Toward Smarter Schedules for Commercial Buildings

Omid Ardakanian, Arka Bhattacharya, David Culler

University of Waterloo
September 2016

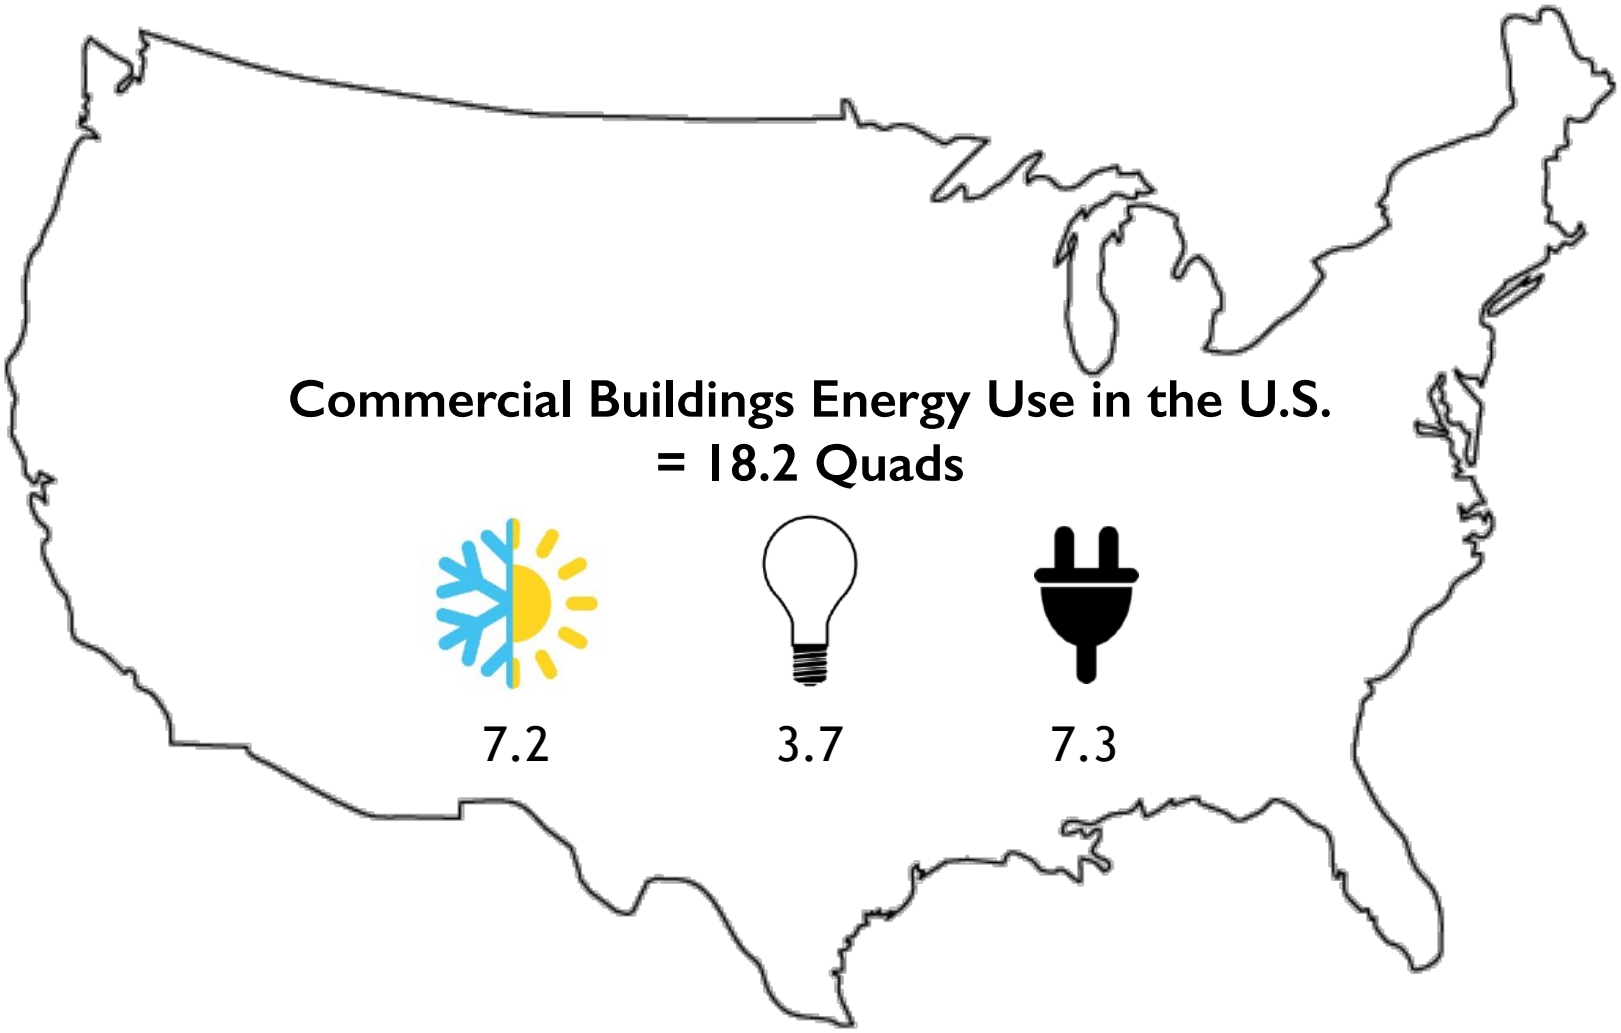

1 Quad of energy ≈ $20 Billion

10% Saving in HVAC Energy Consumption ≈ $14.4 Billion per year

Source: Buildings Energy Data Book (2011), Table 3.1.1

Motivation

- The HVAC systems traditionally run on a single static schedule
  - A lot of energy is wasted in conditioning empty/partially-occupied spaces
  - Occupancy sensors are not commonly available in all zones

- Retrofitting a large commercial building with occupancy sensors is intrusive, overly costly, and error prone given the scale

- Potential for applying smarter (dynamic and customized per-zone) schedules derived from non-intrusive occupancy sensing which relies on coarse-grained measurements available through the building management system
  - Can we trade off occupant comfort for energy efficiency?

Commercial Buildings' HVAC System

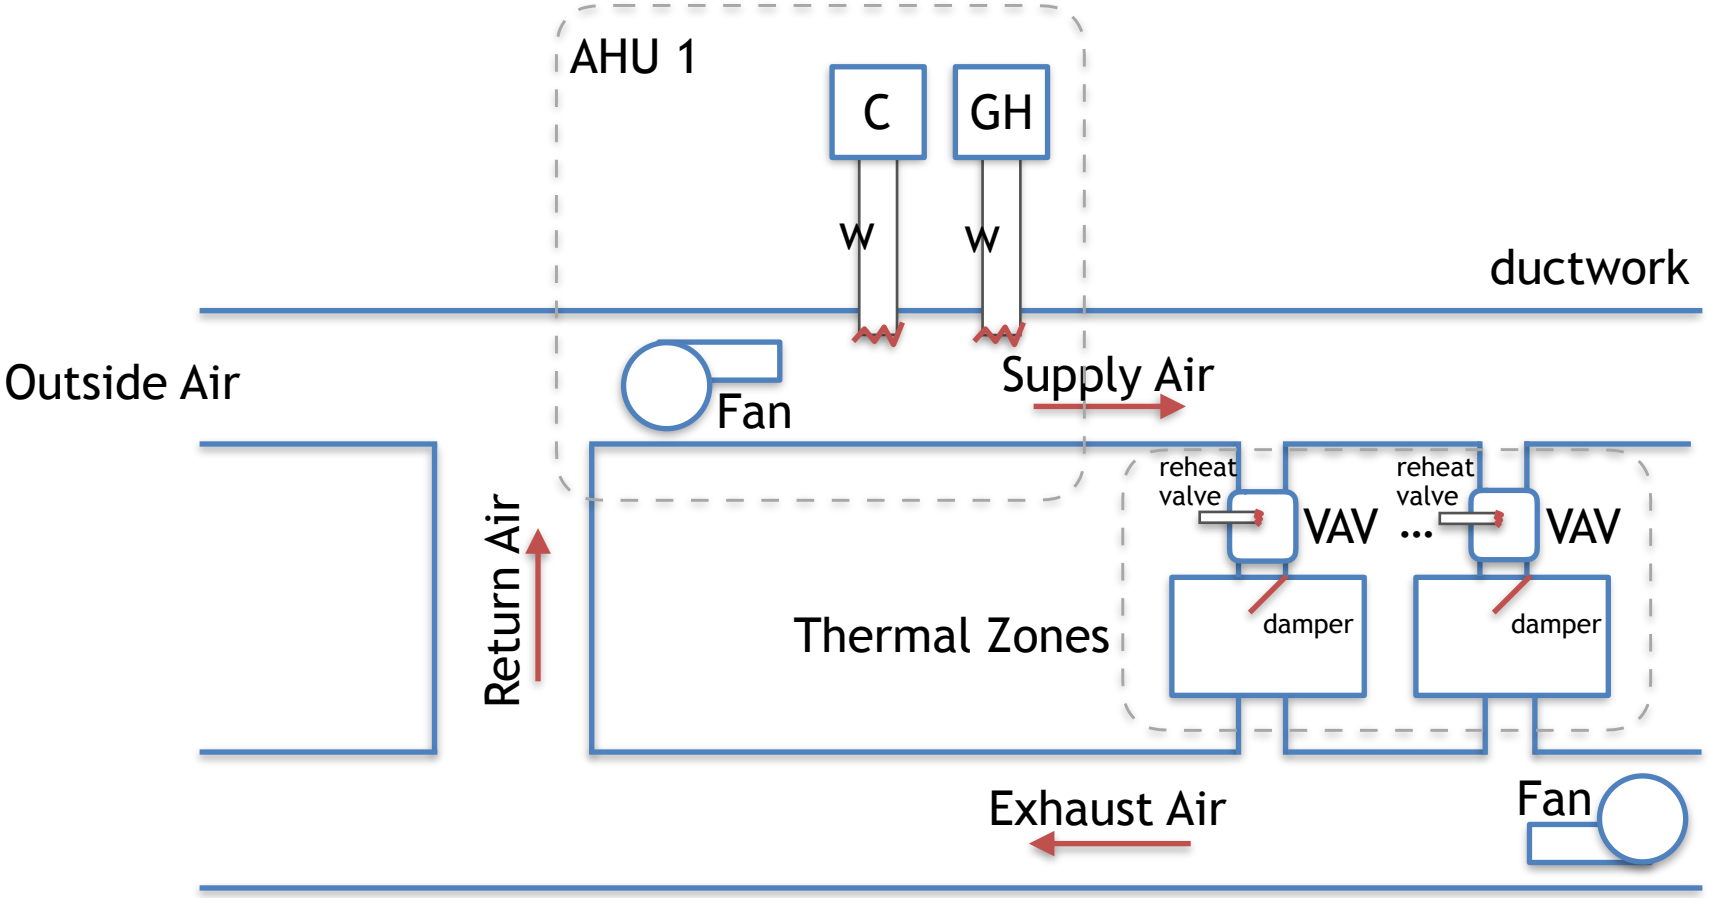

Commercial Buildings' HVAC System

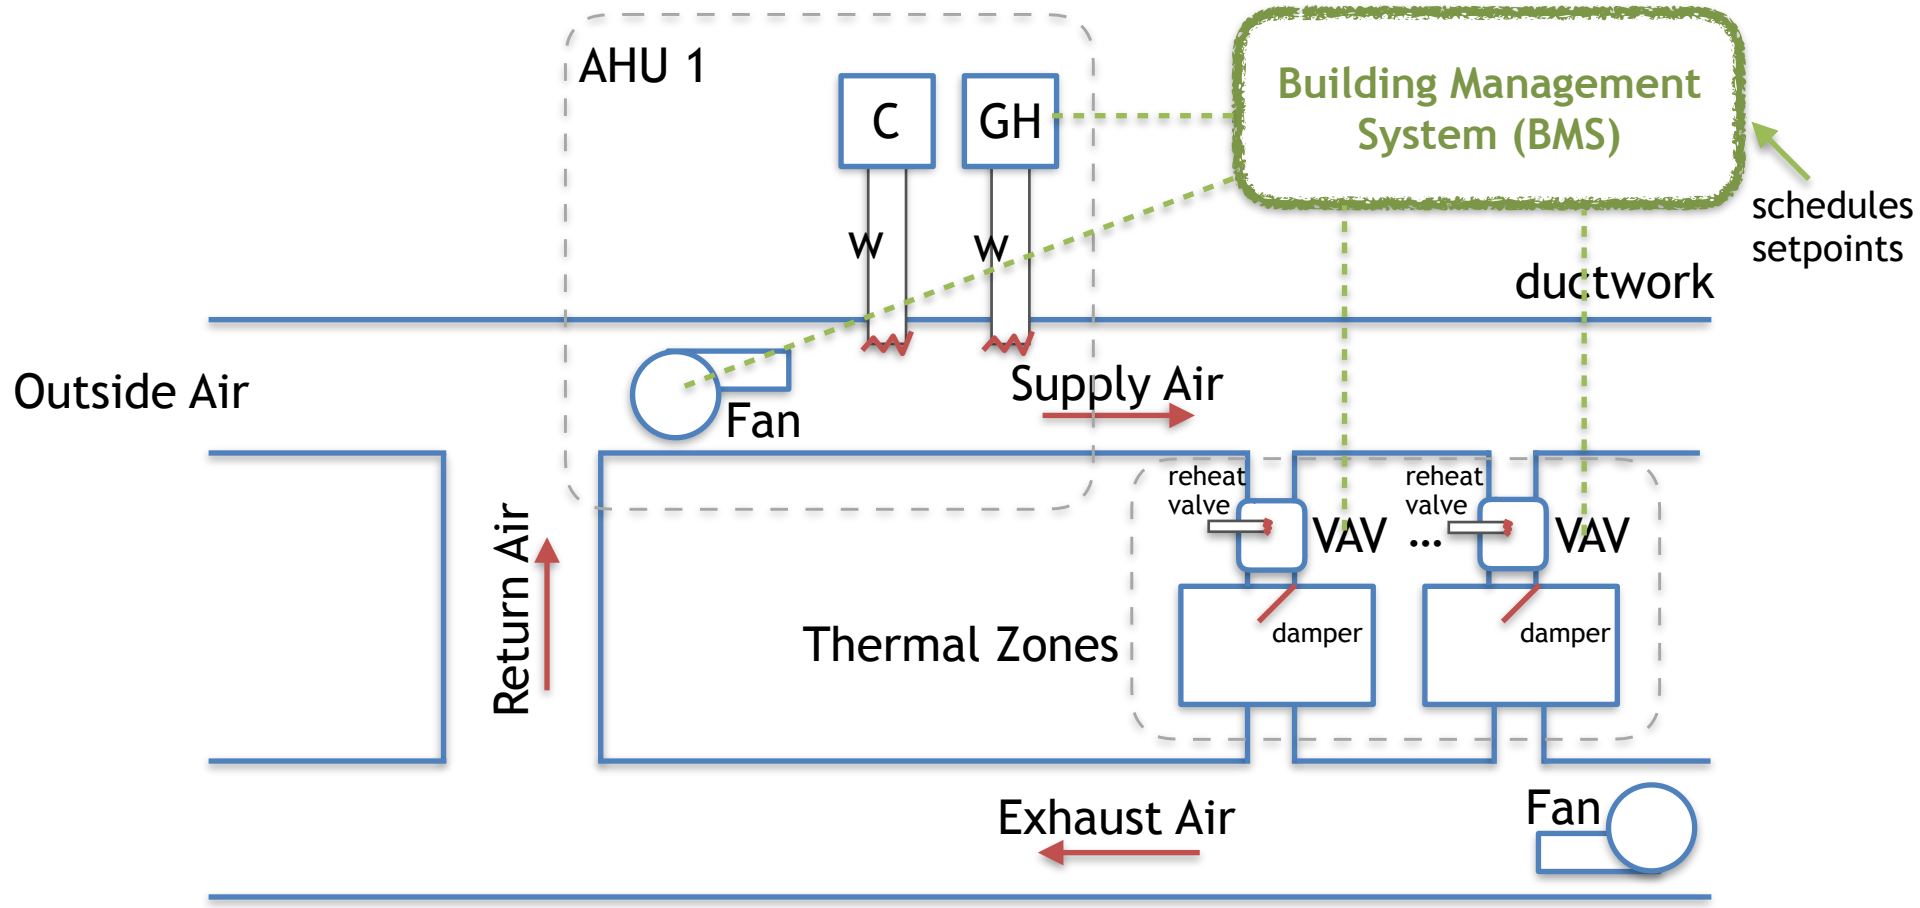

BMS controls indoor climate (i.e. zone temperature and air flow) according to ASHRAE standards using nested control loops:

- AHU controls: supply air velocity and temperature
- VAV controls: damper position and reheat valve position

The Basic Idea

- Reheating air at terminal zones is bad
  - Why do you need to reheat?
    - AHU supplies air at 57F which is too cold for occupants
  - Why does AHU supply air at 57F?
    - To account for hottest zones
    - To account for losses through ductwork
- This leads to avoidable energy loss
  - Reheating the supplied air to the zones must be limited to when they are occupied

Major Heat Sources in a Zone

- Heat gain from the presence of occupants
- Heat gain from (possibly) periodic external sources such as solar irradiance
- Heat gain from internal sources such as servers, displays, and incandescent lights
- Conduction heat transfer through the external building envelope as well as walls, windows, and open doors

Several confounding effects contribute to the change of heat load in a zone. The VAV system responds to this.

VAV responds to changes in the occupancy state of a zone if occupants inject enough heat into it

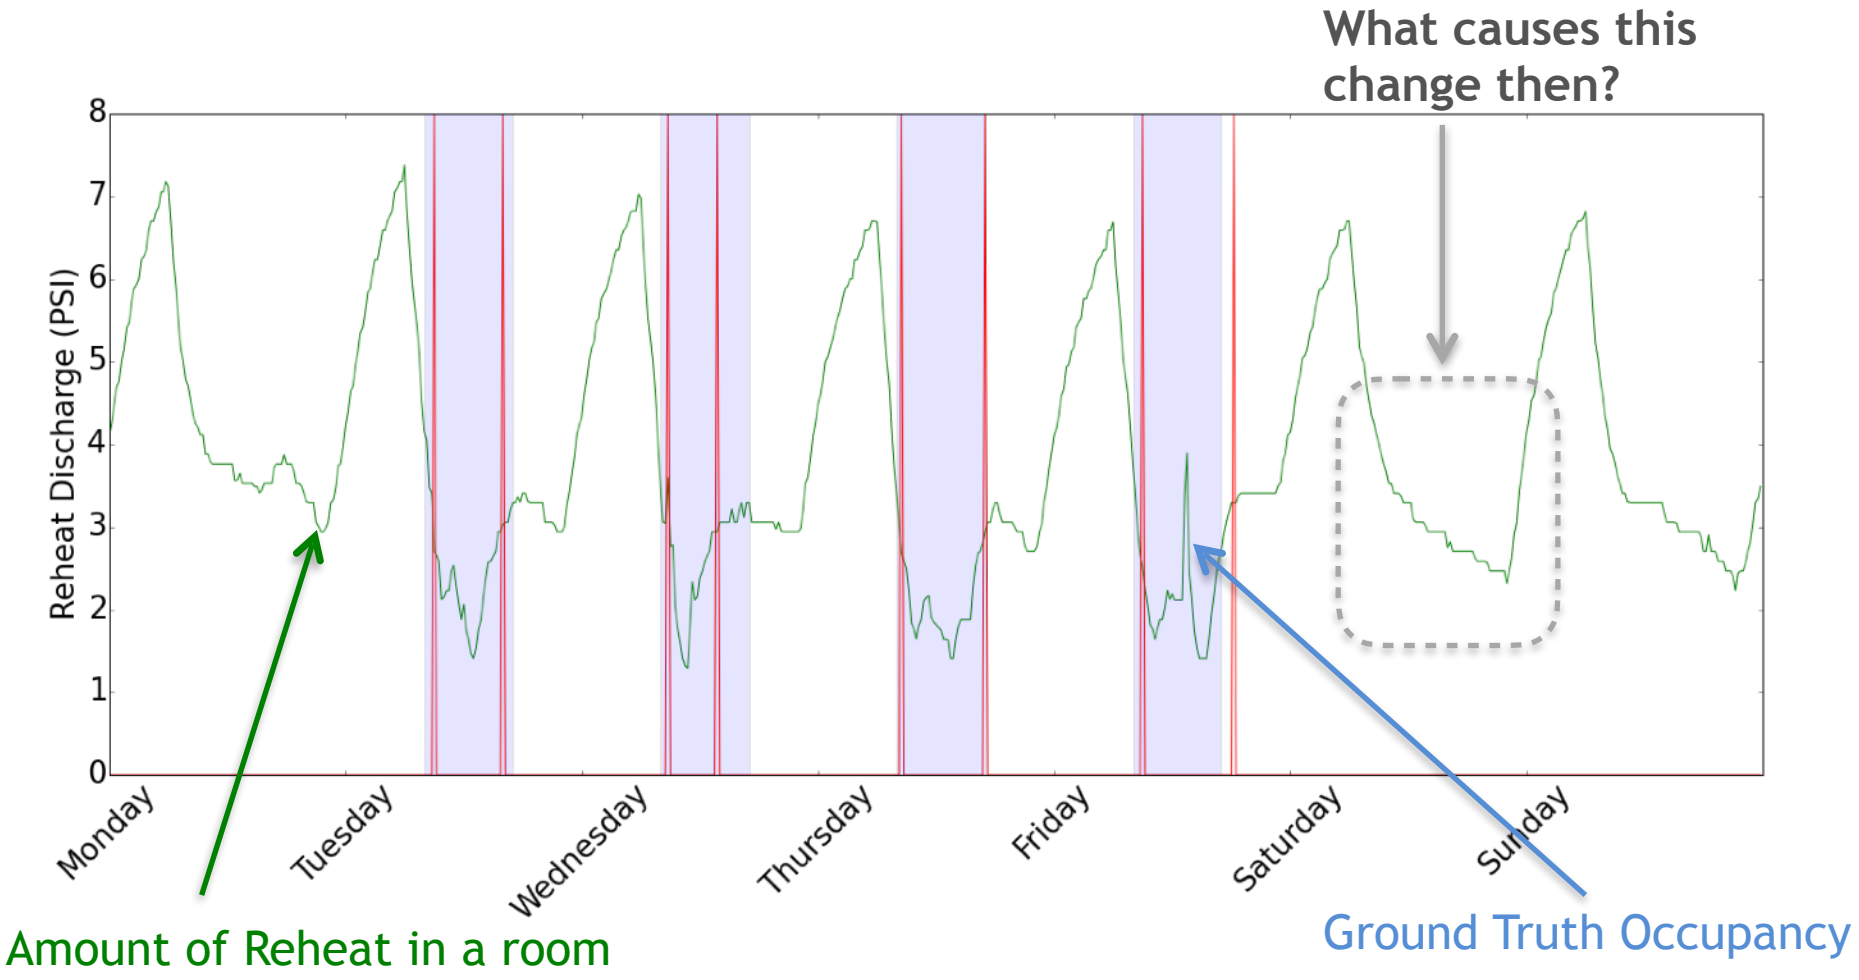

There is always one sensor (*occupancy indicative sensor*) which picks up this response

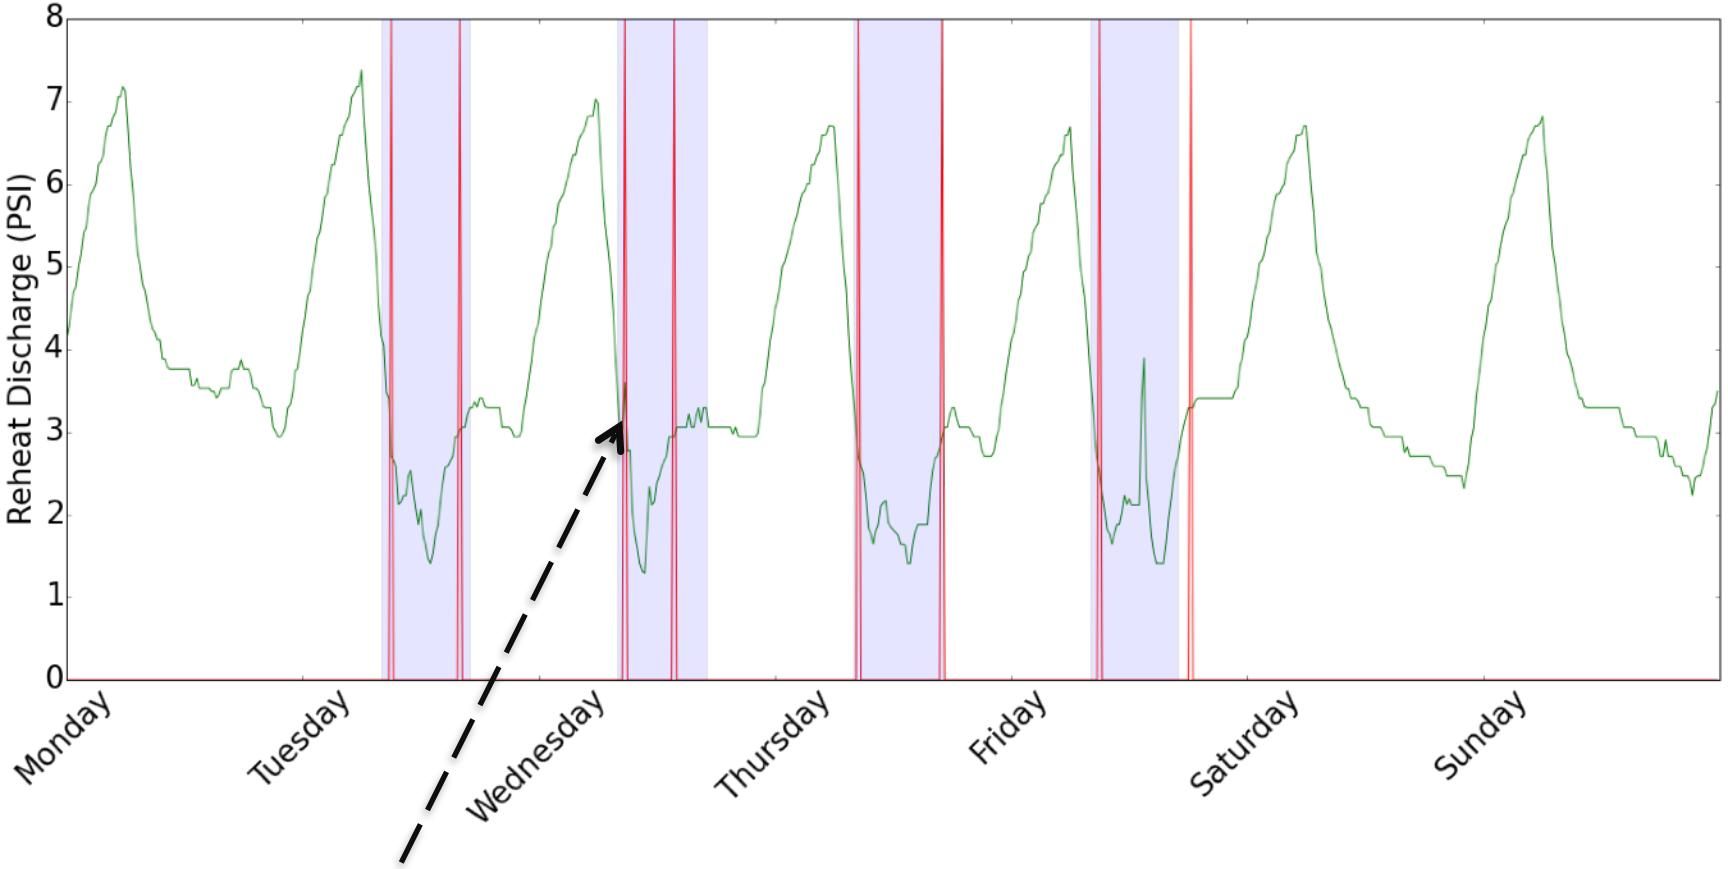

Reheat valve or damper reacts shortly after an occupant enters the room

Our Approach

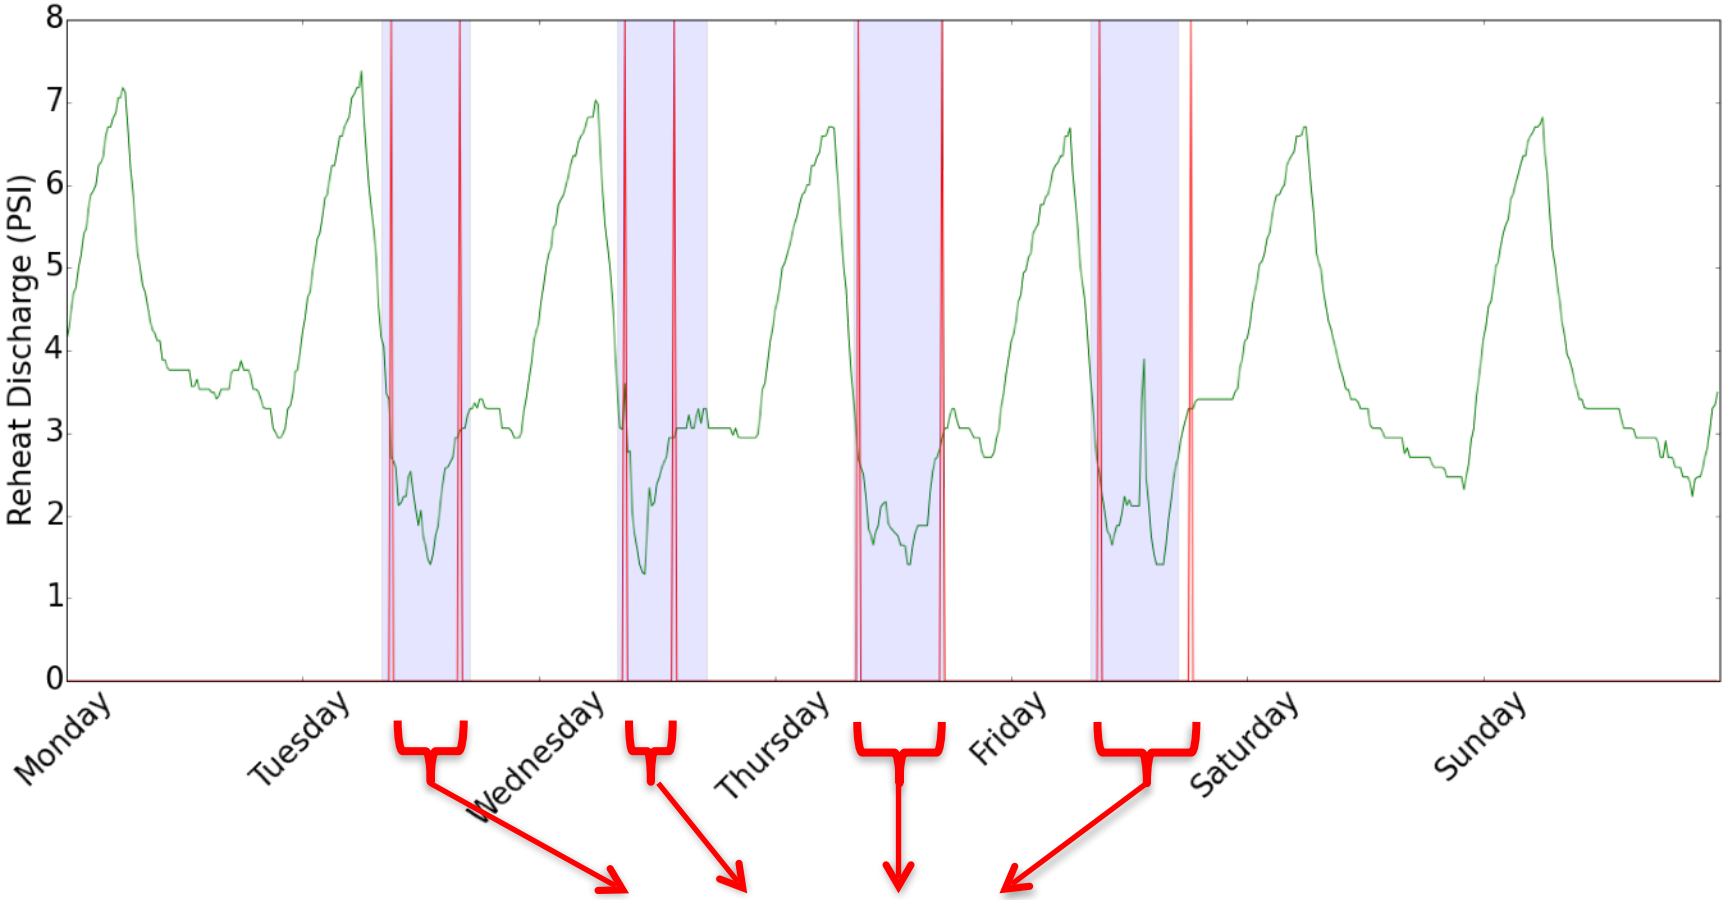

Systematically determine those edges (Canny Edge Detector) and apply schedules that take the inferred occupancy into account

Some other salient points

- We do not want to know the exact minute the occupant(s) came in

- Data is our friend
  - Assuming zone occupancy is stationary, we can get a good estimate of occupancy using statistical techniques if we collect enough weeks worth of data and pick up just enough of the edges

Testbed

- Three large UC Berkeley campus buildings with Building Management Systems installed by different vendors, containing 117, 109 and 270 zones respectively

- The occupancy indicative sensors used for the three buildings are
  - the pneumatic control sensors in Building 1
  - the air flow sensors in Building 2
  - the reheat sensors in Building 3

- Building 2 implemented a setback strategy which allows the temperature to drift when the zones are presumably empty. The nighttime setbacks were in effect every day from 7pm–5am

Time Series Analysis Techniques

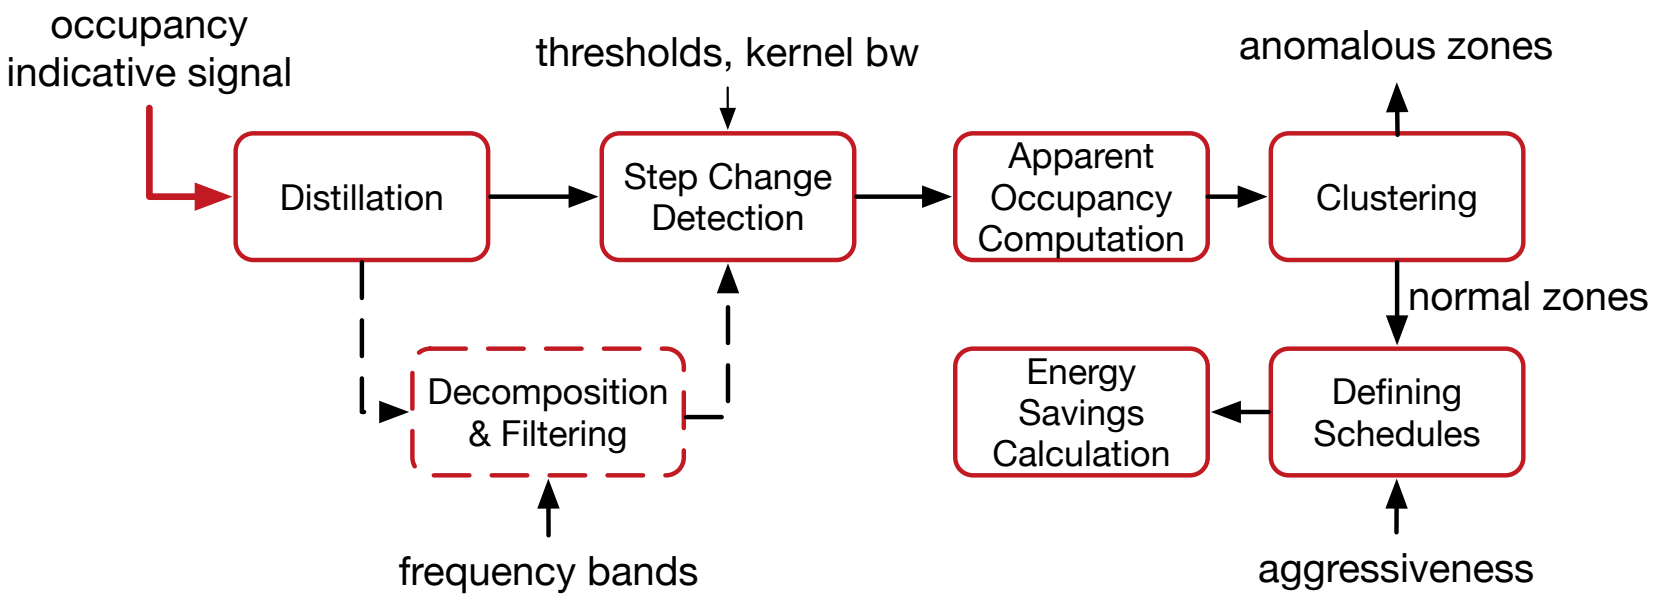

Empirical Mode Decomposition

- A non-stationary signal can be decomposed into a small number of non-uniform oscillatory components termed *intrinsic mode functions (IMFs)*

- *The IMFs admit Hilbert transform; hence, the notion of instantaneous frequency can be defined*

- *Certain IMFs can be extracted from the signal*

- *This technique is used to remove noise, along with diurnal and seasonal effects from the occupancy indicative signal*

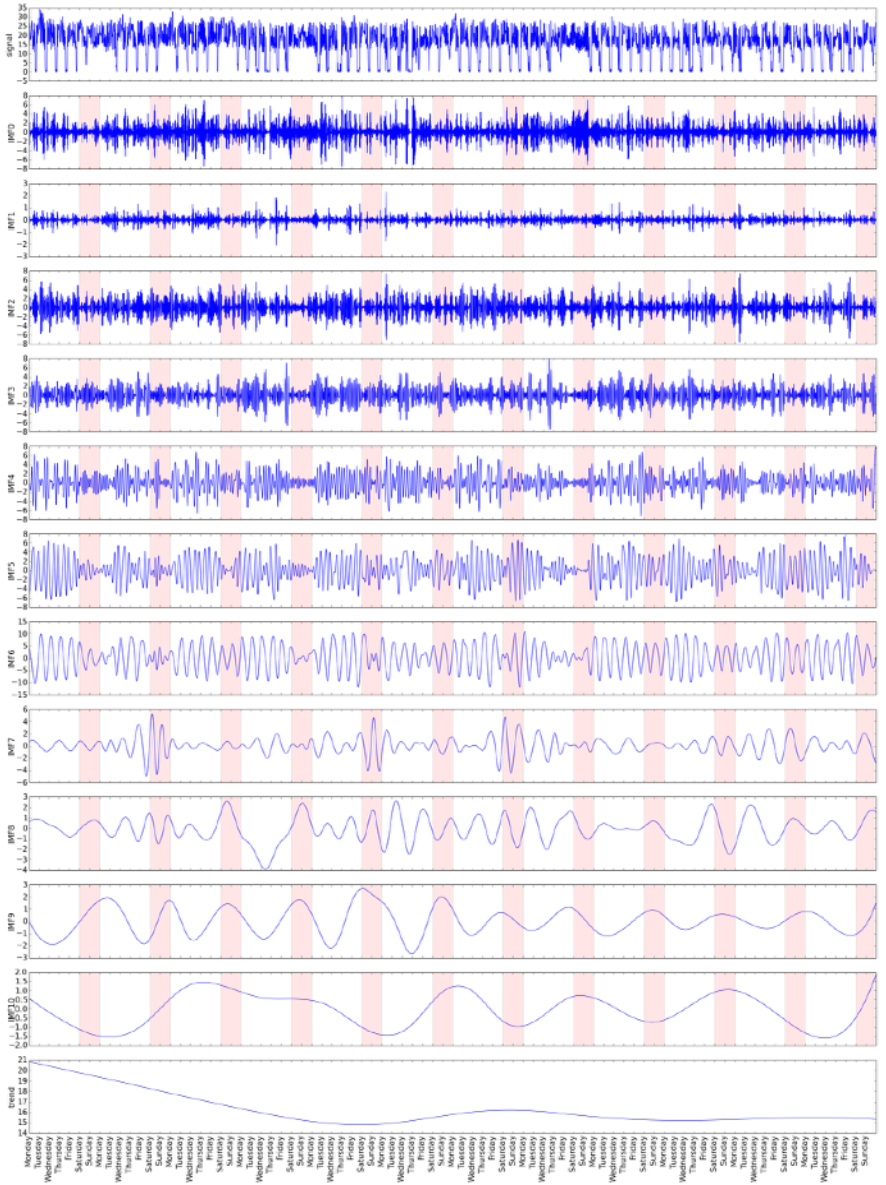

Canny Edge Detection

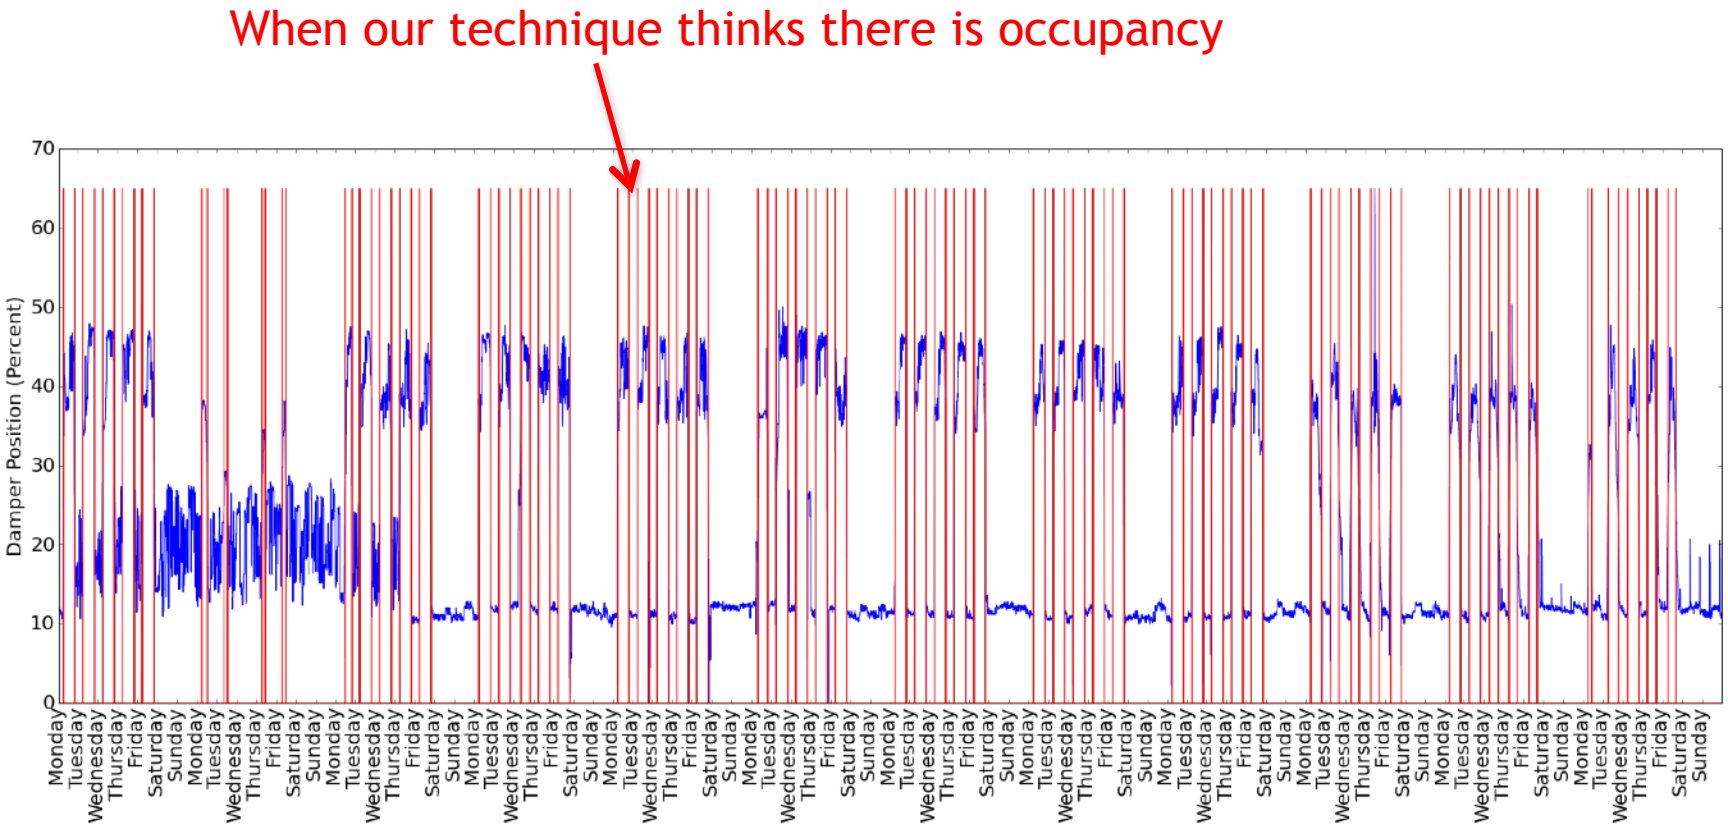

Compute convolution of the signal with the first derivative of a Gaussian kernel
local maxima of the result are upward step-edges and its local minima are downward step edges

Determining Apparent Occupancy of a Zone

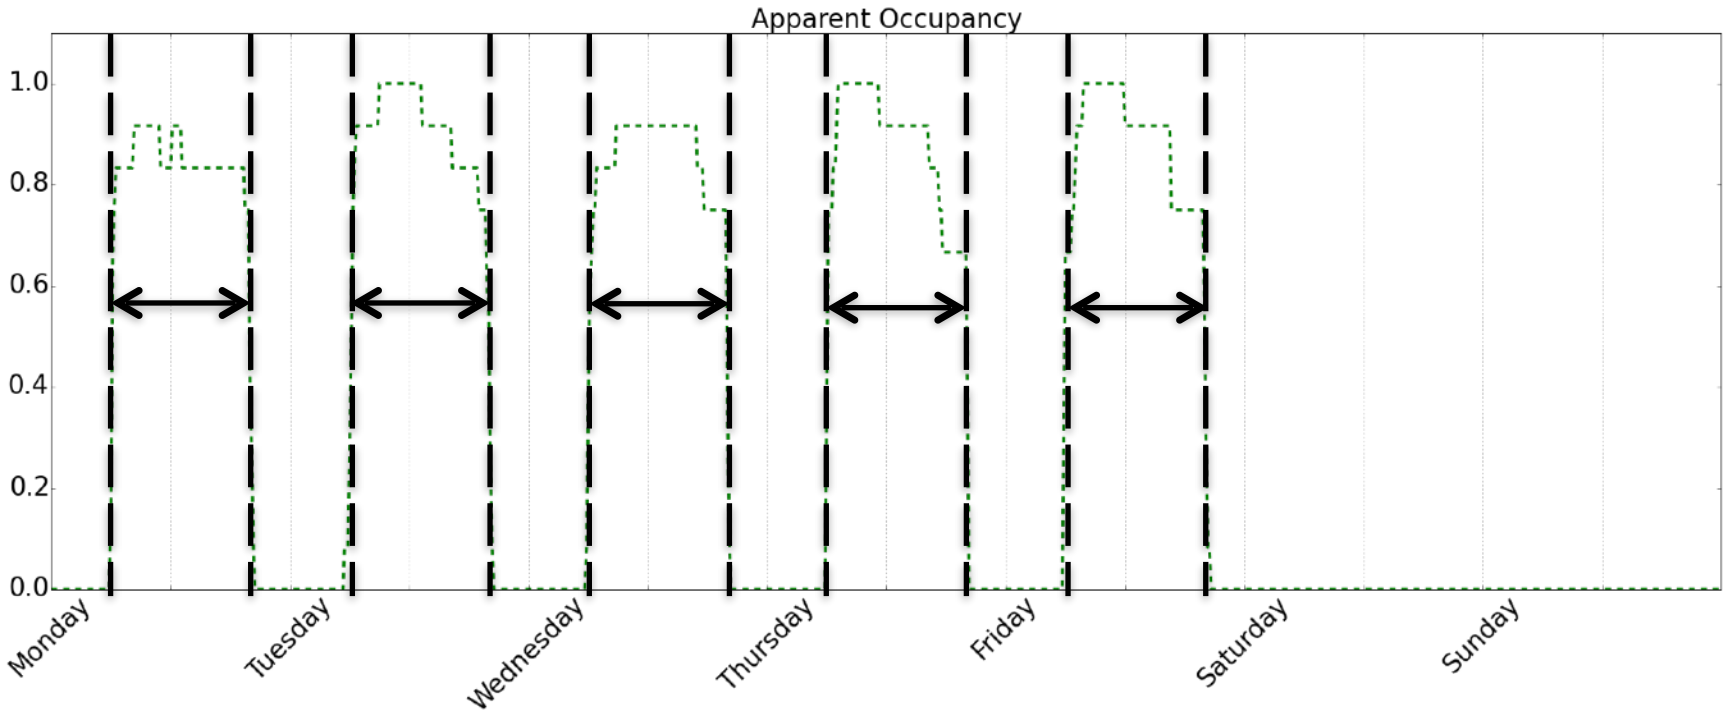

Compute weekly average of the inferred occupancy state of a zone

More Aggressive Schedules

schedule start: *x*th percentile of the start times
schedule end: (100-*x*)th percentile of the end times

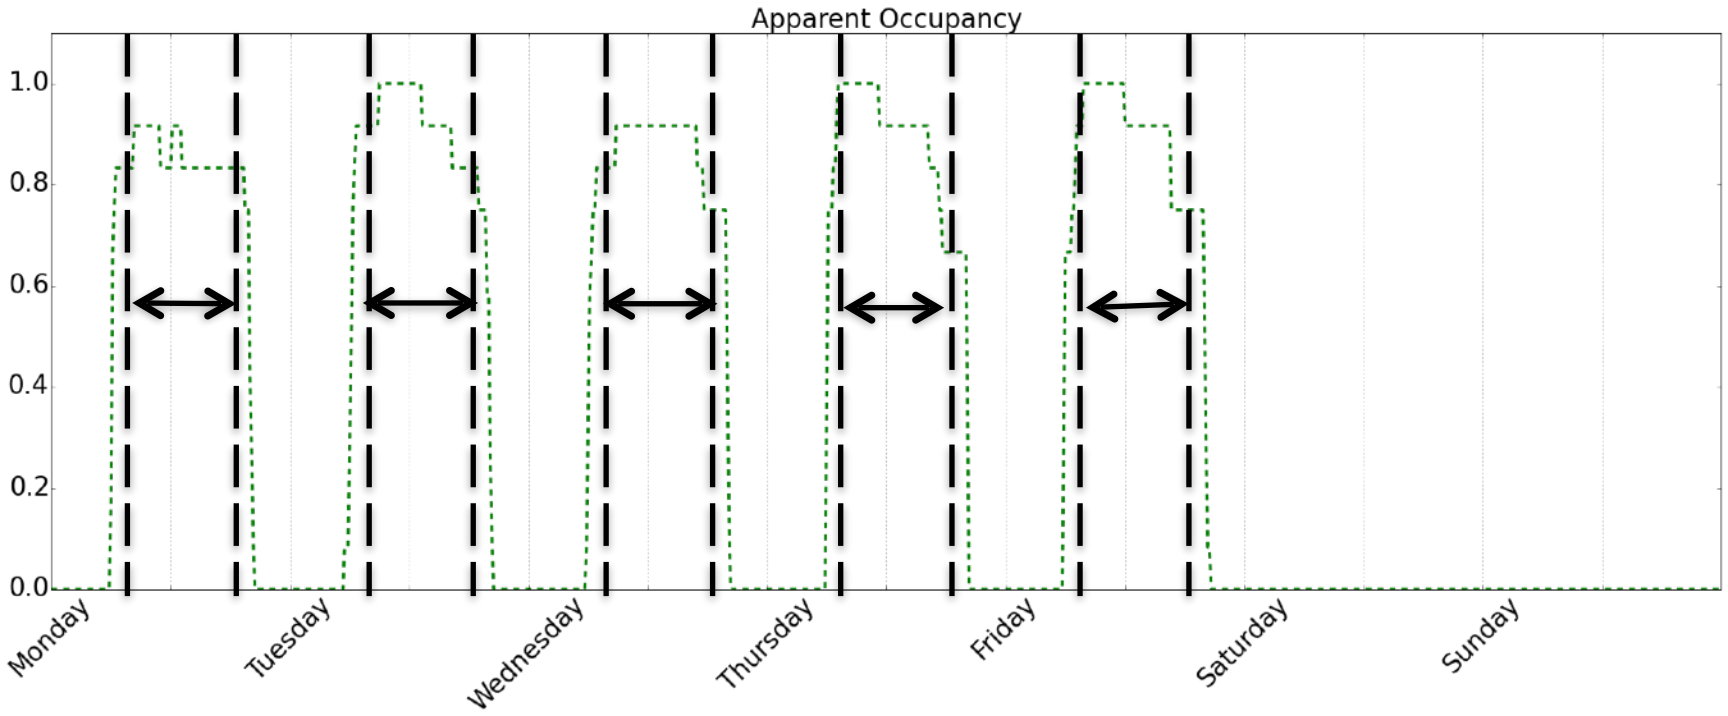

Tradeoff:
Less time the HVAC system is operating, hence more energy savings
However, there are times when an occupant comes in and finds their room unconditioned

How do we save energy?

- Determine schedules for each zone by applying the proposed non-intrusive technique

- Program VAVs so that they do not reheat supply air outside those schedules
  - Saving the nefarious "reheat" energy.

Validation

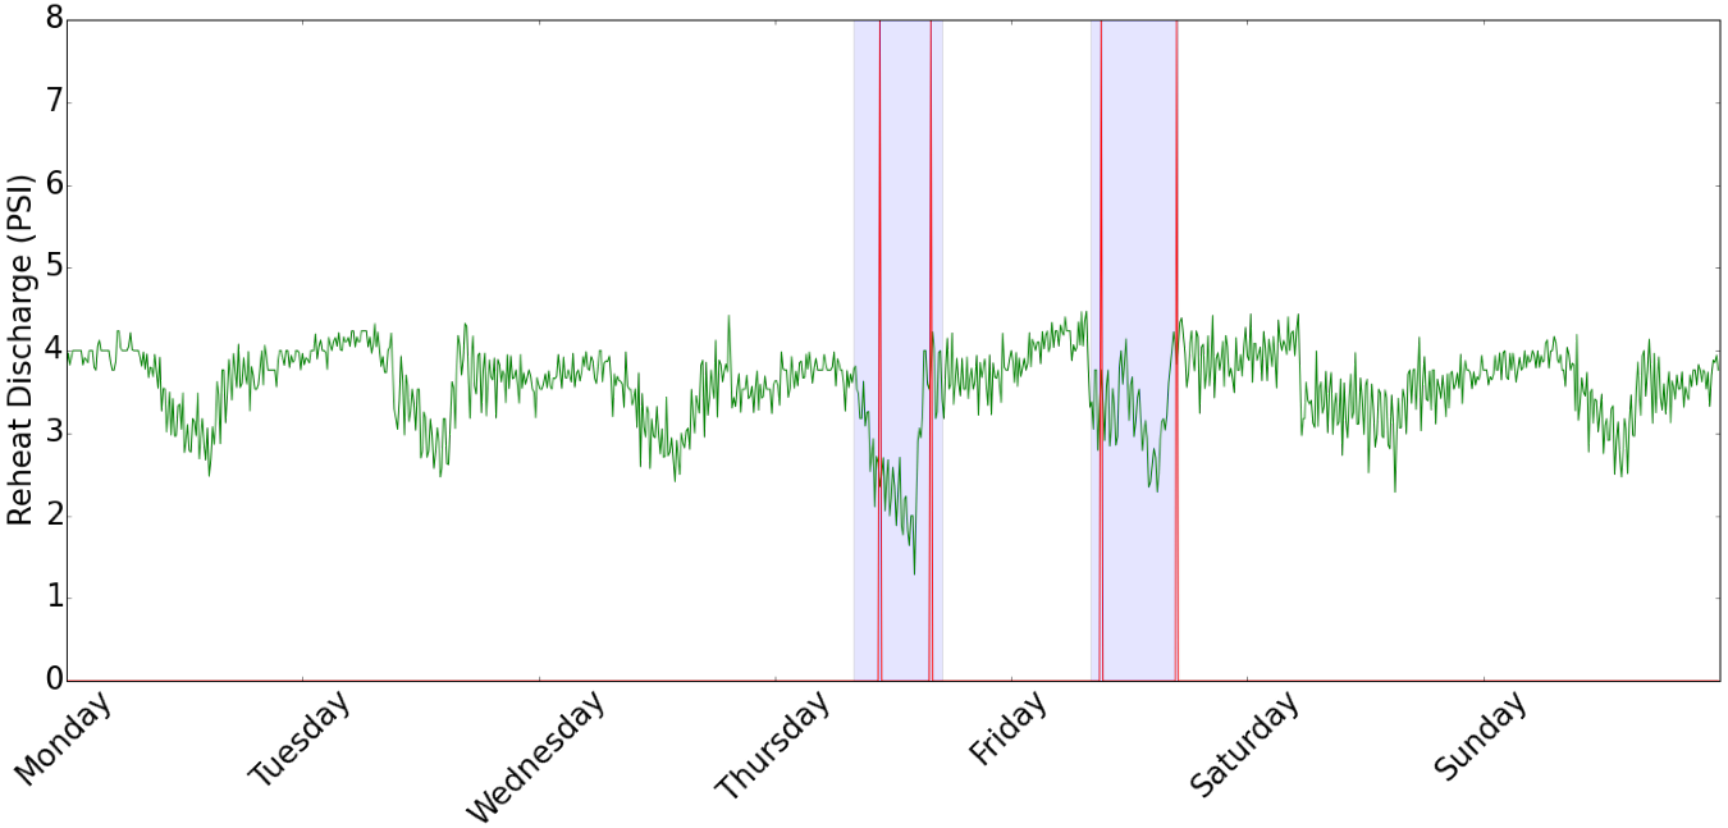

Limited ground truth data:

a) Manually logged occupancy hours of 7 shared and private offices in our testbed

b) Extracted occupancy hours from video recordings (a security camera installed in a lab)

What does overall building occupancy look like?

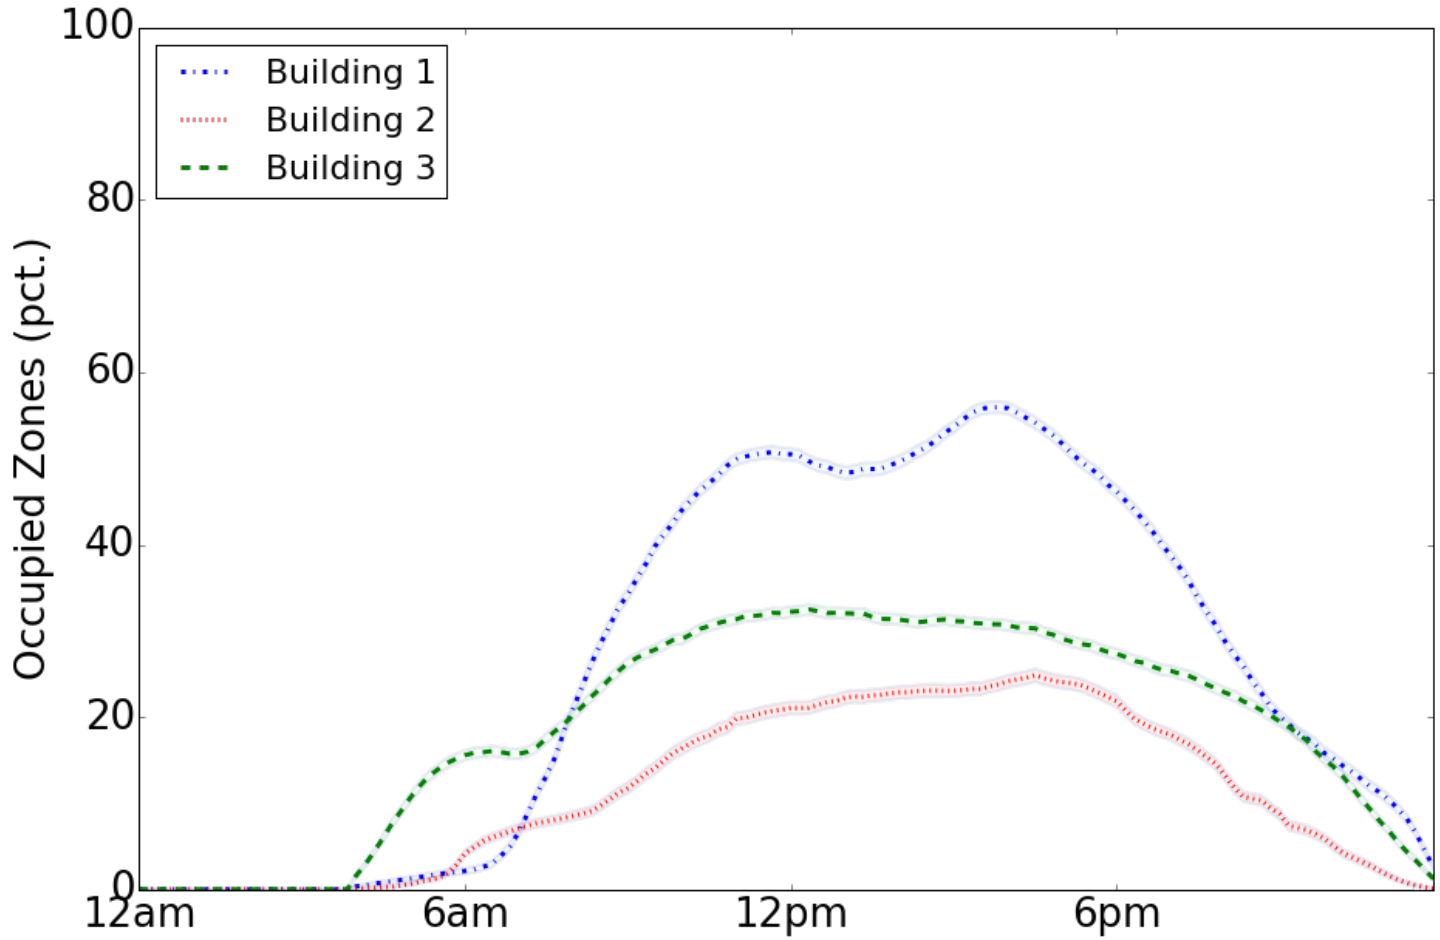

What does each of the individual zones' occupancy look like?

Occupancy profile (Weekdays)

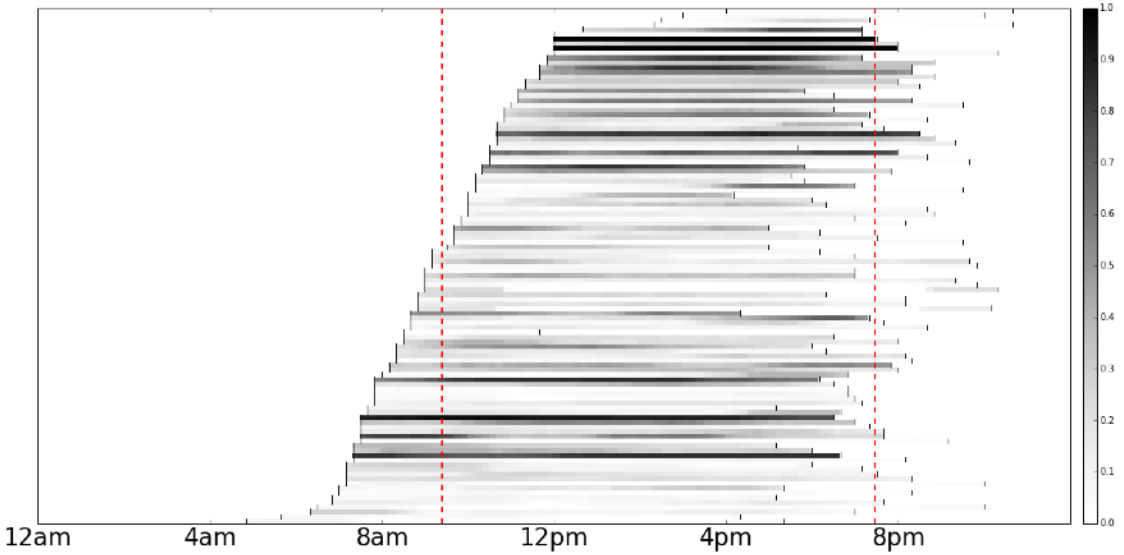

Occupancy profile (Weekends)

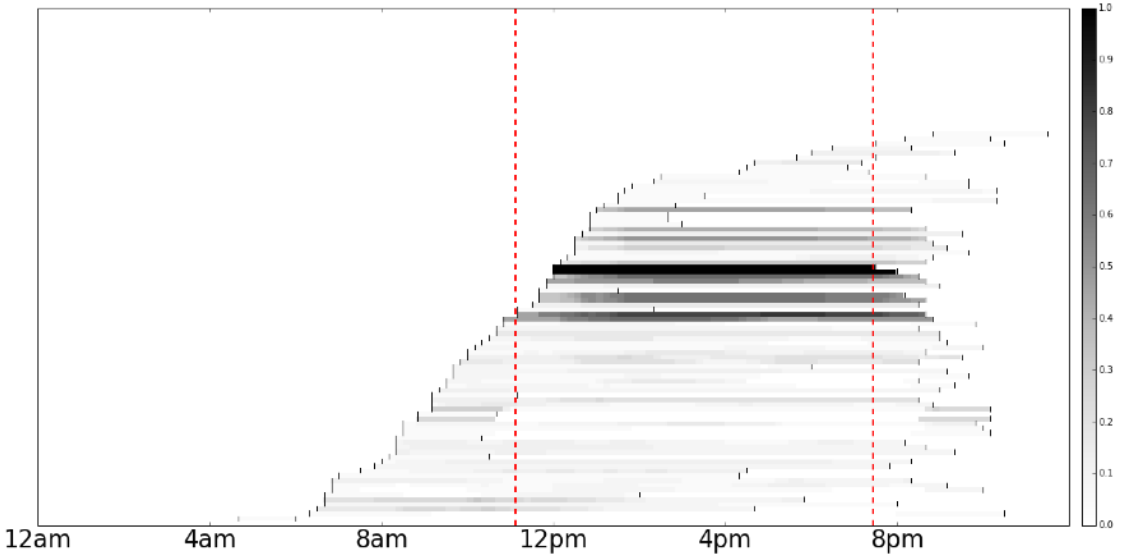

Static vs. Adaptive Schedules

- **Static Schedules** (easy to implement, high occupant comfort violation, moderate to high energy savings on reheat)
  - **Naïve:** predefined schedules for all zones based on the facilities manager's intuition
  - **Learned:** customized per-zone schedules learned over a short period of time
    - Actual zone occupancy must be stationary otherwise we have to repeat this process after a while
- **Adaptive Schedules w/ Sliding Training Window** (difficult to program into a legacy BMS, low occupant comfort violation, high energy savings on reheat)
  - **Weekly:** a customized per-zone schedule for all days of the week
  - **Per-Day:** a customized per-zone schedule for each day of the week
  - **Weekday-Weekend:** a customized per-zone schedule for weekdays and another one for weekends
  - Could achieve between 37%–76% energy savings and between 1%–8% comfort violations across the buildings

How we are actually saving reheat energy?

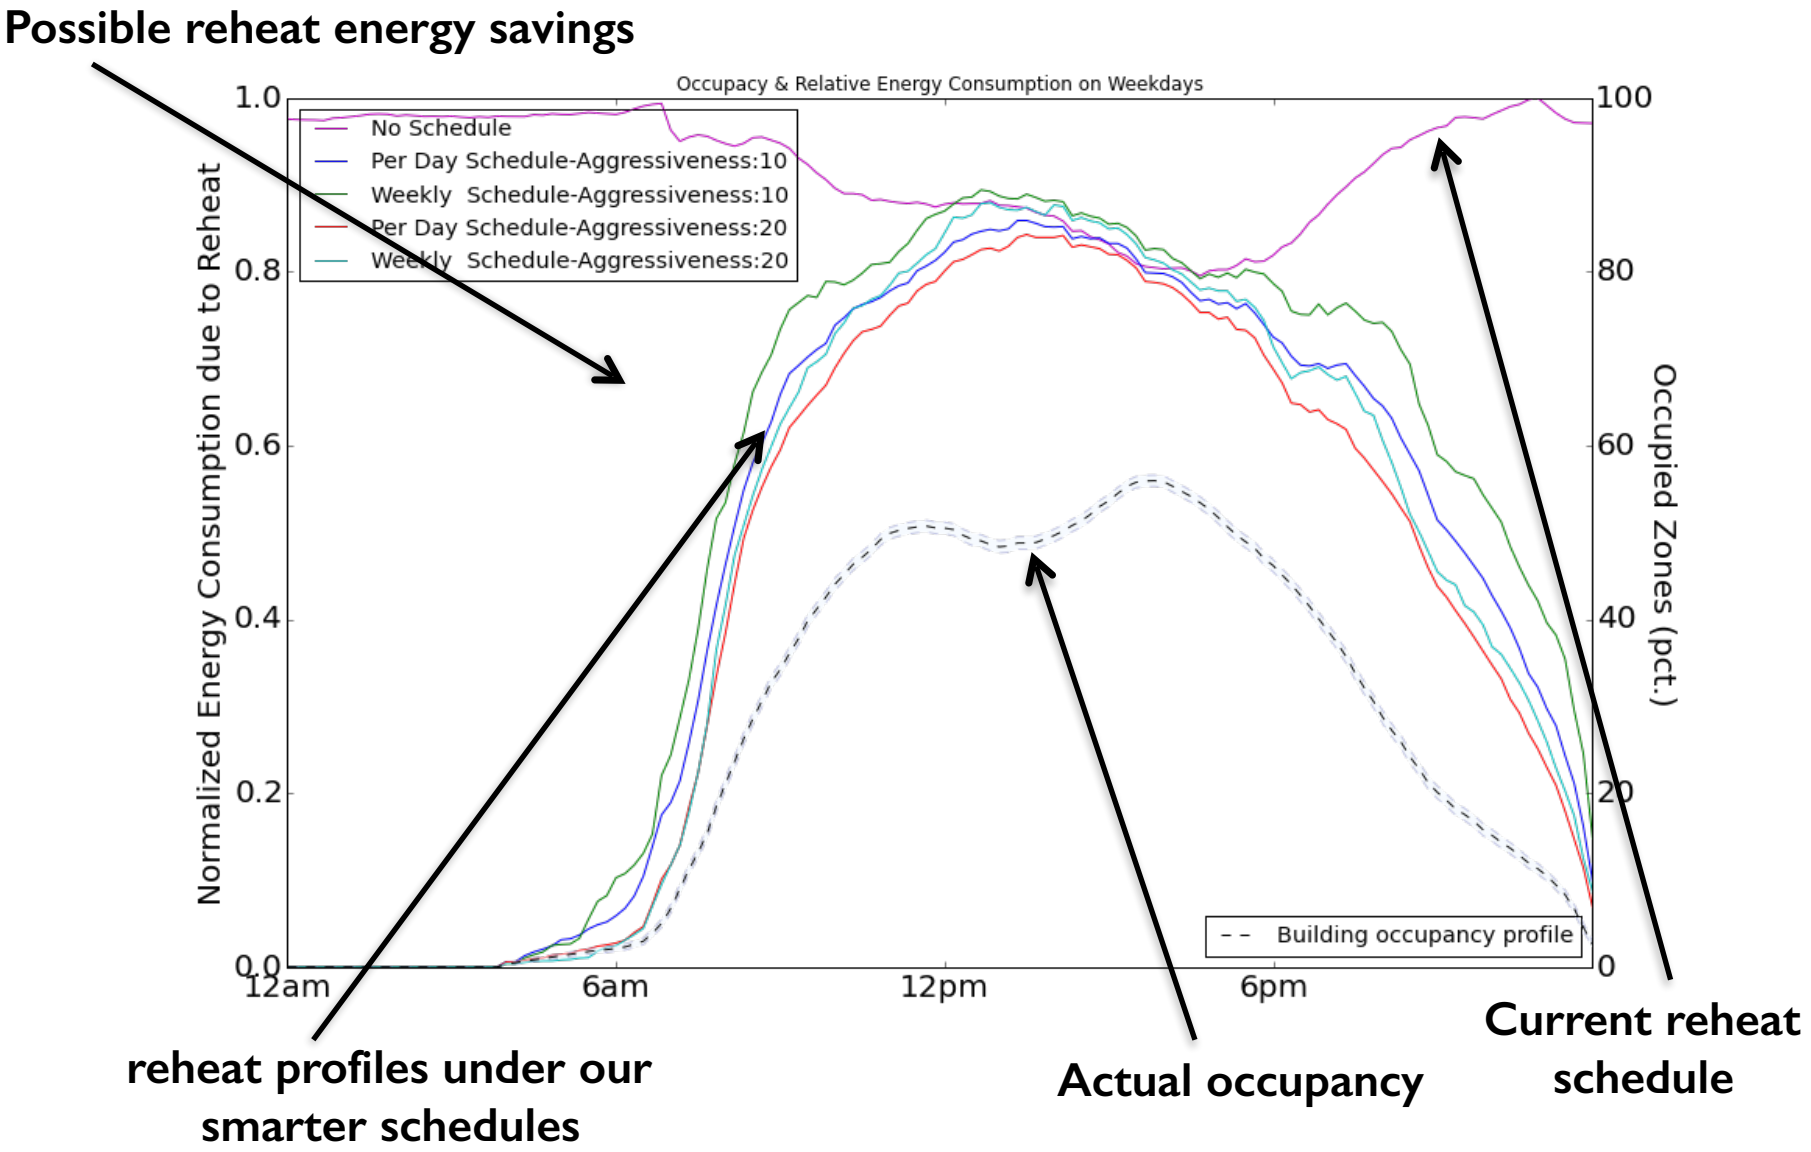

Major Takeaways

- Naïve schedules are not great in terms of occupancy comfort violation
- Learning a schedule on even a small amount of data (~2 weeks) helps save ~50% of reheat energy, with negligible (<2%) estimation errors (occupancy comfort violations)
- If the underlying occupancy patterns are static, then increasing the length of the training window has diminishing returns
  - 1 week training window < 4 week training window
  - 4 week training window ~ 8 week training window
- Having different schedules for each particular day might be an overkill. The resulting occupancy comfort violations might not be worth the sophistication
- Learning weekly schedules over a sliding window generally gives only slightly better results (again maybe not worth the trouble)

Work in Progress

- KETI motes are brought up in one of our test buildings
  - PIR, CO2, Light, …
- Detect zone-level occupancy using sensor fusion
- Compare the obtained occupancy schedules with those inferred from the VAV response
- Can we identify the number of occupants in each zone?