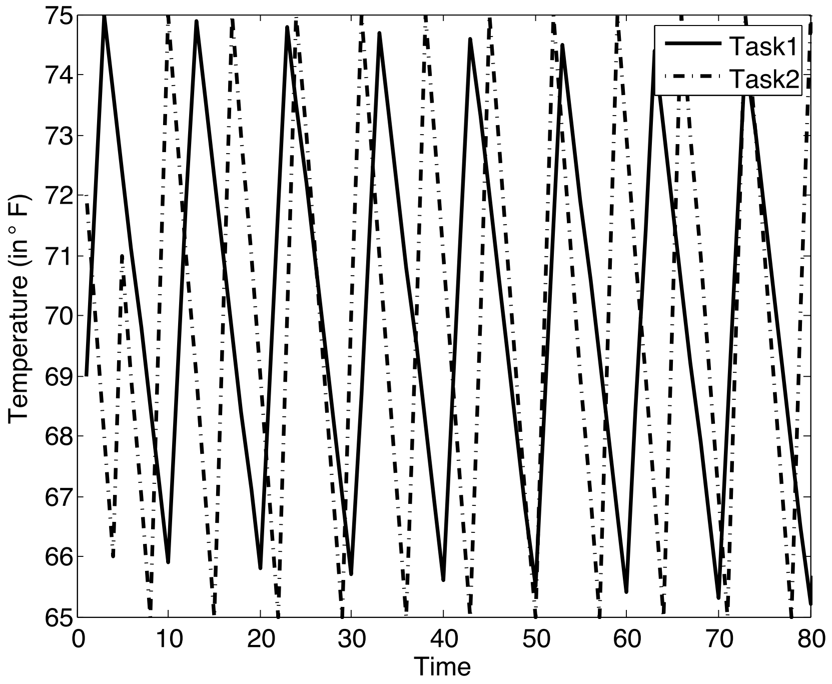

Figure: Peaks occur when tasks run independently

Simulation: Peak Reduction

Peak Reduction of 50%

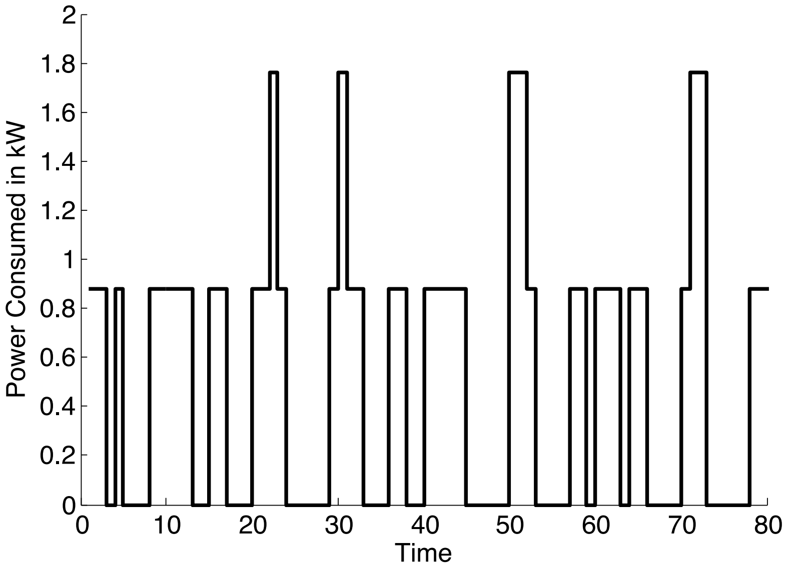

Figure: Power Consumption: Independent Tasks

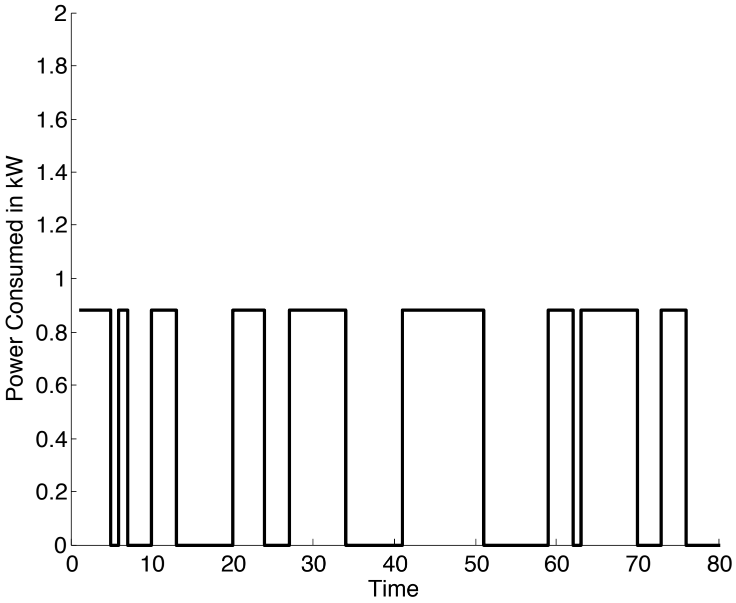

Figure: Power Consumption: Lazy Scheduling

- Peak Power = 1.758 kW
- Total Energy = 50.11 kWh

- Peak Power = 0.879 kW
- Total Energy = 45.72 kWh

Simulation: Lazy Scheduling

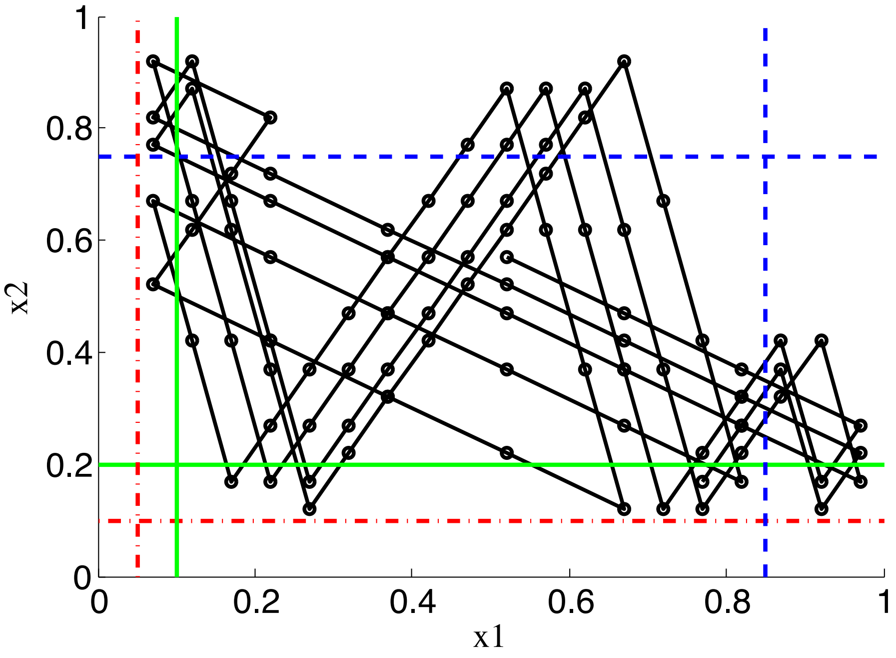

Figure: State-Space Trajectory

- Tasks remain within thresholds
- Modes only change in a *lazy* manner (near thresholds)

Simulation: Lazy Scheduling for 3 Tasks

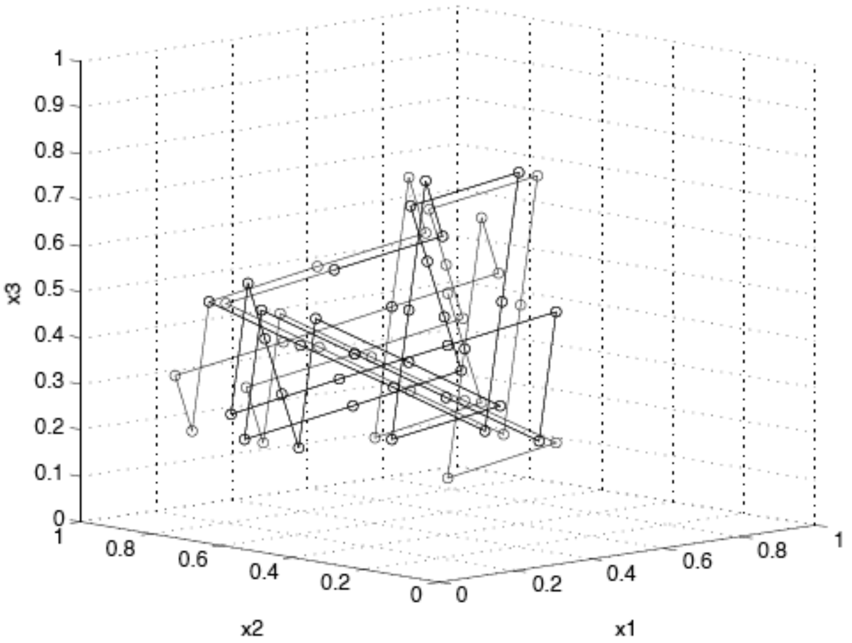

Figure: State-Space Trajectory

Green Scheduling – So far…

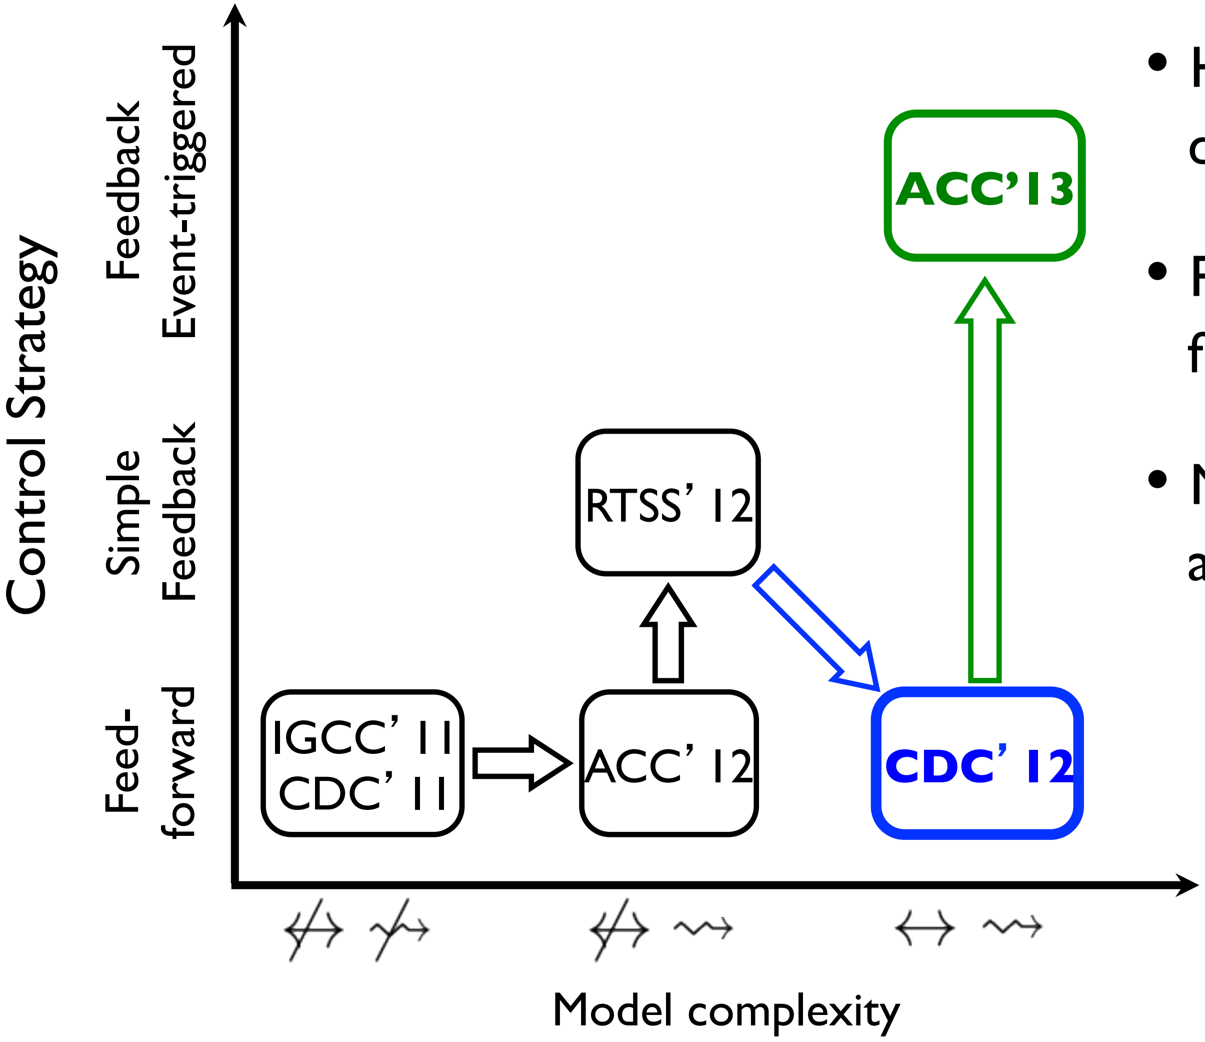

- Hierarchical, distributed control
- Pricing signal & financial cost
- More realistic applications

↔ with & without
↮ load interaction
⇝ with & without
⇝̸ disturbances

APPLICATION:

ENERGY-EFFICIENT OPERATION OF MULTIPLE CHILLER PLANTS

Peak Demand Reduction

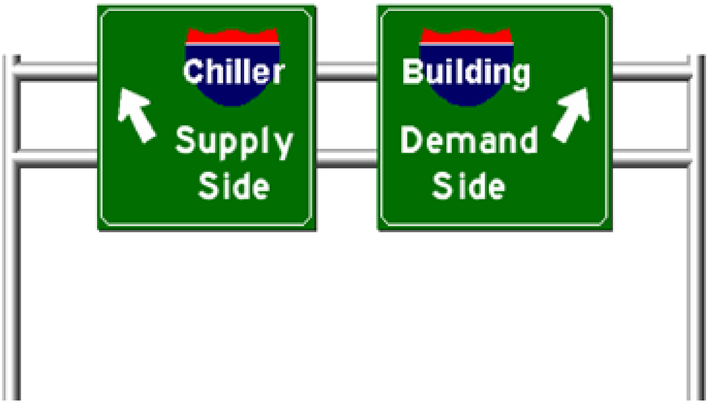

- ✓ Chillers
- ✓ Cooling Towers
- ✓ Thermal Energy Storage
- ✓ Condensers
- ✓ Underground mains …

- ✓ Variable Air Volume Box (VAV)
- ✓ Air Handling Units
- ✓ Occupant comfort
- ✓ Solar heat gain …

Chiller Plant Operation at Penn

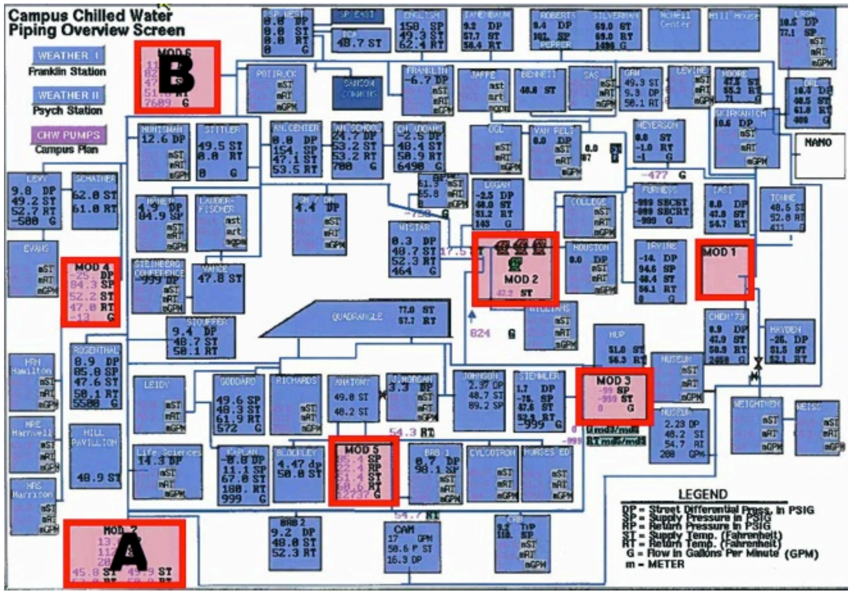

Chilled water distribution system
at the University of Pennsylvania

Chiller Plant A

- Over **4M gallons** of chilled water (42 F) pumped into the campus.
- Plant A consumes **26 MW** at full capacity
- Plant A and Plant B account for > **30%** of total peak power consumption (108 MW)

COP vs. PLR

- Approximate COP by quadratic function of PLR

$$COP = a_0 + a_1\ PLR + a_2\ PLR^2$$

Example:

If cooling load = 1000 RT[1],

chiller capacity = 1250 RT,
then PLR = 1000/1250 = 0.8,

and COP obtained from
COP-PLR curve.

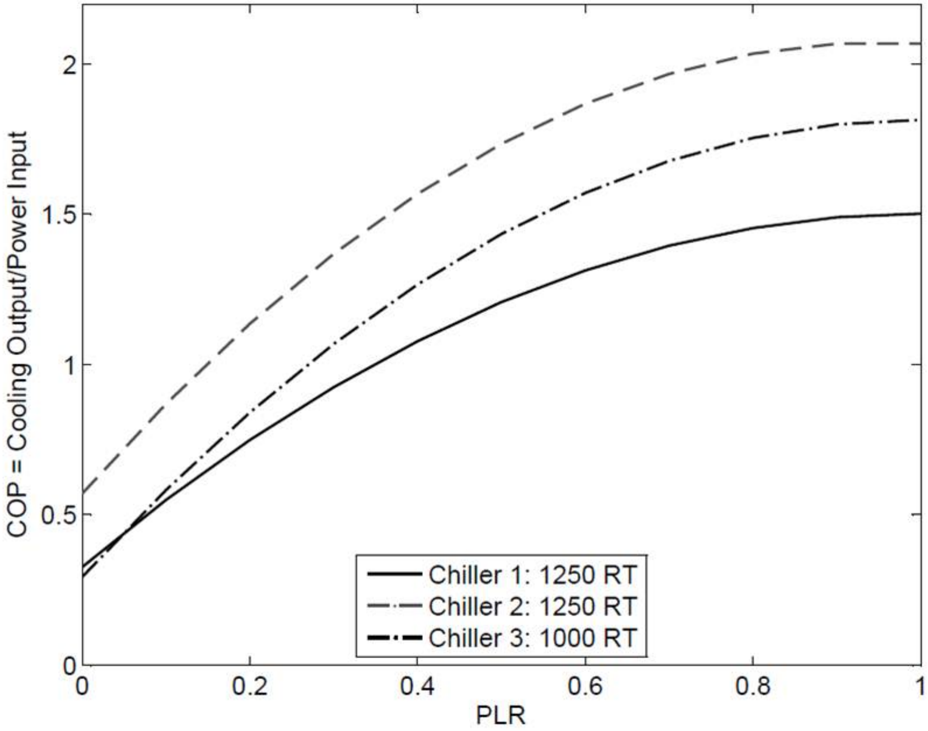

[1] Refrigeration Ton is the amount of heat that must be removed to melt 1 Ton of ice in 24h. 1 RT = 3517W

Thermal Energy Storage (TES)

Thermal Energy Storage is used for demand shifting and off loading chillers.

- **Long term**: > 10 hours
- **Short term**: < 2 hours

TES at Cornell University

Short-term Thermal Energy Storage can improve the COP of the chiller plant.

How TES Improves the COP?

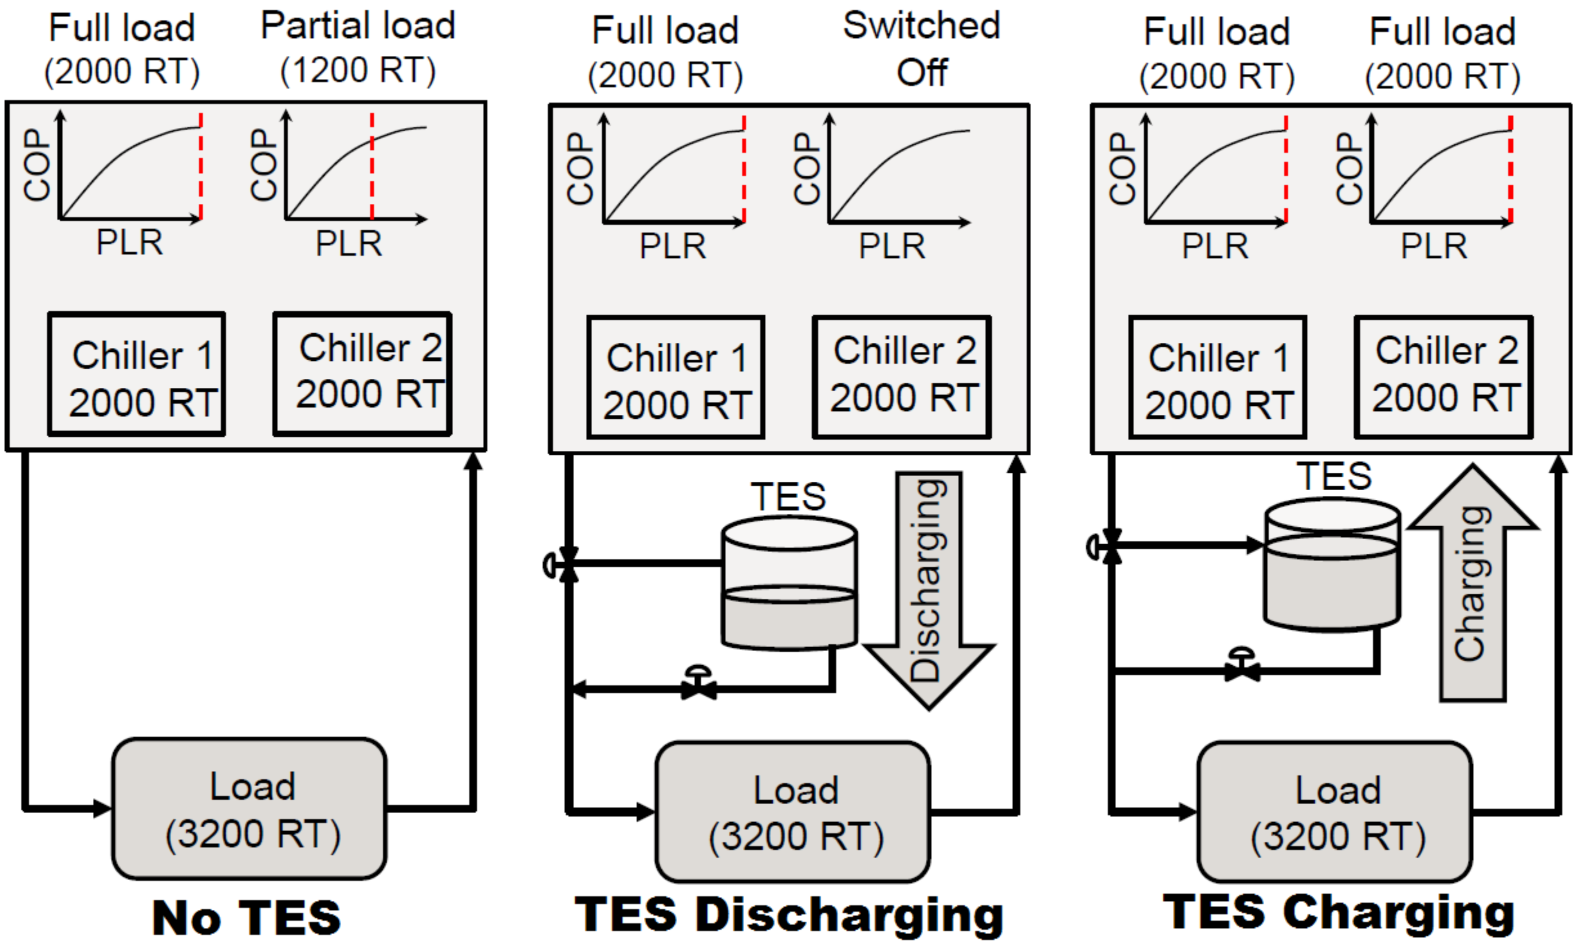

No Thermal Energy Storage

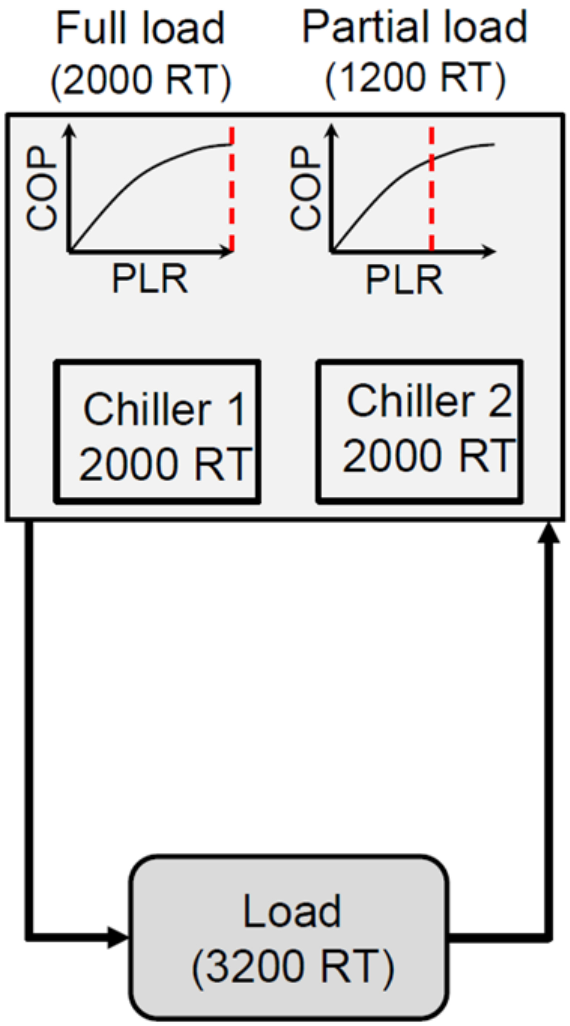

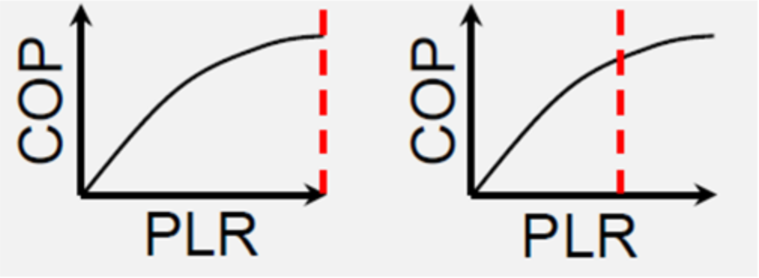

Inefficient

TES Discharging

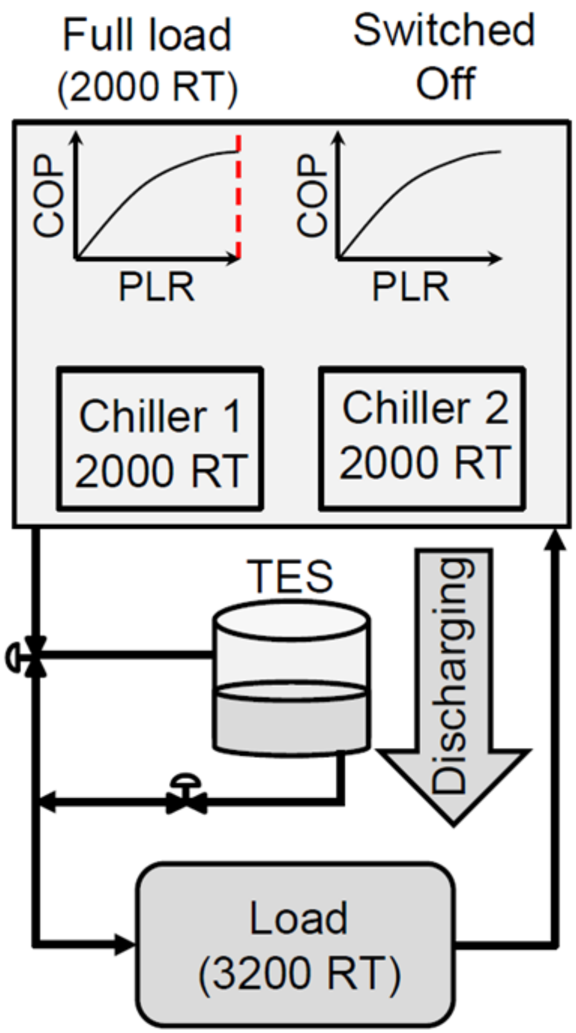

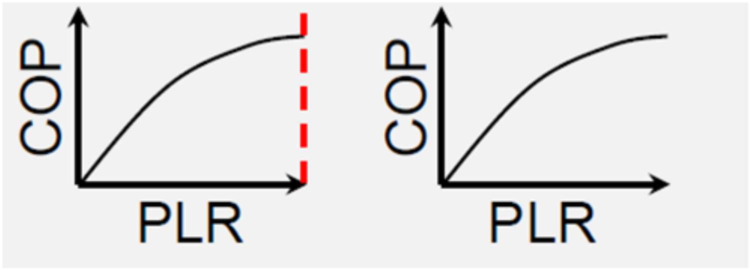

Efficient!

TES Charging

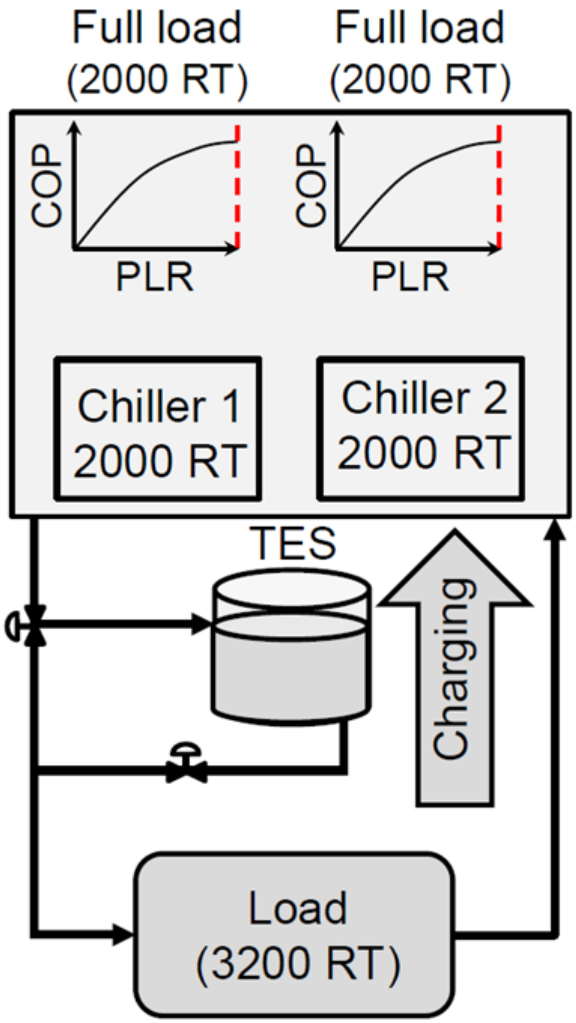

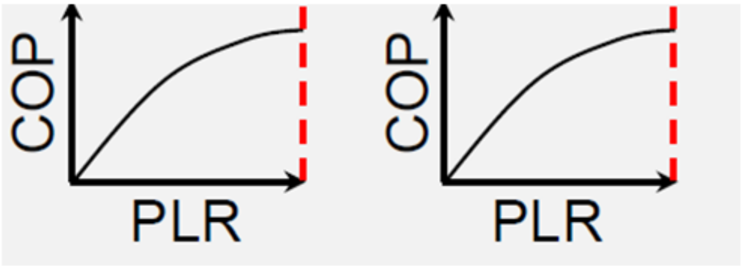

Efficient!

Multiple chiller plants with TES

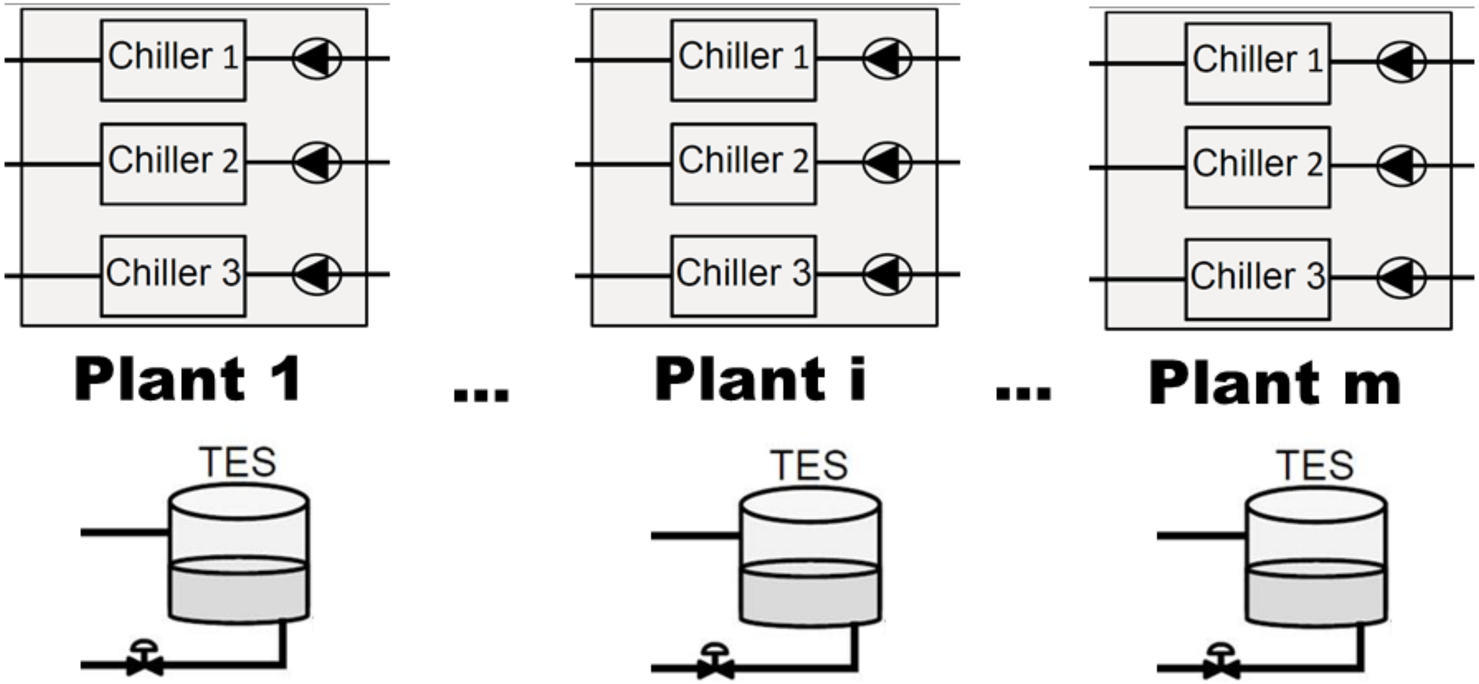

Uncoordinated operation among multiple chiller plants can cause large spikes in total electricity consumption.

Our goal is to coordinate multiple chiller plants to reduce their aggregate peak power consumption.

GS for Chiller Plants

- Consider m > 1 chiller plants.
- Each plant has a short term TES system.
- Compute modes for each hour h, $h = 0,1,\ldots,H$, based on load forecasting.

- **Charging mode**: water level increases with rate $a_{i,h} > 0$;
- **Discharging mode**: water level decreases with rate $b_{i,h} < 0$.

Behl, M.; Nghiem, T. X.; Mangharam, R., *Green Scheduling for Energy-Efficient Operation of Multiple Chiller Plants*, RTSS 2012.

Single chiller plant with TES

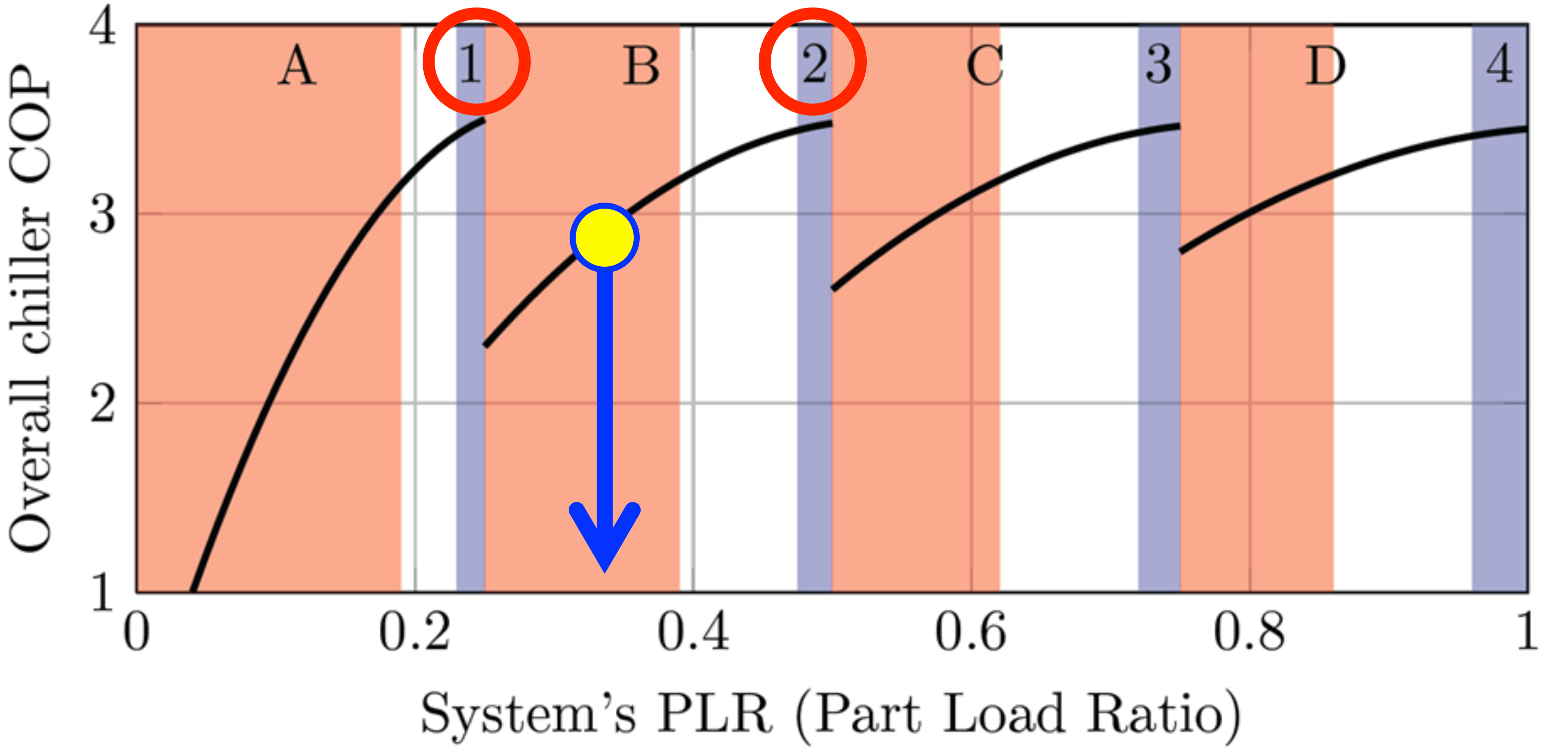

- Say PLR = 0.34 i.e., in region B
- Plant can operate in optimal regions 1 and 2 if a TES is available.

GS for Chiller Plants

Schedule the operating modes $u_i(t)$, $t \geq 0$, of all plants such that:

- Safety Constraint:
  $l_i \leq x_i(t) \leq h_i \quad \forall t, i$

- Peak Constraint:
  $$\sum_{i=1}^{n} u_i(t) \leq k(t) \quad \forall t$$

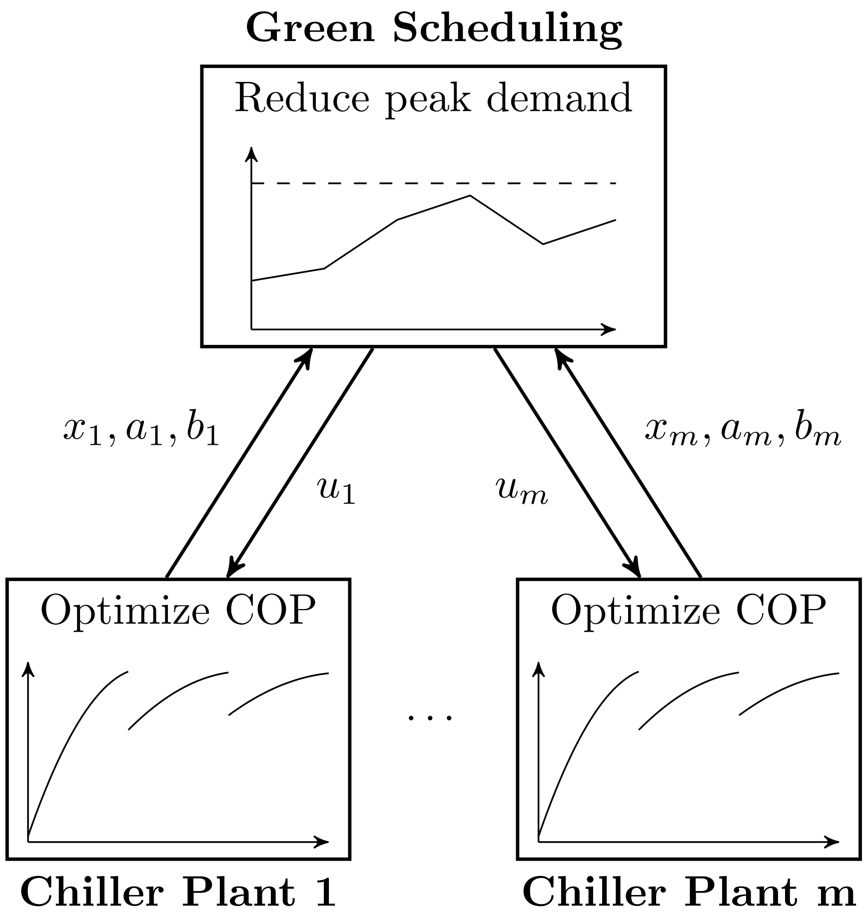

Simulation Setup

Three chiller plants, each containing 3 chillers.

- Lower safety threshold: 3m
- Upper safety threshold: 13.5m
- Time step = 15 mins

Chiller plant configuration in the simulation

Plant 1	3 chillers rated at 1250 RT, 1200 hp
Plant 2	3 chillers rated at 1250 RT, 1200 hp
Plant 3	3 chillers rated at 1000 RT, 900 hp
T_{chws}	5.5°C (42°F)
T_{cwr}	20°C (68°F)
ΔT	10°C

Average Hourly Load Profile

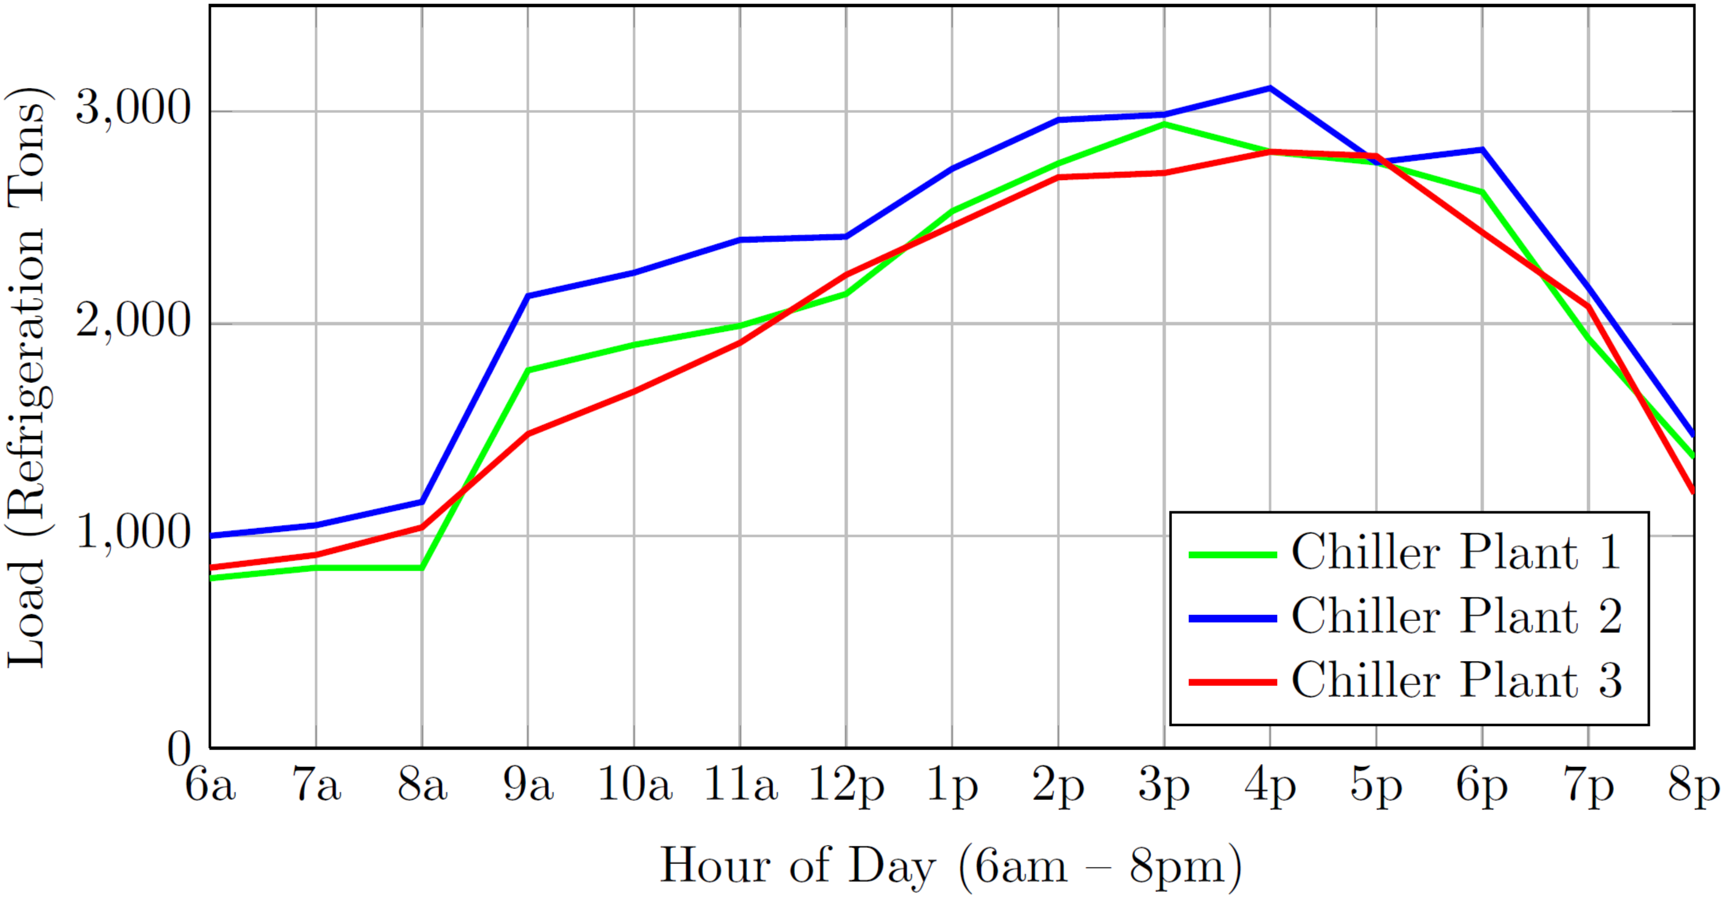

Simulation Scenarios

1. **Case 1:** No Thermal Energy Storage present
2. **Case 2:** Thermal Energy Storage with uncoordinated operation
3. **Case 3:** Thermal Energy Storage with (discrete-time) Green Scheduling

Simulation Result: Power Demand

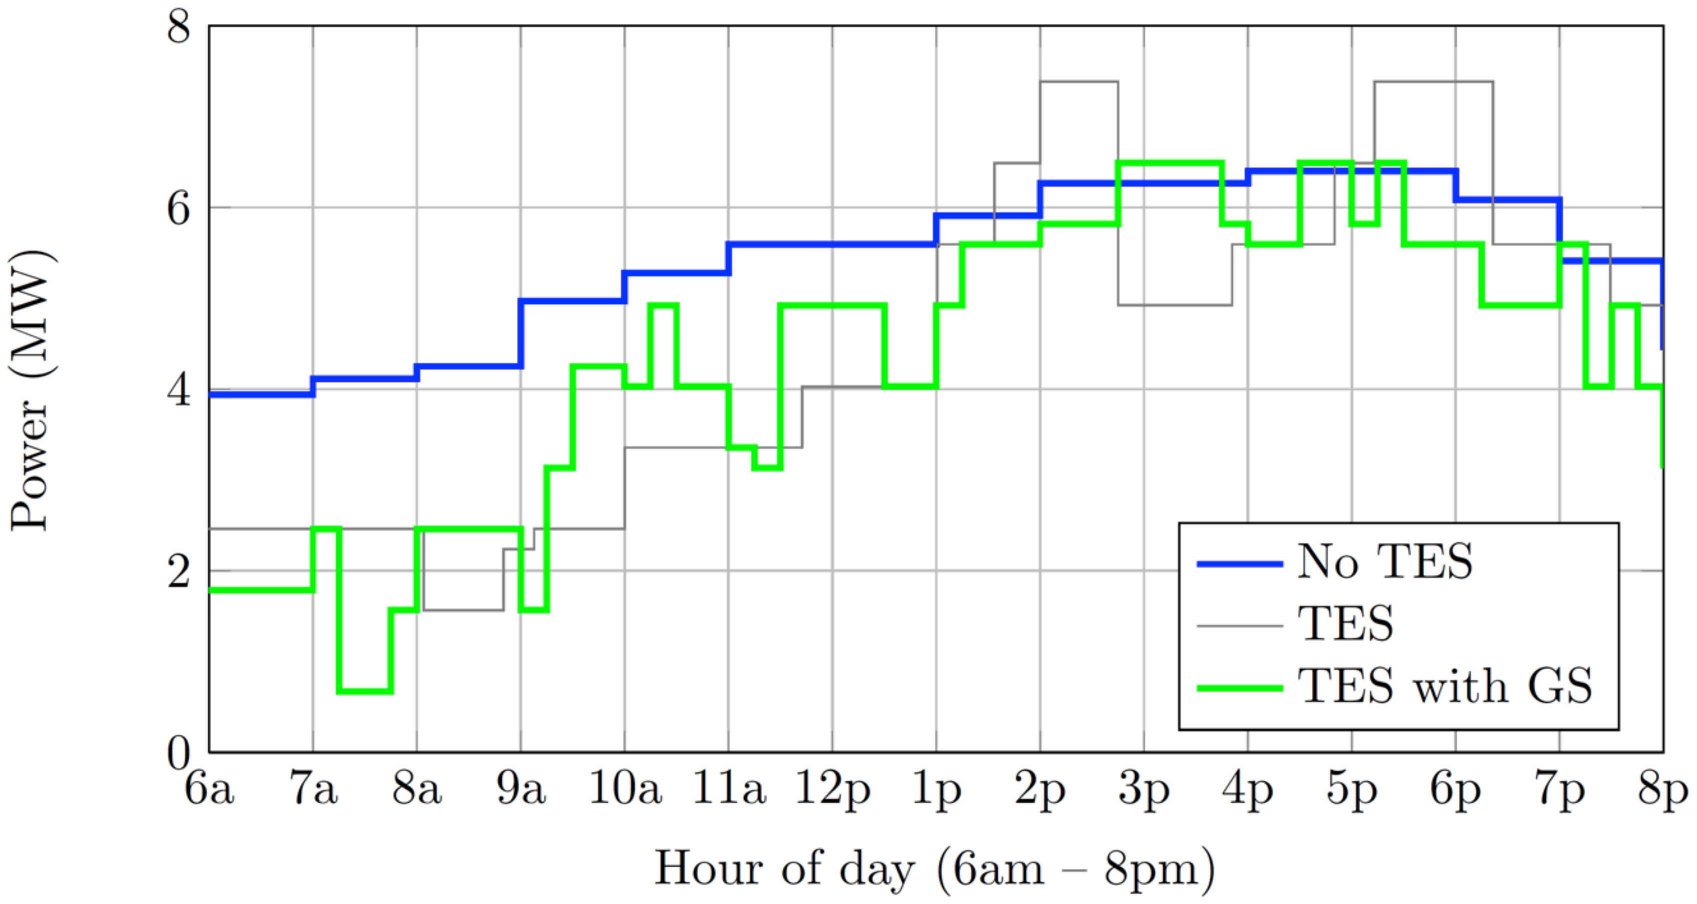

Simulation Result: Power Demand

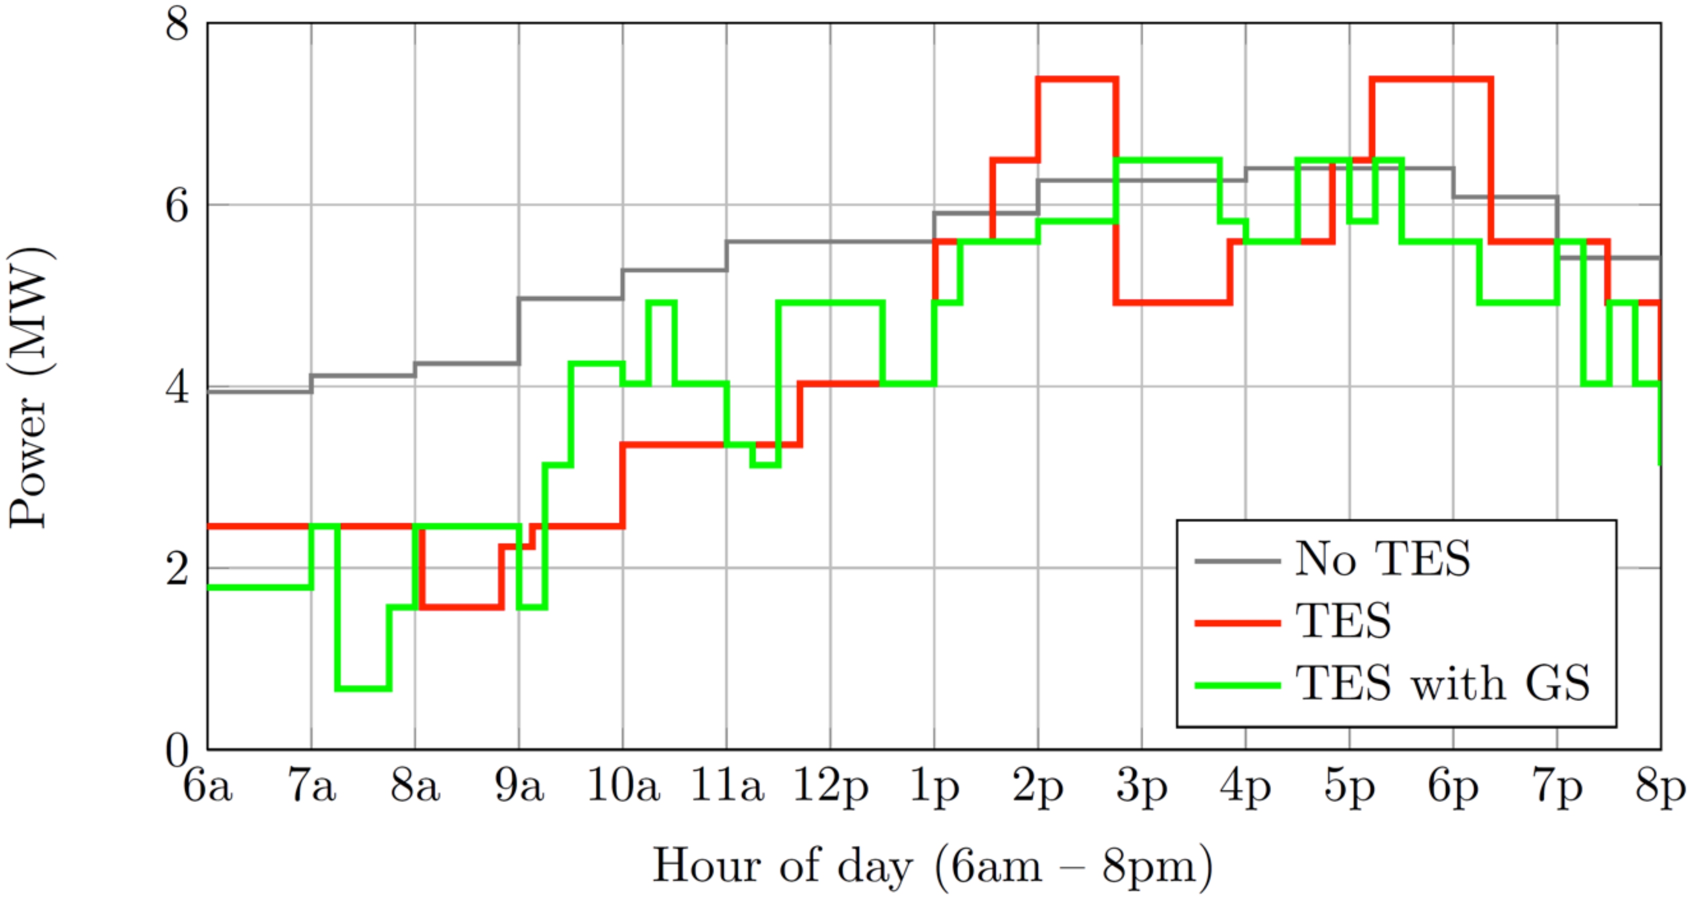

Cumulative Energy Consumption

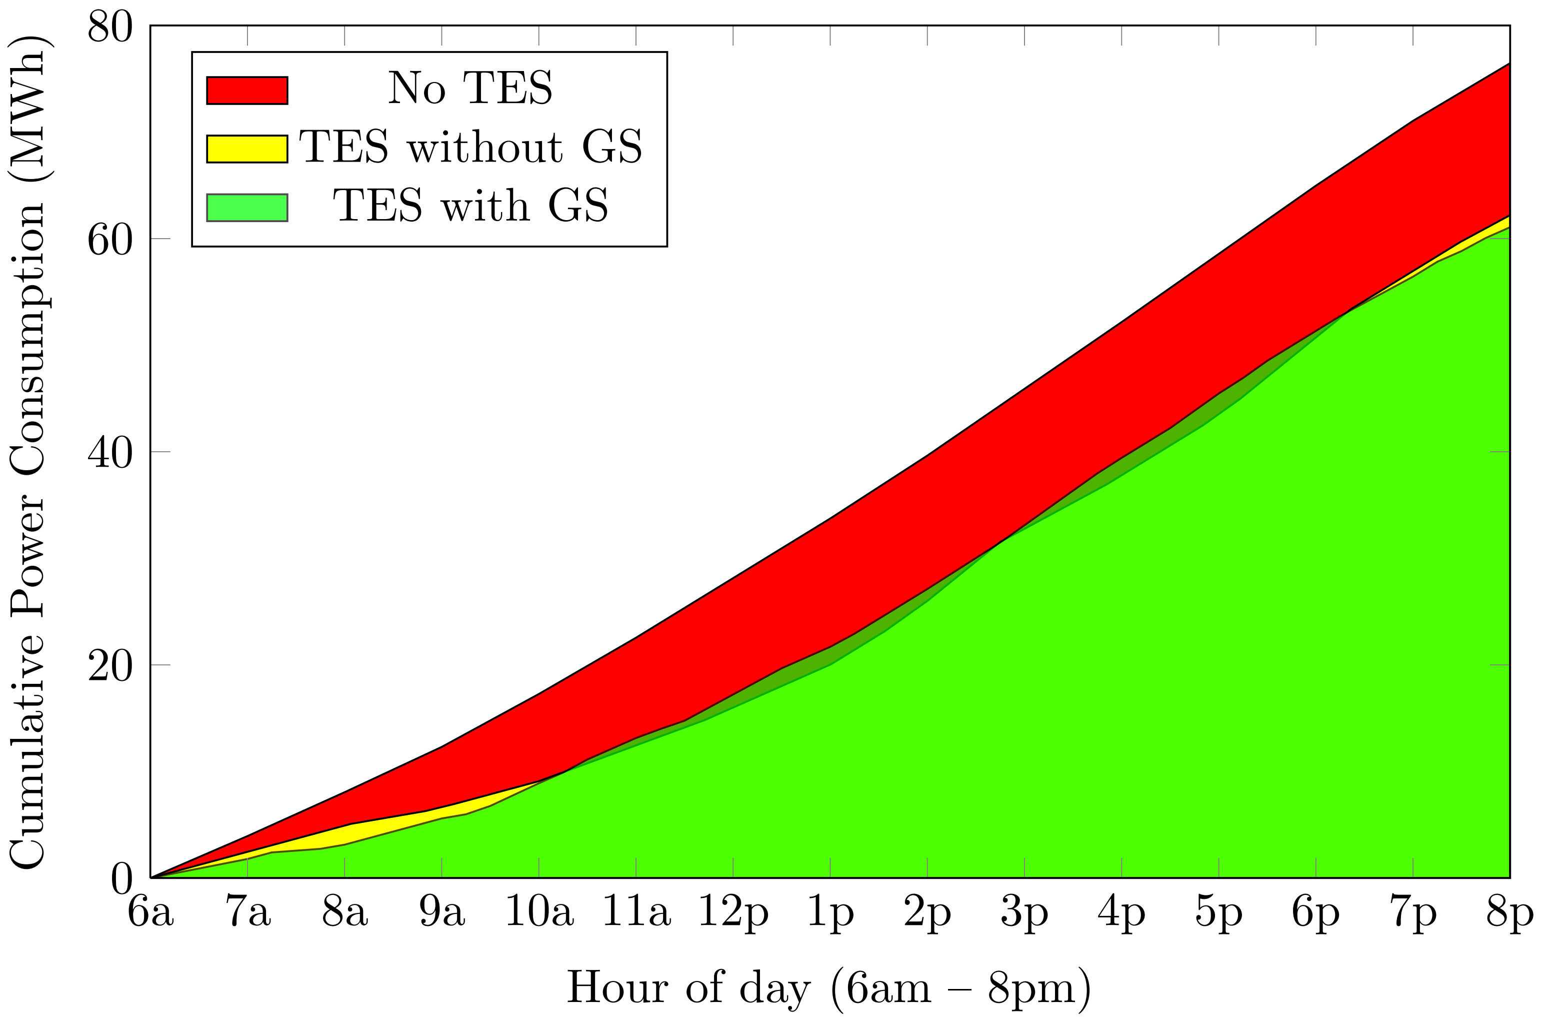

Electricity Pricing Policy

PECO's demand pricing rate structure:

Block	kWh's in Block	Charges (cents per kWh)
First Block	80 × peak	24.94
Second Block	80 × peak	12.67
Third Block	Remaining	8.64

Simulation Result Summary

Green Scheduling leads to the lowest electricity bill.

	Peak (MW)	Energy Consumption (MWh) single day	Expected Monthly Bill ($)	% savings
No TES	6.40	**80.9**	**292,801**	-
On-Off with TES	**7.38**	66.1	274,266	6.33
GS with TES	6.48	61.4	243,461	**16.85**

SCHEDULING SYNTHESIS:

DISCRETE-TIME FEEDBACK SCHEDULING

Discrete-time Green Scheduling

- Consider *discrete-time* dynamics
$$x(t+1) = Ax(t) + B_0 + Bu(t) + Wd(t),\ \forall t \in \mathbb{N}$$
for a finite time horizon *[0, T]*.
- **Safety constraint**: $x(t) \in S \subset R^n\ \forall t$
- Initial state $x(0) \in S$
- $X_f \subseteq S$: set of desired final states at time T: $x(T) \in X_f$
- Define the set of **admissible** discrete-time schedules

Define $\mathbb{U}(k(\cdot), S, x_0, X_f)$ as the set of all discrete-time schedules $u : \{0, 1, \ldots, T-1\} \to \{0, 1\}^m$ such that:

① $\sum_{i=1}^{m} u_i(\tau) \leq k(\tau)$ for all τ; ← **Peak constraint**

② $x(\tau) \in S$ for all τ; ← **Safety constraint**

③ $x(T) \in X_f$. ← **Terminal state is reached**

Robust Backward Reachable Sets

Robust backward reachable set operator:

Given any X ⊆ S at time step t, define operator

$$R_t^{-1}(X) = \{x \in S : \exists u : \sum_i u_i \leq k(t) \wedge f(x,u,d) \in X \,\forall d \in D\}$$

which is the set of safe states that can reach X with some admissible control u, regardless of the constrained disturbance d.

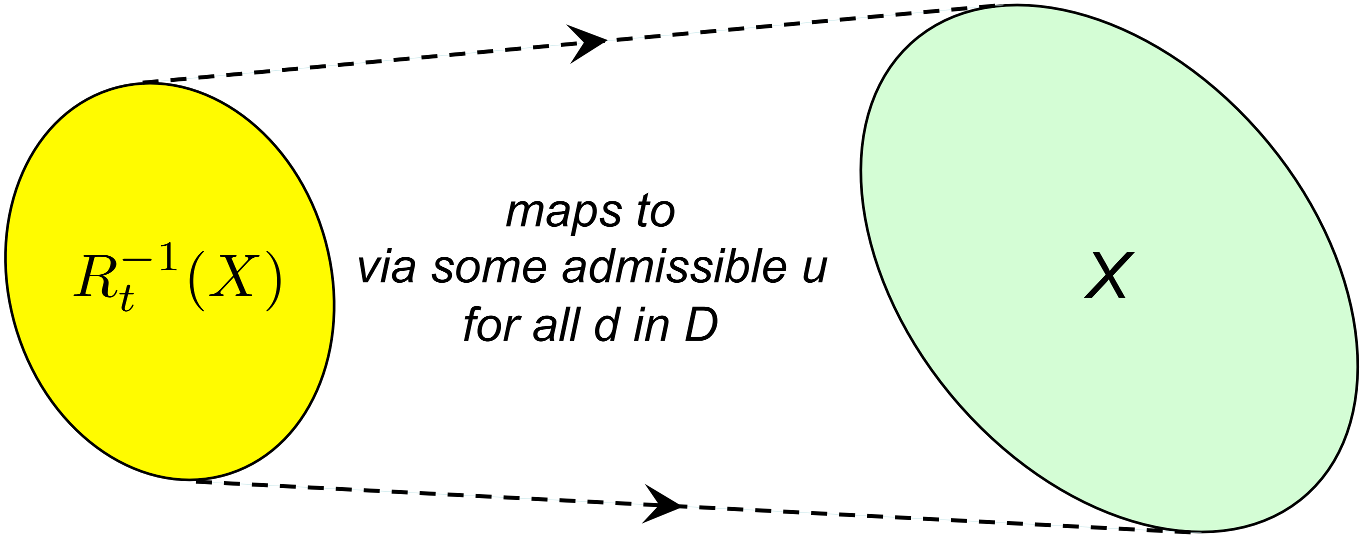

Pre-compute Backward in Time

- Using robust backward reachable set operator, compute a sequence of sets $\{X_t : t = 0,1,\ldots,T\}$ backward in time from $X_T = X_f$: $X_t = R^{-1}{}_t(X_{t+1})$

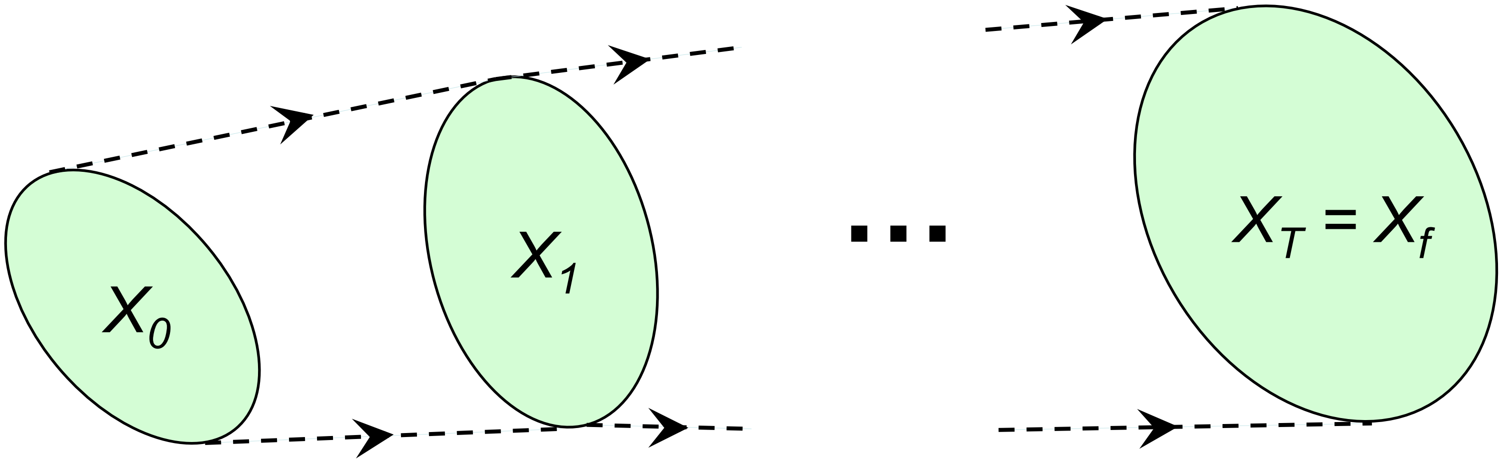

- Each X_t is the set of safe states from which and from time step t, the system state can reach the final set X_f safely under the time-varying peak constraint $k(\cdot)$.

MLE+: Matlab Toolbox for Integrated Modeling and Controls for Energy-Efficient Buildings

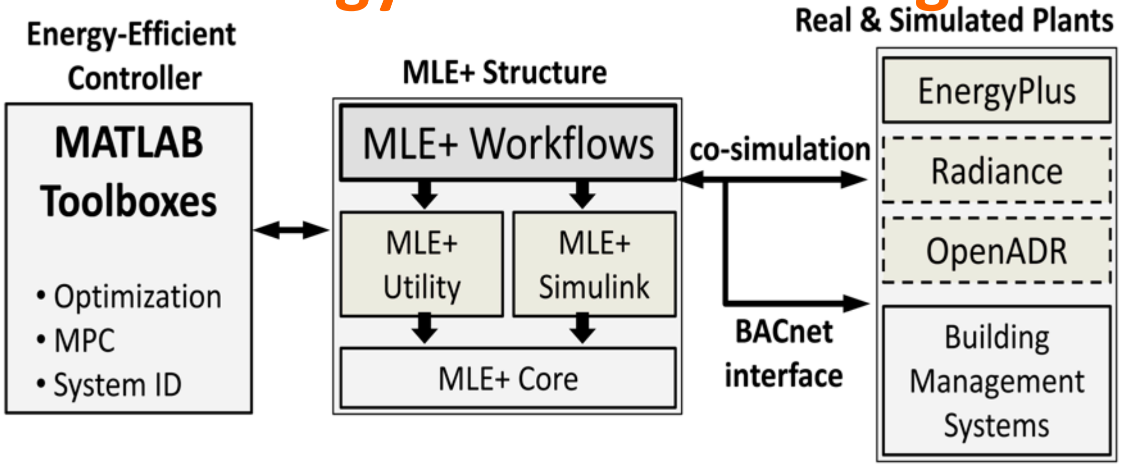

Energy-Efficient Controller

MATLAB Toolboxes

- Optimization
- MPC
- System ID

MLE+ Structure

MLE+ Workflows

MLE+ Utility

MLE+ Simulink

MLE+ Core

co-simulation

BACnet interface

Real & Simulated Plants

EnergyPlus

Radiance

OpenADR

Building Management Systems

Willy Bernal, Madhur Behl,
Truong Nghiem and Rahul Mangharam

1 EnergyPlus Building Model

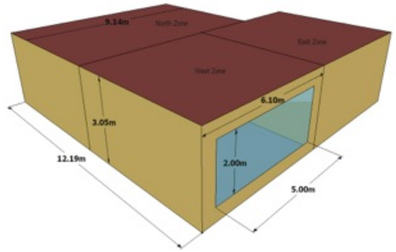

2 System Identification

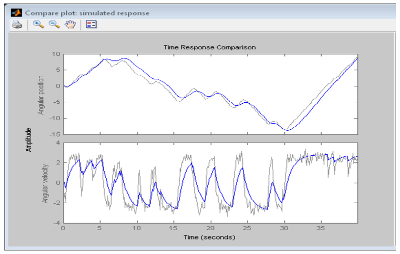

3 Control Design in Matlab

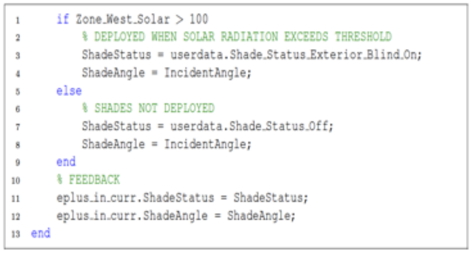

```
    if Zone_West_Solar > 100
        % DEPLOYED WHEN SOLAR RADIATION EXCEEDS THRESHOLD
        ShadeStatus = userdata.Shade_Status_Exterior_Blind_On;
        ShadeAngle = IncidentAngle;
    else
        % SHADES NOT DEPLOYED
        ShadeStatus = userdata.Shade_Status_Off;
        ShadeAngle = IncidentAngle;
    end
    % FEEDBACK
    eplus_in_curr.ShadeStatus = ShadeStatus;
    eplus_in_curr.ShadeAngle = ShadeAngle;
end
```

4 Simulation Results

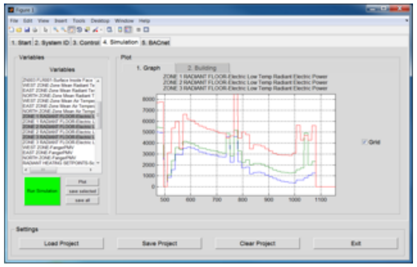

5 Control Deployment

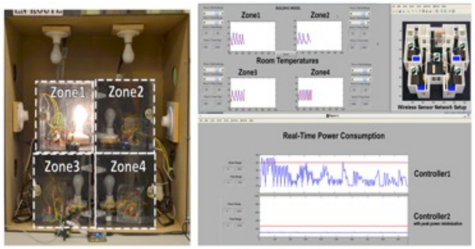

(a) Test bed with 4 zones

(b) Energy dashboard showing zone temperatures and power

Penn Engineering

A U.S. DOE Energy Innovation HUB

Re-energizing buildings for the future.™

mLAB Real-Time & Embedded

MLE+: Design and Deployment Integration of Energy-Efficient Building Controls

Willy Bernal, Madhur Behl, Truong Nghiem and Rahul Mangharam

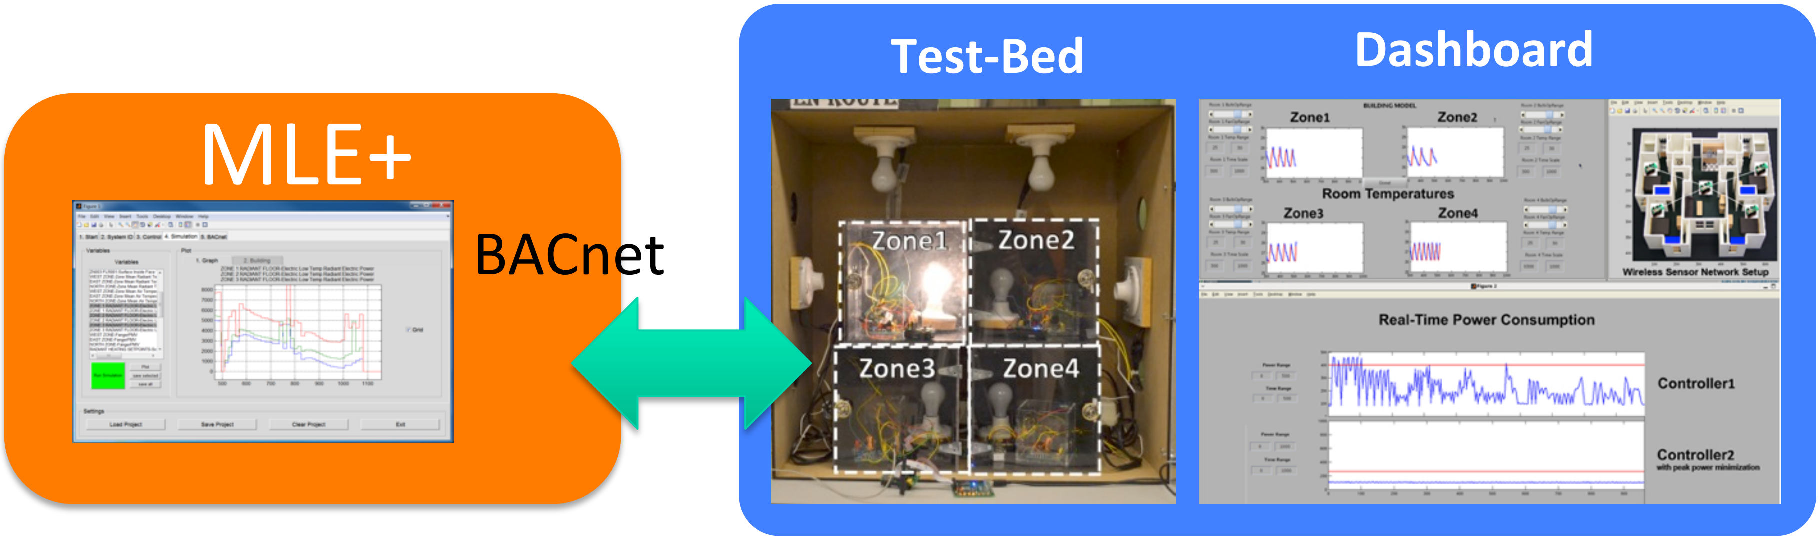

MLE+

BACnet

Test-Bed

Dashboard

Penn Engineering

A U.S. DOE Energy Innovation HUB

Re-energizing buildings for the future.™

mLAB Real-Time & Embedded

Integrated Modeling: HVAC System

HVAC System
- Temperature Setpoints according to Schedule
- Central HVAC Model
- VAV Boxes

CentralHVAC

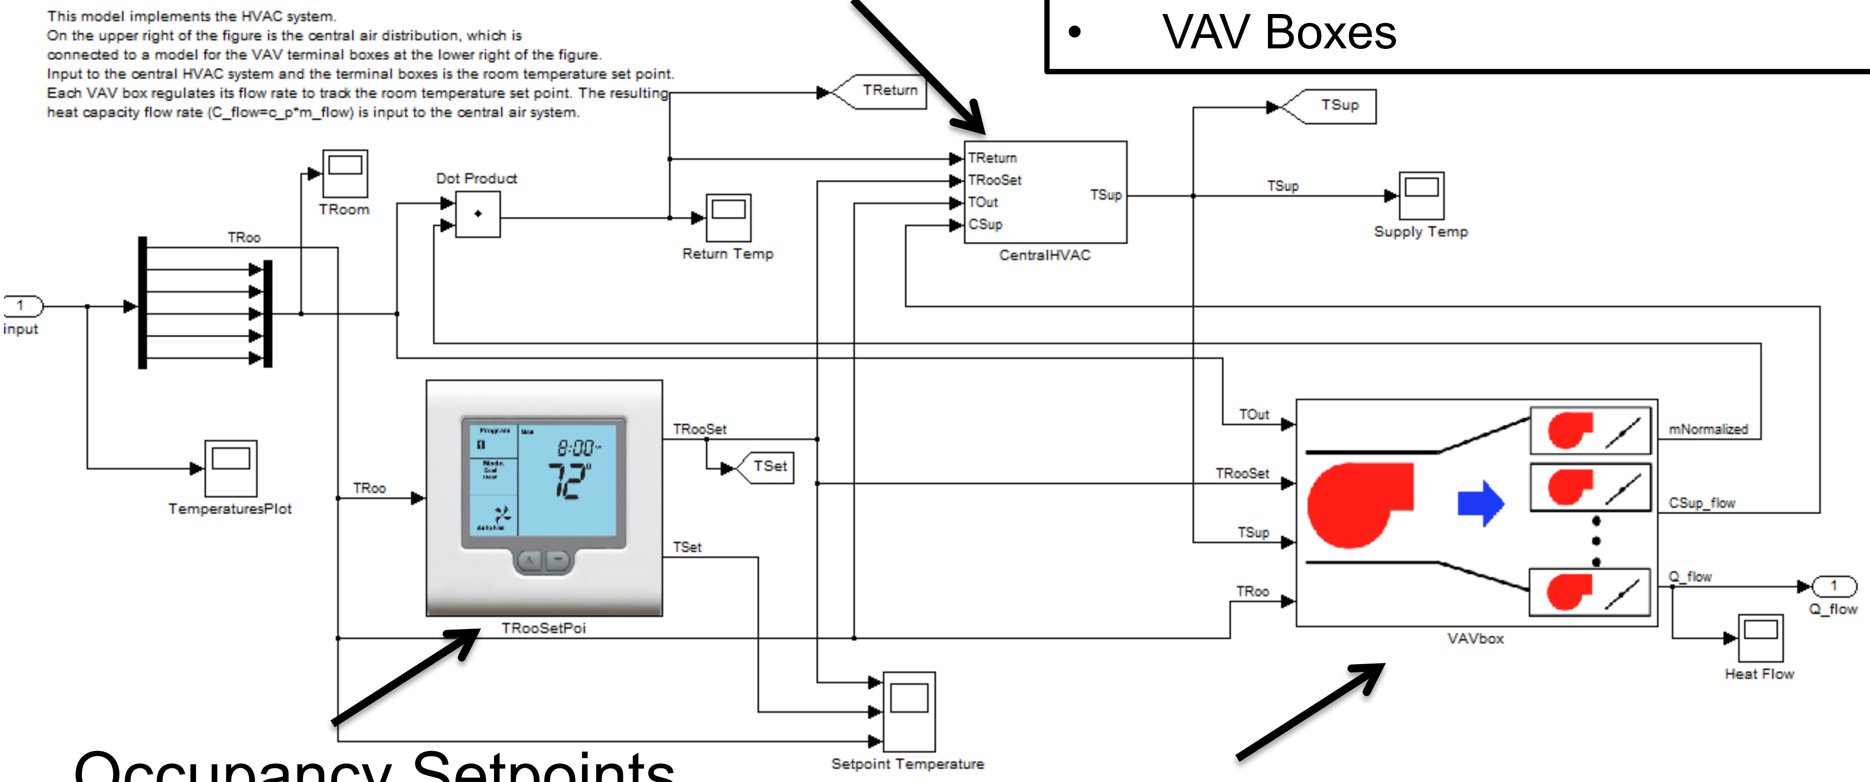

This model implements the HVAC system.
On the upper right of the figure is the central air distribution, which is connected to a model for the VAV terminal boxes at the lower right of the figure.
Input to the central HVAC system and the terminal boxes is the room temperature set point.
Each VAV box regulates its flow rate to track the room temperature set point. The resulting heat capacity flow rate (C_flow=c_p*m_flow) is input to the central air system.

Occupancy Setpoints

VAV Boxes

NREL: Campus-Wide Simulation

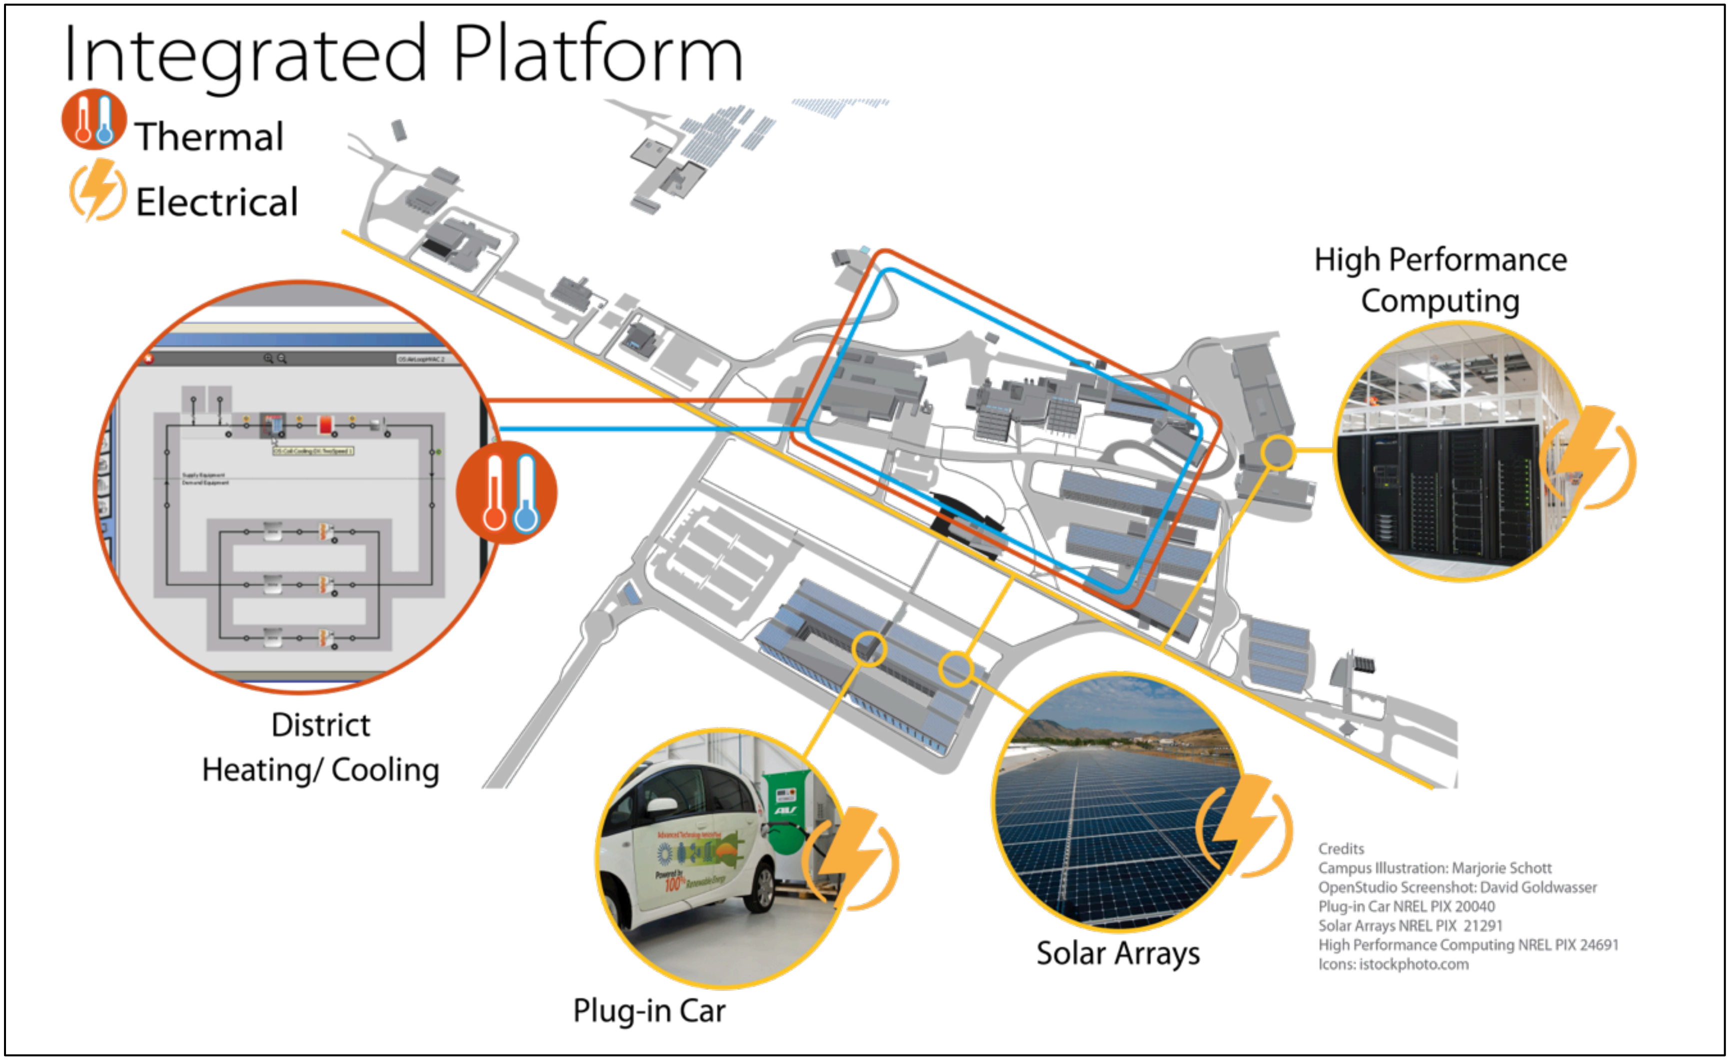

Credits
Campus Illustration: Marjorie Schott
OpenStudio Screenshot: David Goldwasser
Plug-in Car NREL PIX 20040
Solar Arrays NREL PIX 21291
High Performance Computing NREL PIX 24691
Icons: istockphoto.com

Courtesy of NREL

Campus Simulation: Two Building Example

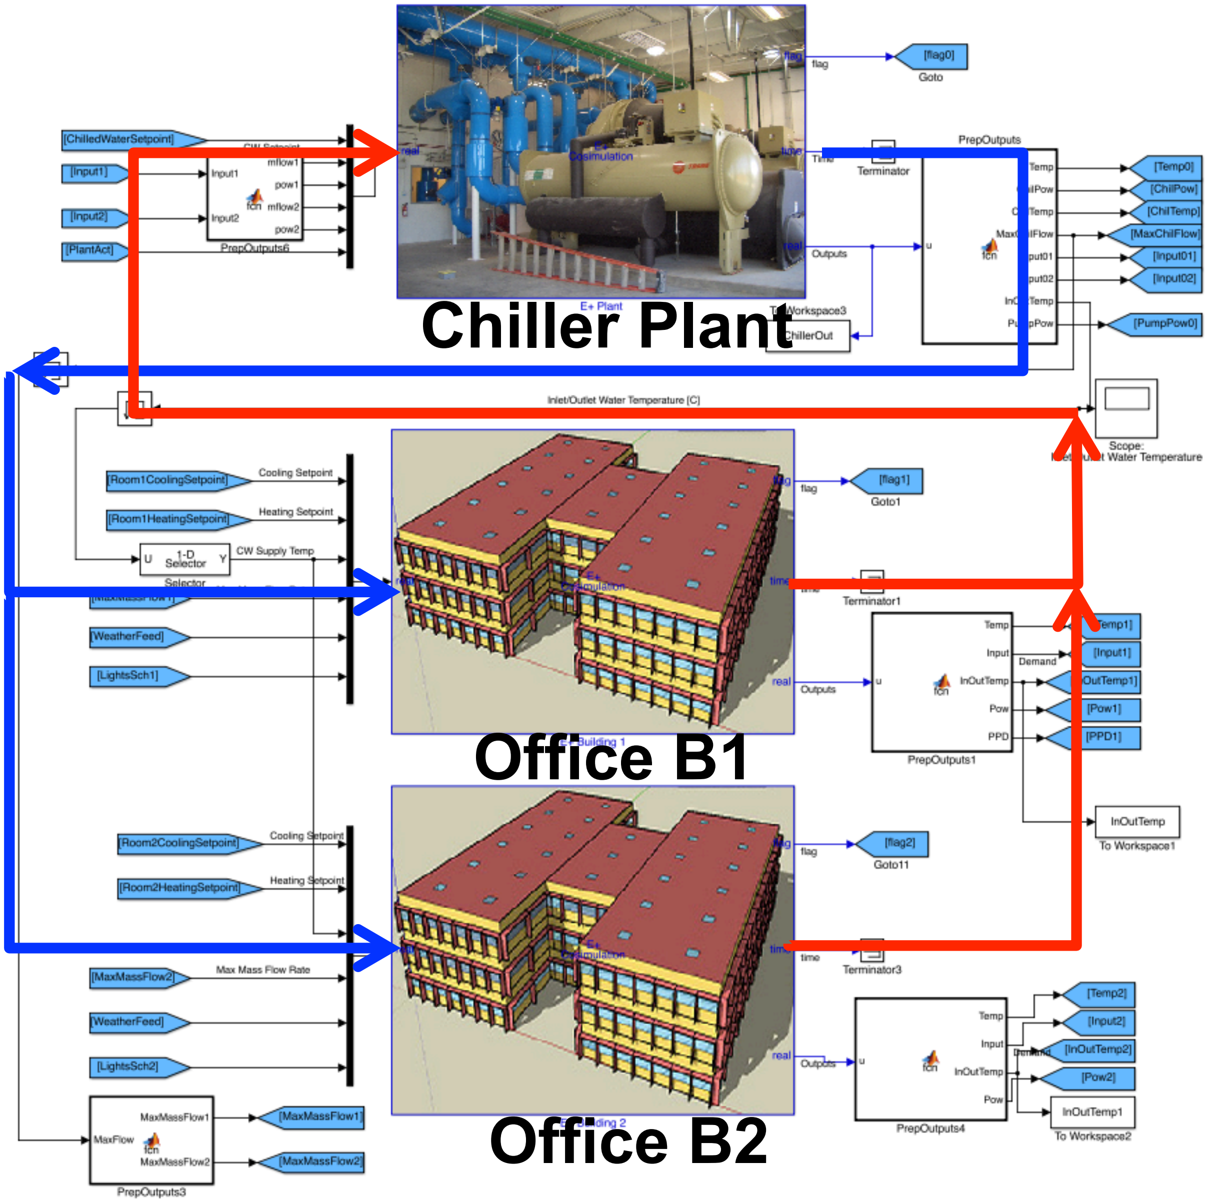

Return

Supply

Overview

- ✓ Air-cooled Electric Chiller.
  - ○ 10kW
- ✓ 2 Small Office Buildings.
  - ○ 50 people
  - ○ 5 zone
  - ○ 400m²
- ✓ Shared Chilled Water Loop.

MLE+ Cloud Analytics: Overview

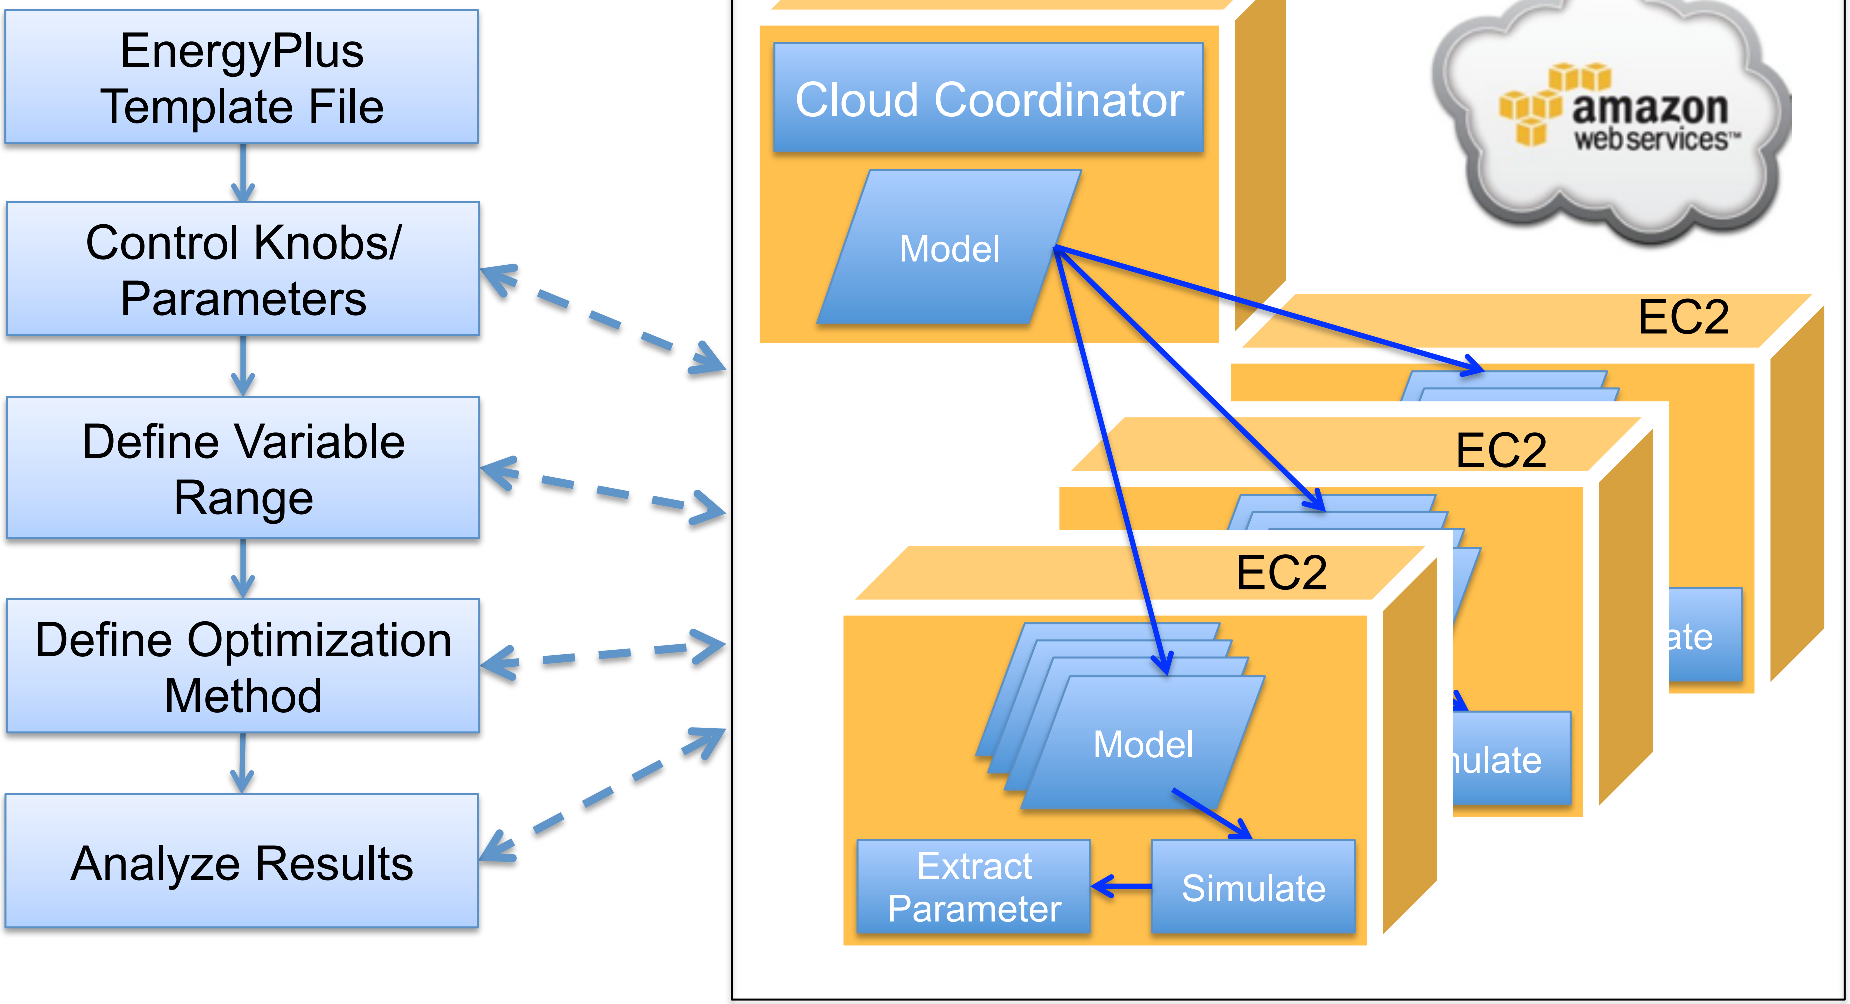

MLE+ Coordinator (Local)

- EnergyPlus Template File
- Control Knobs/ Parameters
- Define Variable Range
- Define Optimization Method
- Analyze Results

Master

Cloud Coordinator

Model

EC2

EC2

EC2

Model

Extract Parameter

Simulate

Cyber-Physical Systems

- **Intersection of Computation, Controls & Communication**
  - Safety-critical and Life-critical systems

- **Tightly coupled with (Messy) Physical Plants**
  - Interesting Domains: Medical, Energy, Automotive…

- **New Interesting Problems involving:**
  - Scheduling and Control
  - From Verified Models to Verified Code of Closed-loop Systems

Research Focus

Networked Cyber-Physical Systems

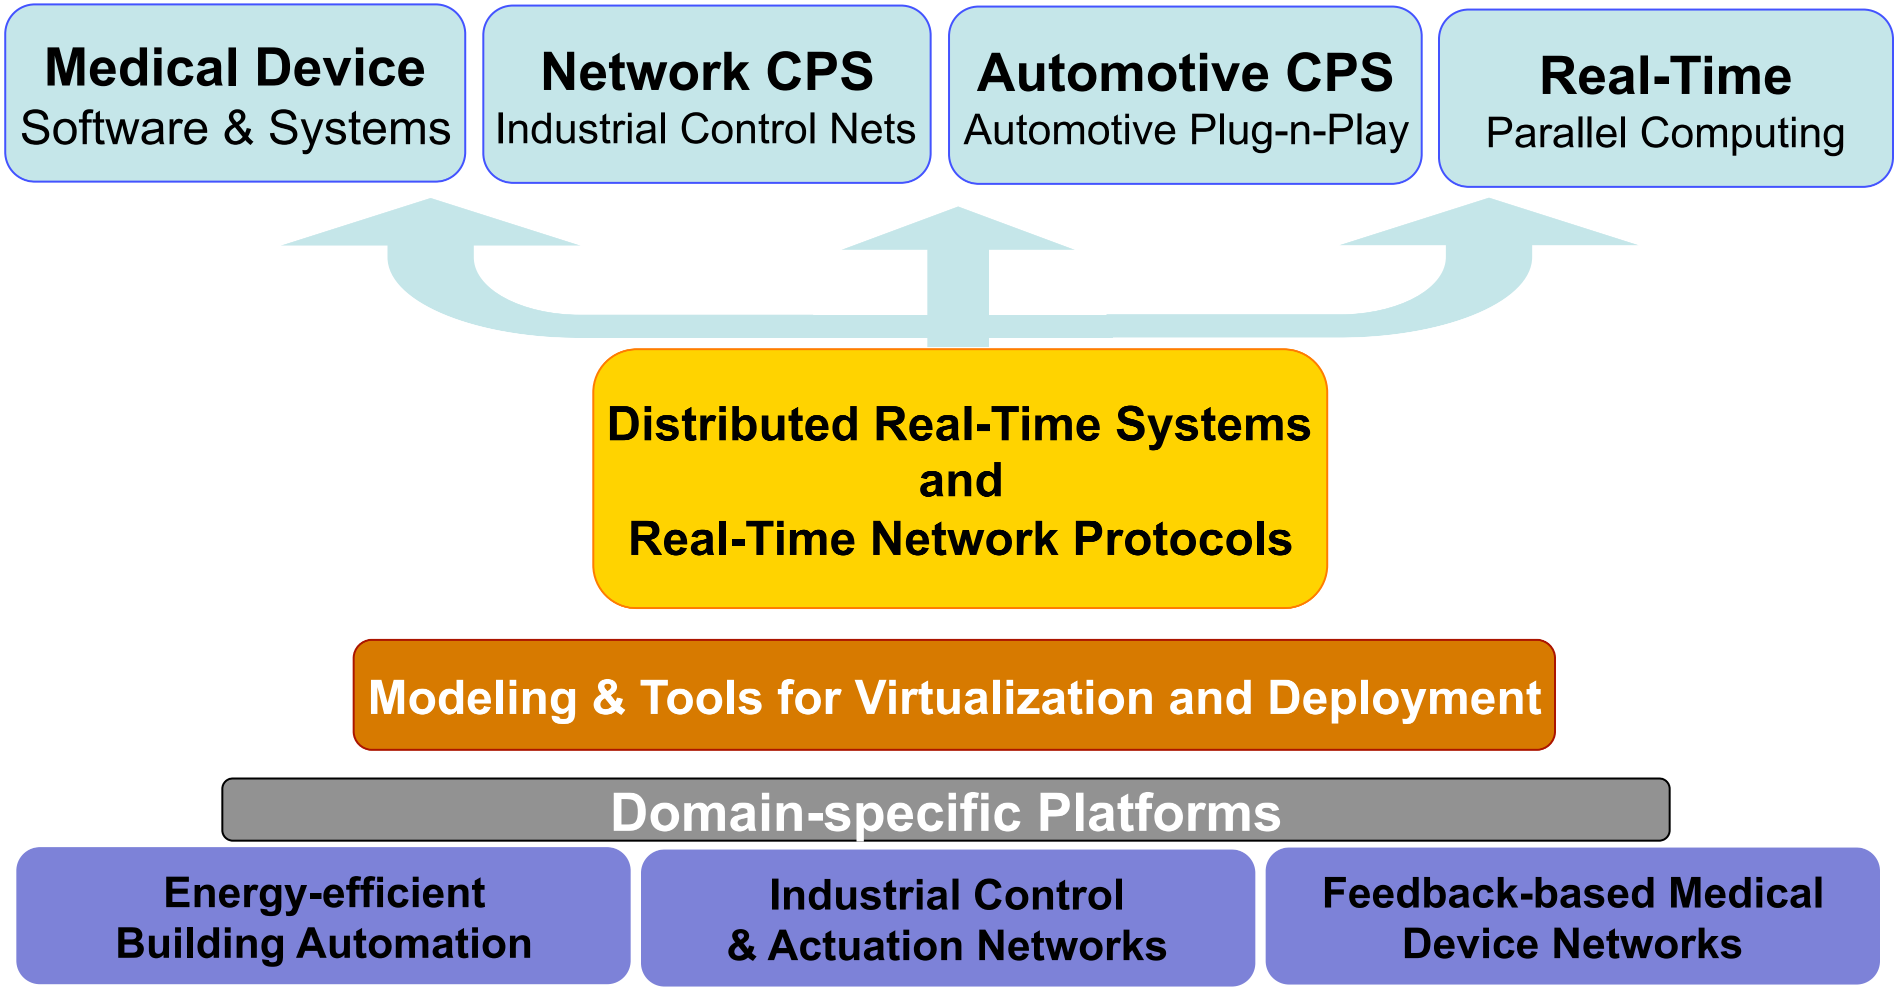

Medical Device
Software & Systems

Network CPS
Industrial Control Nets

Automotive CPS
Automotive Plug-n-Play

Real-Time
Parallel Computing

Distributed Real-Time Systems
and
Real-Time Network Protocols

Modeling & Tools for Virtualization and Deployment

Domain-specific Platforms

Energy-efficient
Building Automation

Industrial Control
& Actuation Networks

Feedback-based Medical
Device Networks

Penn Engineering

mLAB Real-Time & Embedded