Lessons From An Adaptive House

Michael Mozer[+][*]
Robert Dodier[#]
Debra Miller[*]
Marc Anderson[*]

Josh Anderson[+]
Dan Bertini[#]
Matt Bronder[*]
Michael Colagrosso[*]
Robert Cruickshank[#]
Brian Daugherty[*]
Mark Fontenot[∈]
Okechukwu Ikeako[+]
Paul Kooros[+]

Diane Lukianow[+]
Tom Moyer[=]
Charles Myers[+]
Tom Pennell[*]
James Ries[+]
Erik Skorpen[+]
Joel Sloss[+]
Lucky Vidmar[*]
Matthew Weeks[+]

University of Colorado

[*]Department of Computer Science
[+]Institute of Cognitive Science
[#]Department of Civil, Environmental, and Architectural Engineering
[+]Department of Electrical and Computer Engineering
[=]Department of Mechanical Engineering
[∈]Department of Aerospace Engineering

http://www.cs.colorado.edu/~mozer/adaptive-house

Home automation

Homes might be programmed to

- close drapes at night
- turn down sterero volume when phone rings
- flash porch lights if baby is crying
- offer recipes to go with the ingredients in the cupboard

Vision of the future

...Imagine that the owner of a new home does not plan on using his lower level much at night. He can have a technician at the central station program his system so that the temperature is lowered to 60° at 10 p.m. But later, a home theater is installed in the basement, and many late weekend evenings are spent watching DVDs. The owner can simply call the technician and request that the program be changed so that the lower level remains a comfortable 72° on Friday and Saturday nights. (*Electronic House*)

André Malepart of Honeywell visits Hank and Darlene Johnson regularly to "fine tune" their home automation system.

The failure of home automation

“Pressing situations”, *Electronic House*, February 2005

> One of the best things about owning a lighting control system is being able to turn every light off by pressing one button. But as convenient as this is, it can get you into a lot of trouble. Take Gary Cox of Boise, ID. For some reason he kept pressing the ALL OFF button on his LiteTouch lighting system when he didn’t mean to...Tired of being left in the dark, he called the firm that had installed the system to help him kick the habit. Their solution was to make it so that Gary had to hit the ALL OFF button twice to enact the function.

The failure of home automation

“Pressing situations”, *Electronic House*, February 2005

> One of the best things about owning a lighting control system is being able to turn every light off by pressing one button. But as convenient as this is, it can get you into a lot of trouble. Take Gary Cox of Boise, ID. For some reason he kept pressing the ALL OFF button on his LiteTouch lighting system when he didn’t mean to...Tired of being left in the dark, he called the firm that had installed the system to help him kick the habit. Their solution was to make it so that Gary had to hit the ALL OFF button twice to enact the function.

“Oh, shut up, house!”, *SF Chronicle*, November 2004

> The Microsoft Home...can turn the dishwasher on, but it still won’t fill it with dirty dishes or empty out the clean ones. It can tell you which sweater goes with which pair of pants, but it won’t hang the pants up for you. In other words, in its current incarnation, the smart house is more nag than household helper.
>
> ...All this takes programming—something that may be simple enough for the engineers who put together the Microsoft Home but is no such thing for those of us who have been stymied by today’s “smart” electronics (the programmable thermostat comes to mind) that come with every known option but an on-off switch.

The failure of home automation

“Pressing situations”, *Electronic House*, February 2005

> One of the best things about owning a lighting control system is being able to turn every light off by pressing one button. But as convenient as this is, it can get you into a lot of trouble. Take Gary Cox of Boise, ID. For some reason he kept pressing the ALL OFF button on his LiteTouch lighting system when he didn’t mean to...Tired of being left in the dark, he called the firm that had installed the system to help him kick the habit. Their solution was to make it so that Gary had to hit the ALL OFF button twice to enact the function.

“Oh, shut up, house!”, *SF Chronicle*, November 2004

> The Microsoft Home...can turn the dishwasher on, but it still won’t fill it with dirty dishes or empty out the clean ones. It can tell you which sweater goes with which pair of pants, but it won’t hang the pants up for you. In other words, in its current incarnation, the smart house is more nag than household helper.
>
> ...All this takes programming—something that may be simple enough for the engineers who put together the Microsoft Home but is no such thing for those of us who have been stymied by today’s “smart” electronics (the programmable thermostat comes to mind) that come with every known option but an on-off switch.

Interview with Martha Stewart, *Wired,* August 1998

> Q: Any thoughts on smart houses? How about having your refrigerator talk to your stereo?
> A: I don’t want my refrigerator talking to me period. I don’t want it telling me that I am low on meatballs. I do have a brain.

State of the art in lighting control

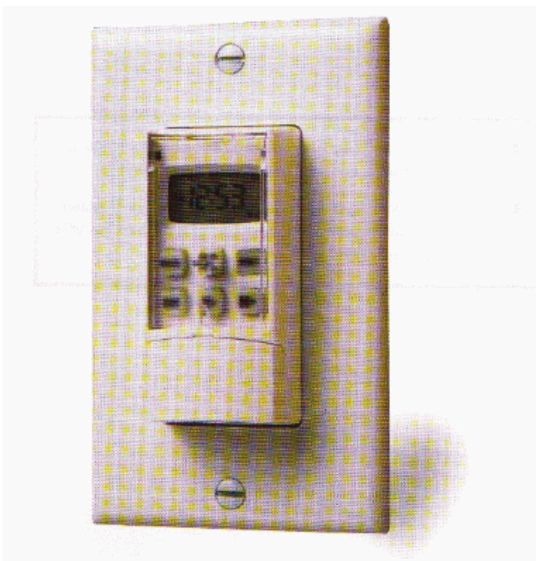

Up at 7 a.m., in bed at 11 p.m.—if your schedule falls into a predictable routine like this, the **SS7C 7-Day Wall Switch Timer** from Intermatic offers a simple way to put the bedroom lights on automatic pilot.

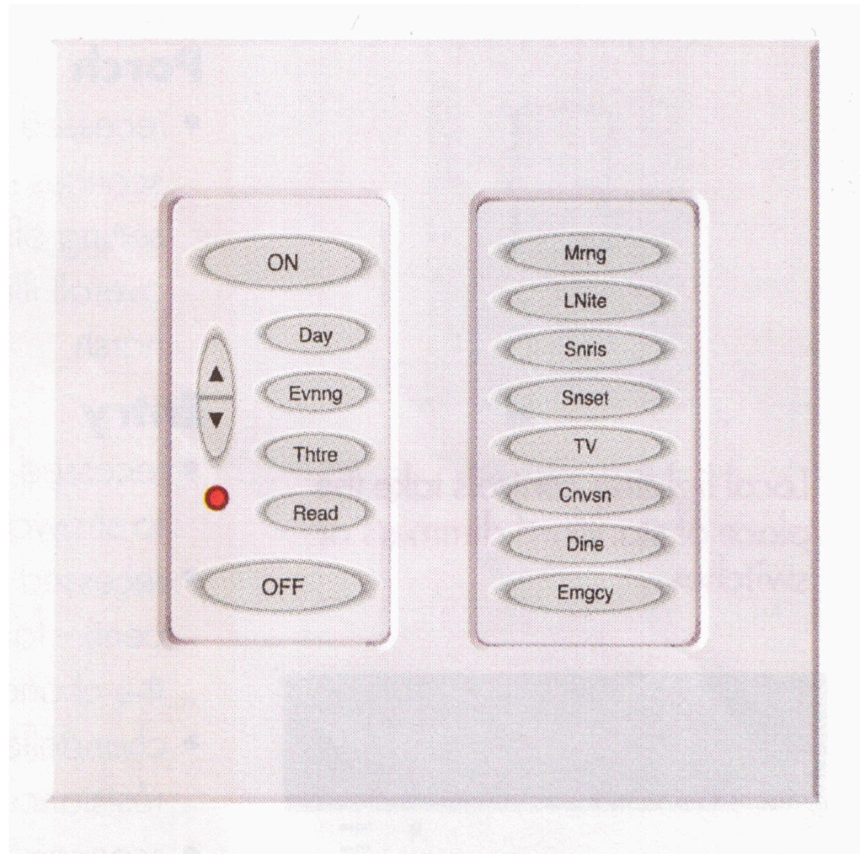

The Brillance II keypad from Lightolier

The adaptive house

Not a programmable house, but a house that *programs itself*.

House *adapts* to the lifestyle of the inhabitants.

House monitors environmental state and senses actions of inhabitant.

House learns inhabitants' schedules, preferences, and occupancy patterns.

House uses this information to achieve two objectives:
(1) anticipate inhabitant needs
(2) conserve energy

Domain: home comfort systems

- air heating
- lighting
- water heating
- ventilation

Tremendous potential cost/energy savings

single set back period on furnace	9–18%
multiple set back periods	25–30%
set back (electric) water heater	25%
shift majority of electric use off peak	20–40%

The adaptive house

Residence in Marshall, Colorado, outside of Boulder

Some of the gang

Great room

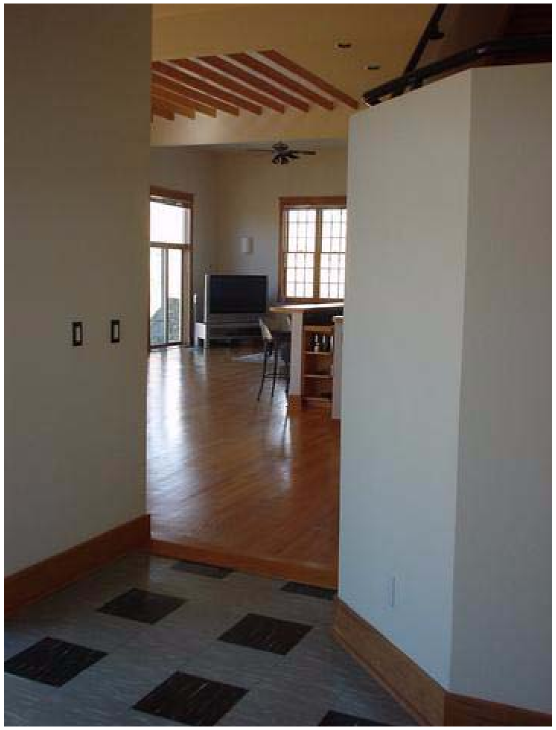

Bedrooms and bathrooms

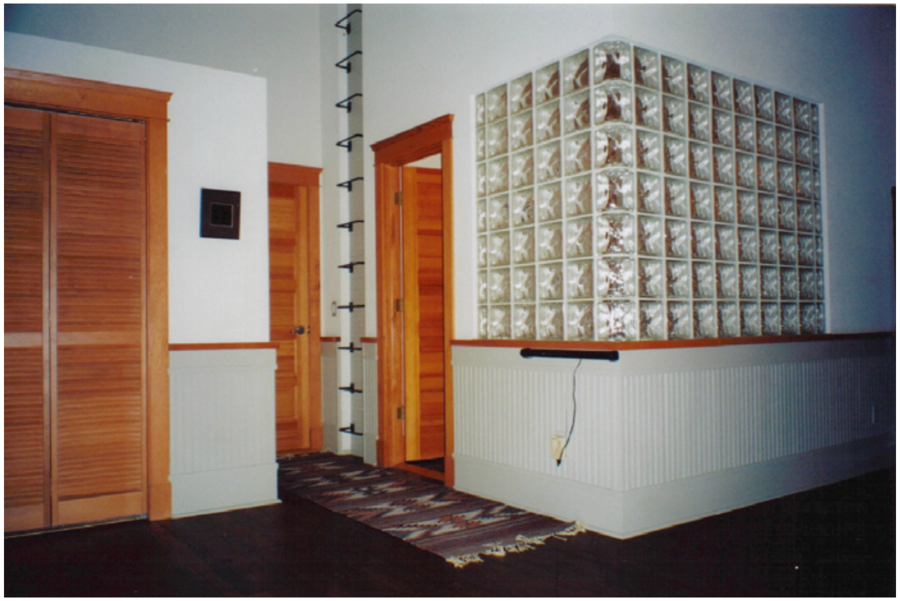

Sensors


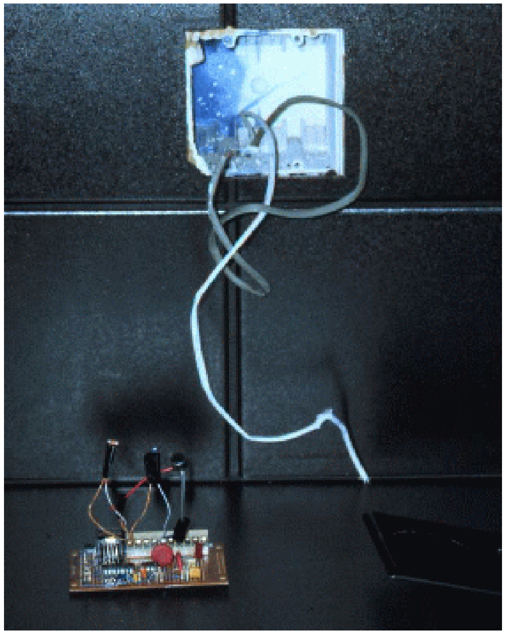

Sensors

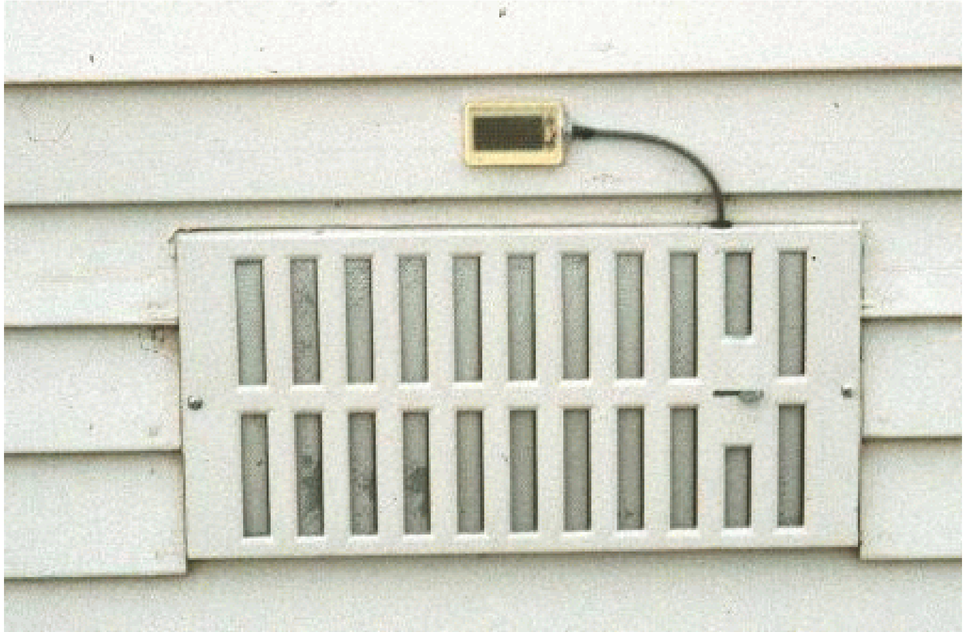

Water heater

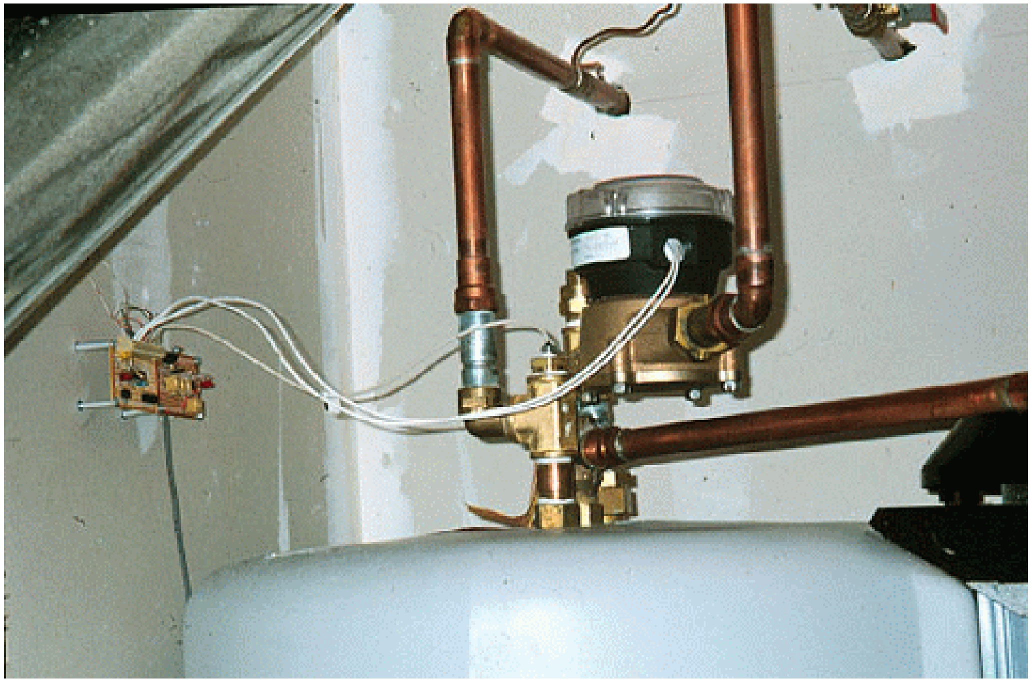

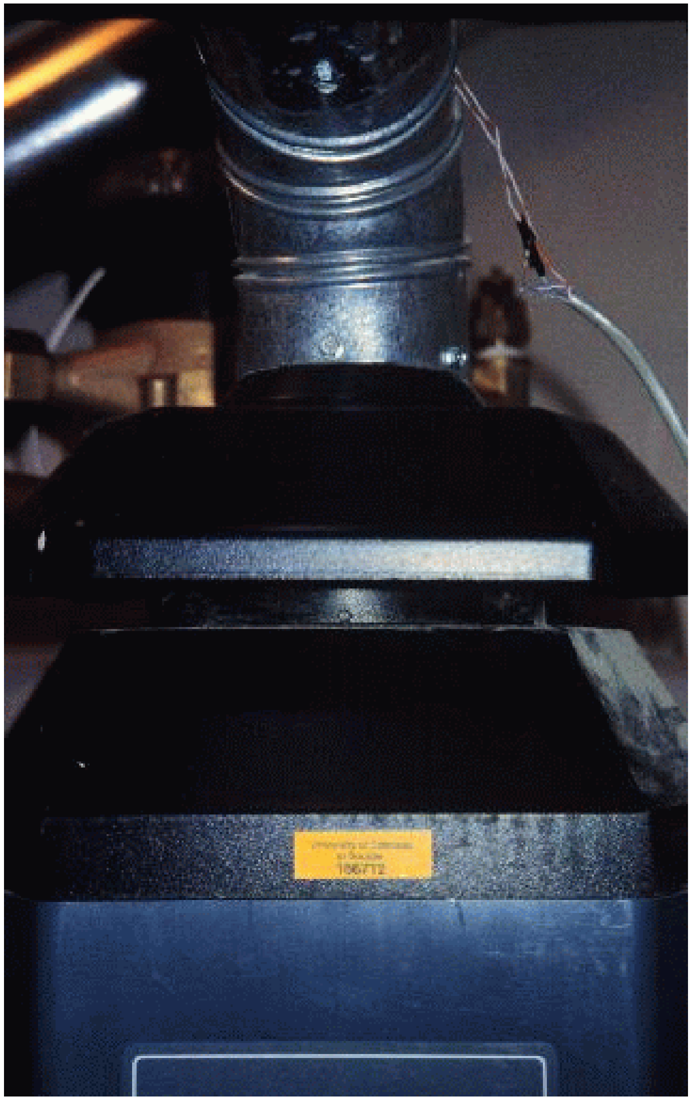

Furnace

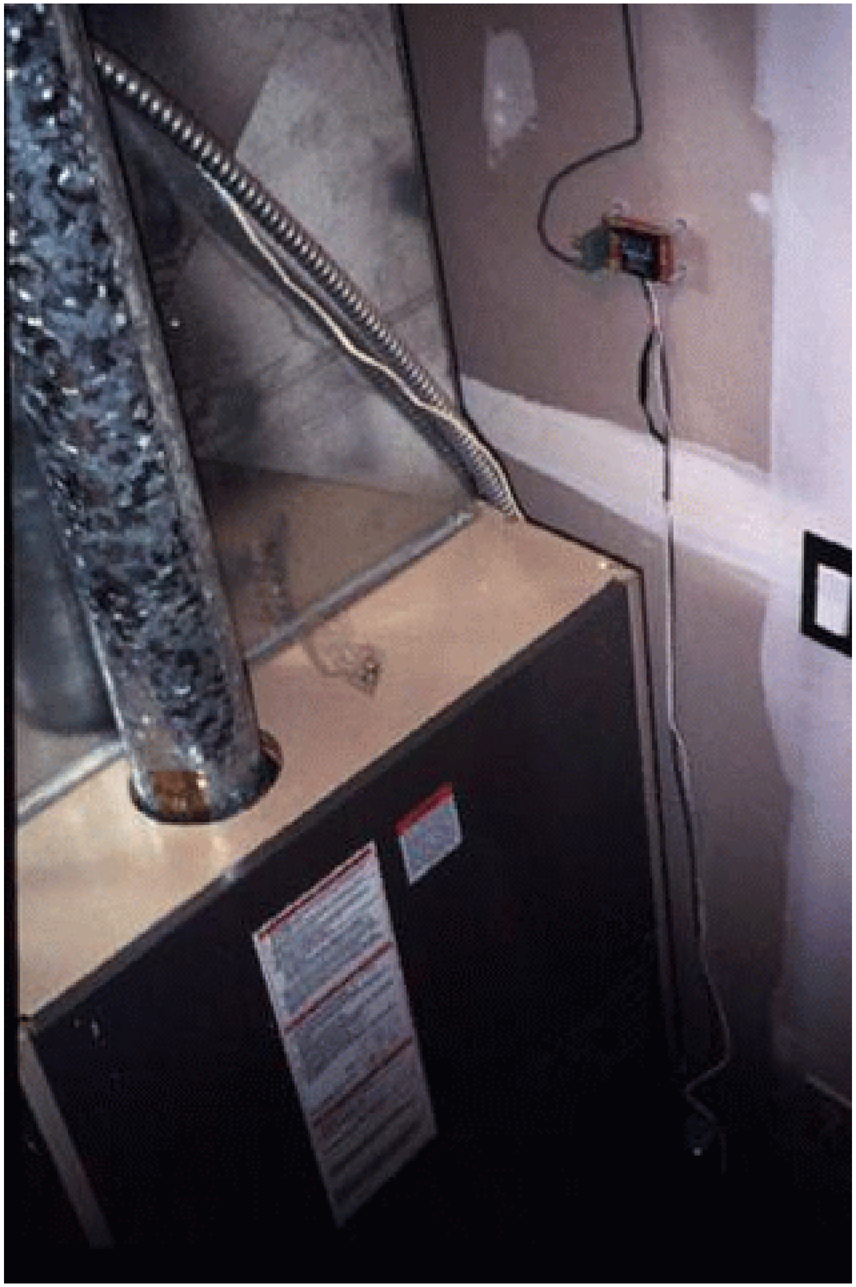

Controls

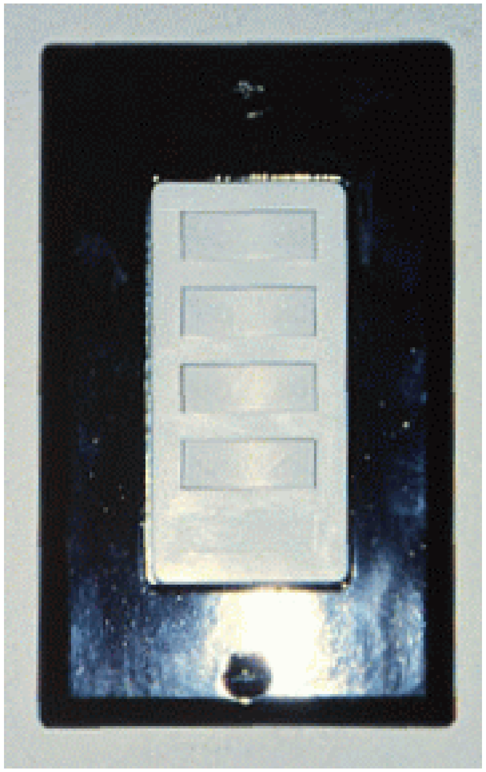

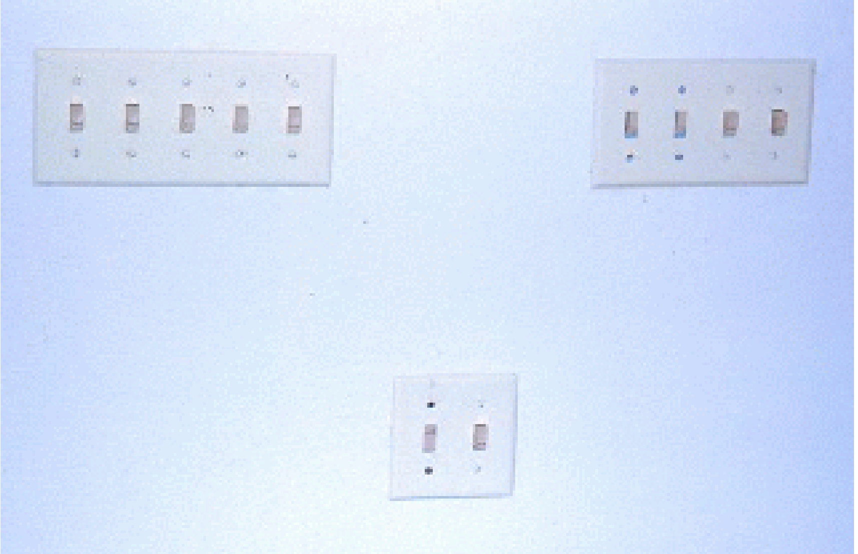

Computers

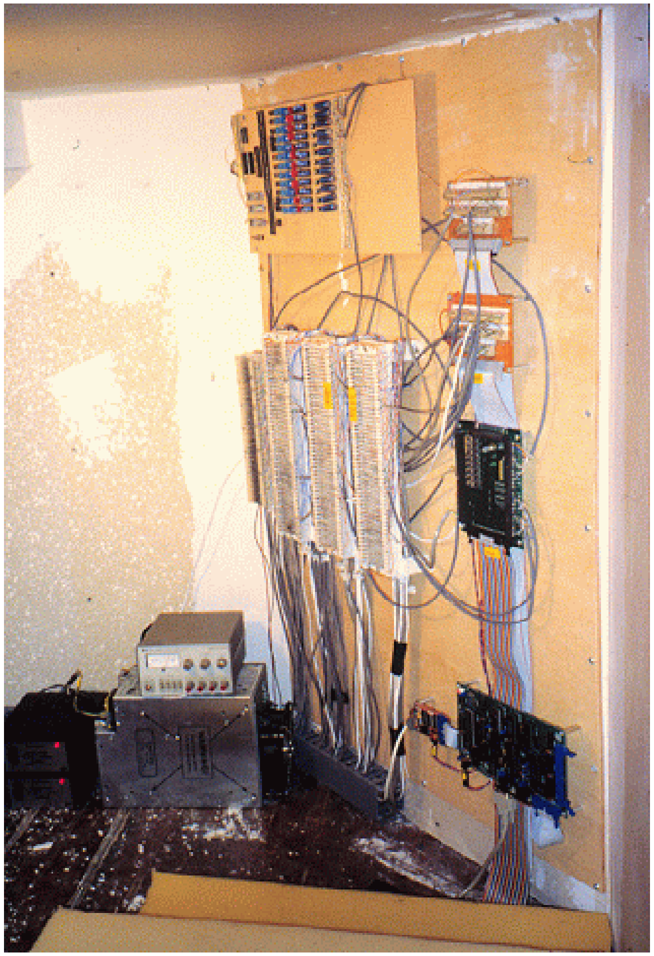

Training signals

Actions performed by inhabitant specify *setpoints*

→ anticipation of inhabitant desires

Gas and electricity costs

→ energy conservation

An optimal control framework

Each constraint has an associated cost:

discomfort cost if inhabitant preferences are neglected

energy cost depends on device and intensity setting

The optimal control policy minimizes

$$J(t_0) = \mathrm{E}\left[\lim_{\kappa\to\infty} \frac{1}{\kappa} \sum_{t=t_0+1}^{t_0+\kappa} d(\mathbf{x}_t) + e(\mathbf{u}_t)\right]$$

where t **= index over nonoverlapping time intervals**
t_0 **= current time interval**
$\mathbf{u}_t$ **= control decision for interval** t
$\mathbf{x}_t$ **= environmental state during interval** t

ACHE
(Adaptive Control of Home Environments)

Separate control system for each task

- air temperature regulation
  - furnace
  - space heaters
  - fans
  - dampers
  - blinds
- lighting regulation
  - wall sconces
  - overhead lights
- water temperature regulation
  - hot water heater

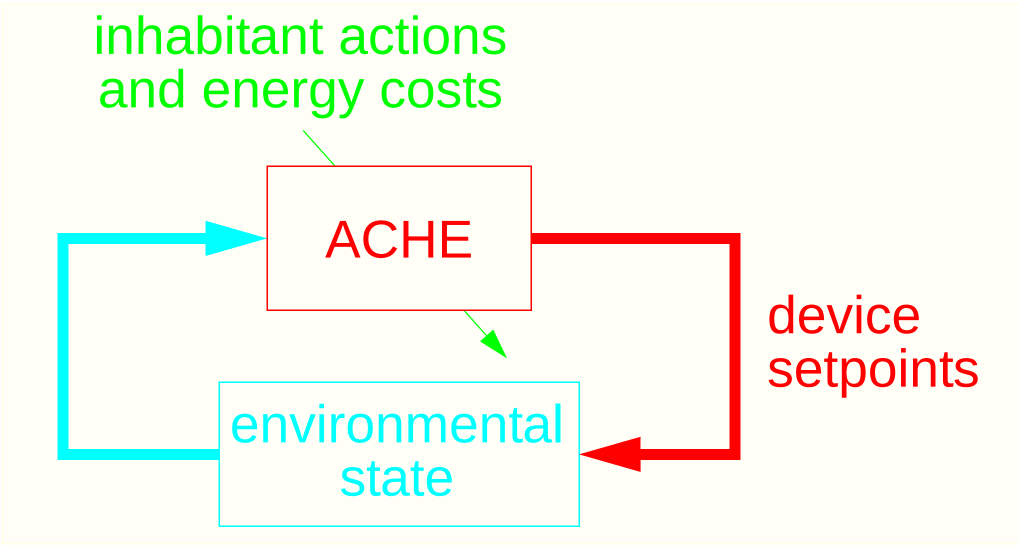

General architecture of ACHE

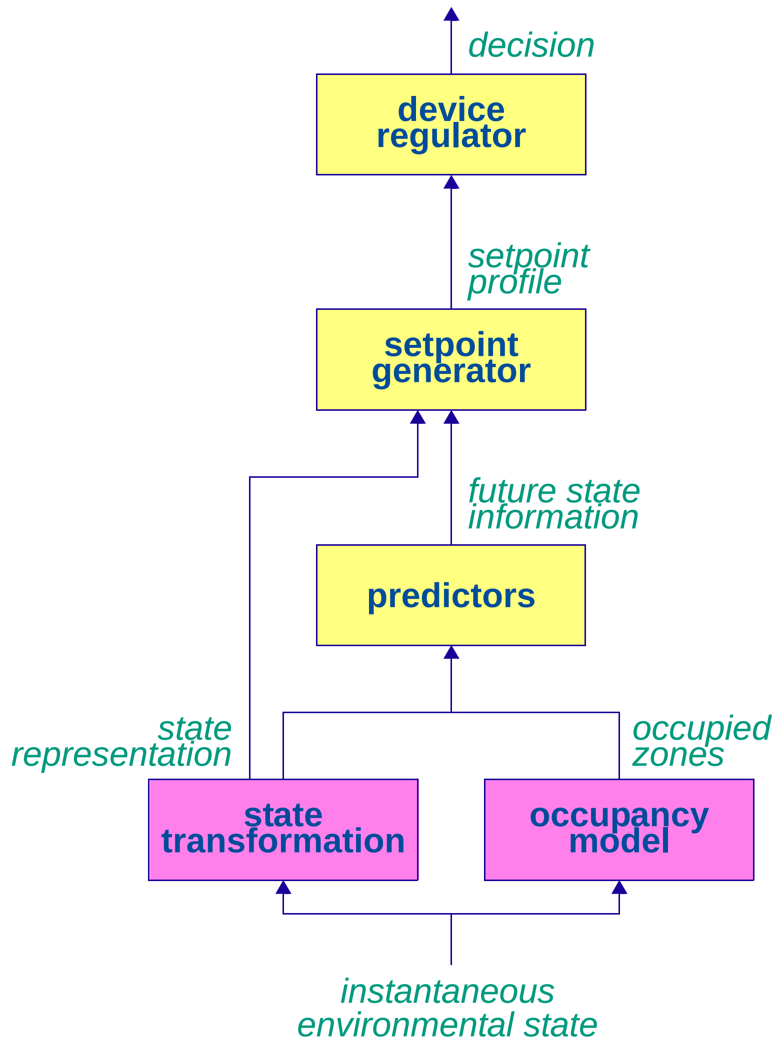

Knowledge encapsulation

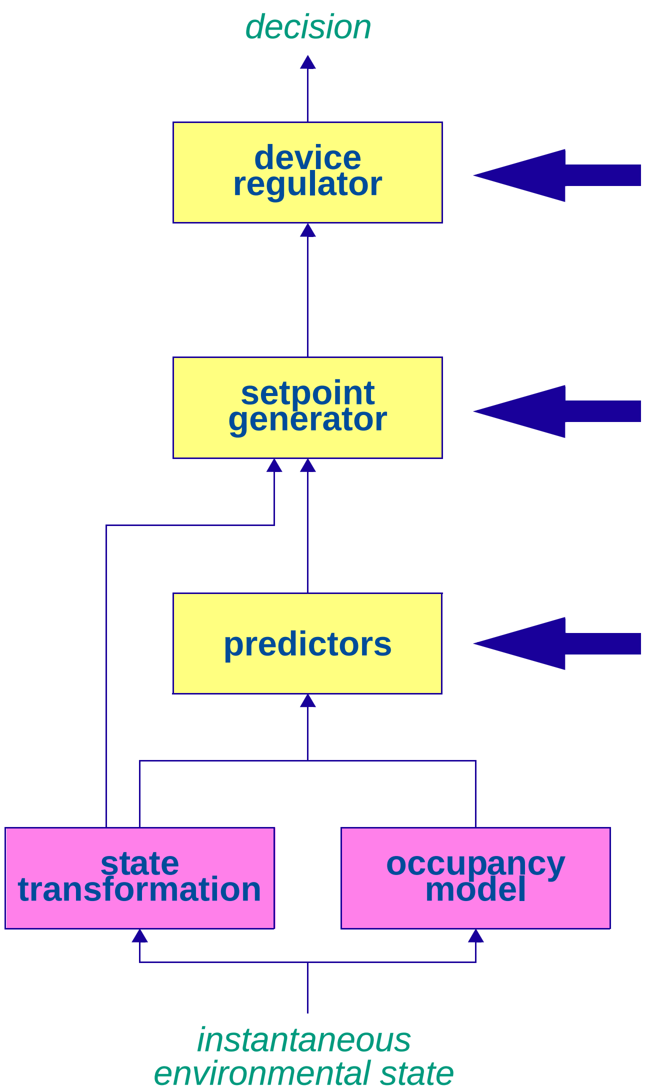

knowledge about energy devices and physical layout of house

knowledge about inhabitants' comfort needs and preferences

knowledge about inhabitants' schedules

Training procedures

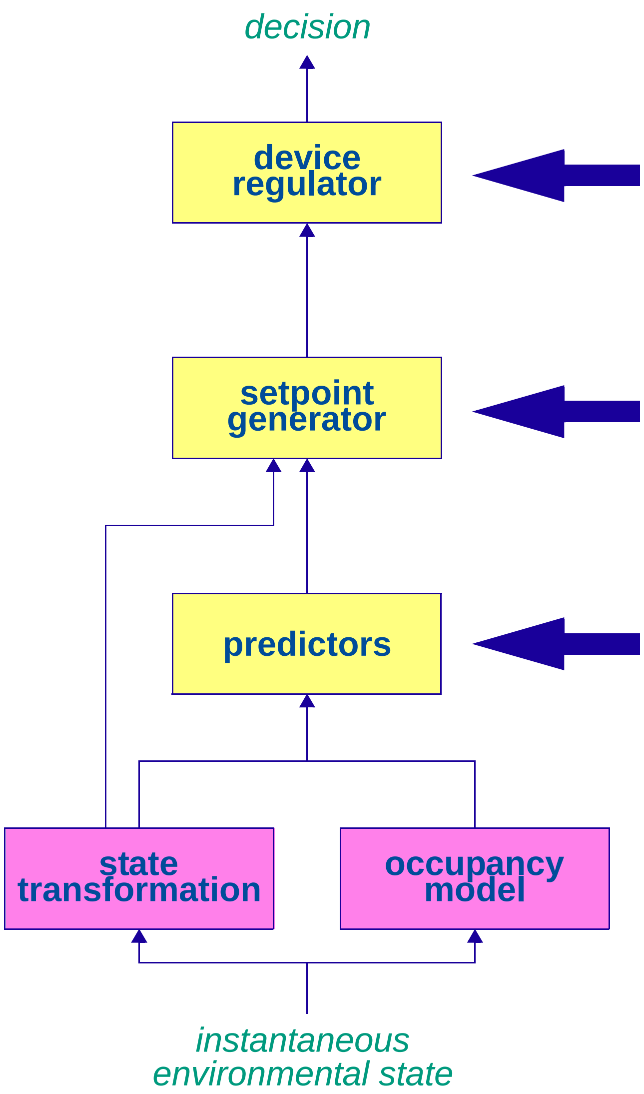

system identification using neural networks

reinforcement learning with look-up tables or memory-based approaches

neural networks trained with self-supervised learning

Lighting control

What makes lighting control a challenge?

Twenty-two banks of lights, each with 16 intensity levels; seven banks of lights in great room alone

Motion-triggered lighting does not work

Lighting moods

Two constraints must be satisfied simultaneously
- maintaining lighting according to inhabitant preferences
- conserving energy

Range of time scales involved

Sluggishness of system

Sequential decision problem

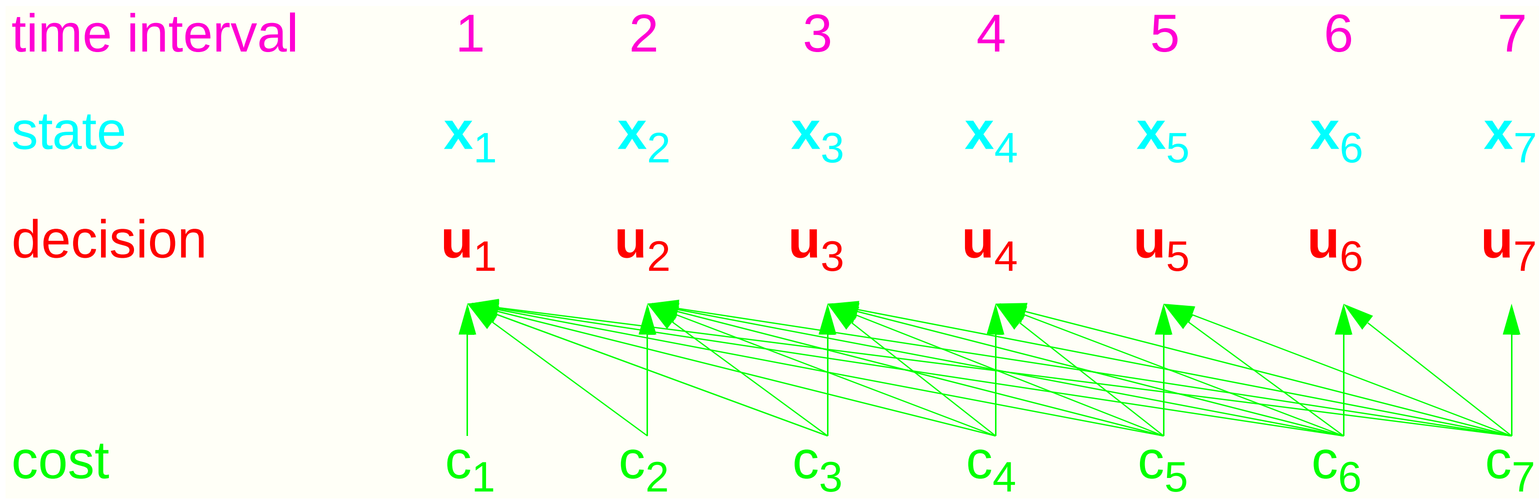

To learn, must determine which decisions are responsible for observed costs (*temporal credit assignment*).

Time scale dilemma

Control decisions must be responsive to changing environmental conditions.

Therefore, time intervals must be brief (~200 ms).

But shorter time intervals make learning more difficult.

Resolving the time scale dilemma

Event-based segmentation

Detect salient events such as zone entry, change in outdoor light level.
Window of time between events treated as basic interval.
Lighting control decision made when event occurs.

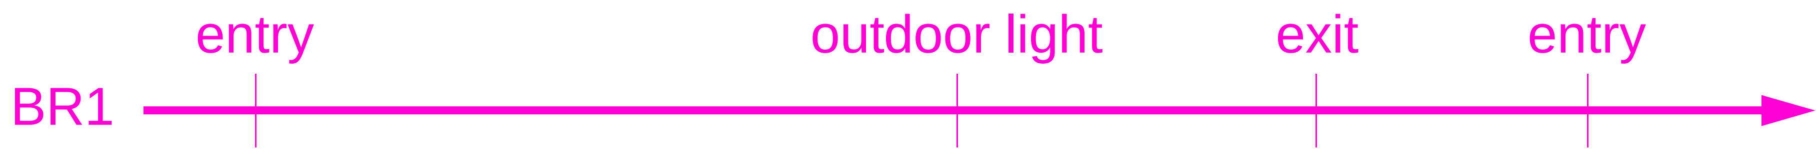

Temporal credit assignment problem greatly simplified.

Motivated by orienting response in biological systems.

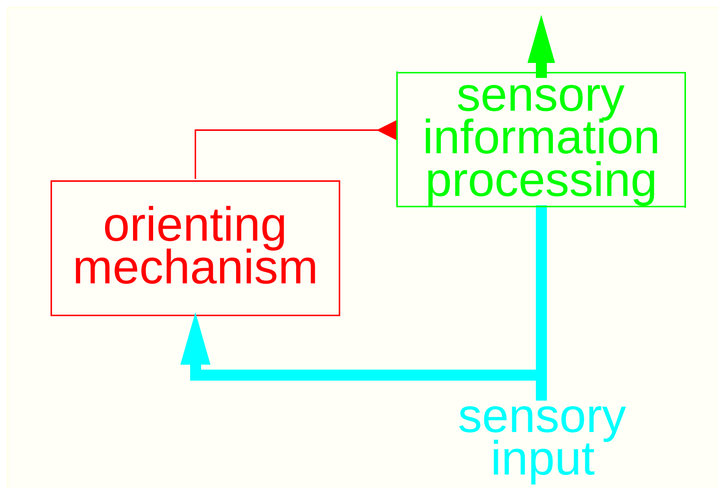

Resolving the sluggishness dilemma

***Anticipator*: Neural network that predicts which zone(s) will become occupied in the next two seconds**

Input

1, 3, and 6 second average of motion signals	(36)
instantaneous and 2 second average of door status	(20)
instantaneous, 1 second, and 3 second average of sound level	(33)
current zone occupancy status and durations	(16)
time of day	(2)

Output

p(zone *i* becomes occupied in next 2 seconds **|** currently unoccupied) (8)

Runs every 250 ms

Training anticipator

Occupancy model provides training signal

Two types of errors

miss
state(t – 2000 ms)
state(t – 1750 ms)
...
state(t – 250 ms)
→ zone *i* becomes occupied

false alarm
state(t) → zone *i* vacant

Training procedure

Given partially trained net, collect misses and false alarms.
Retrain net when 200 additional examples collected.
TD algorithm for misses

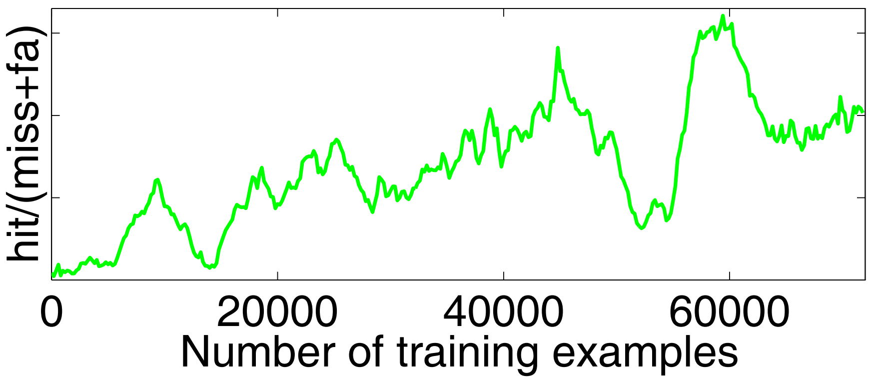

Examples of anticipator performance

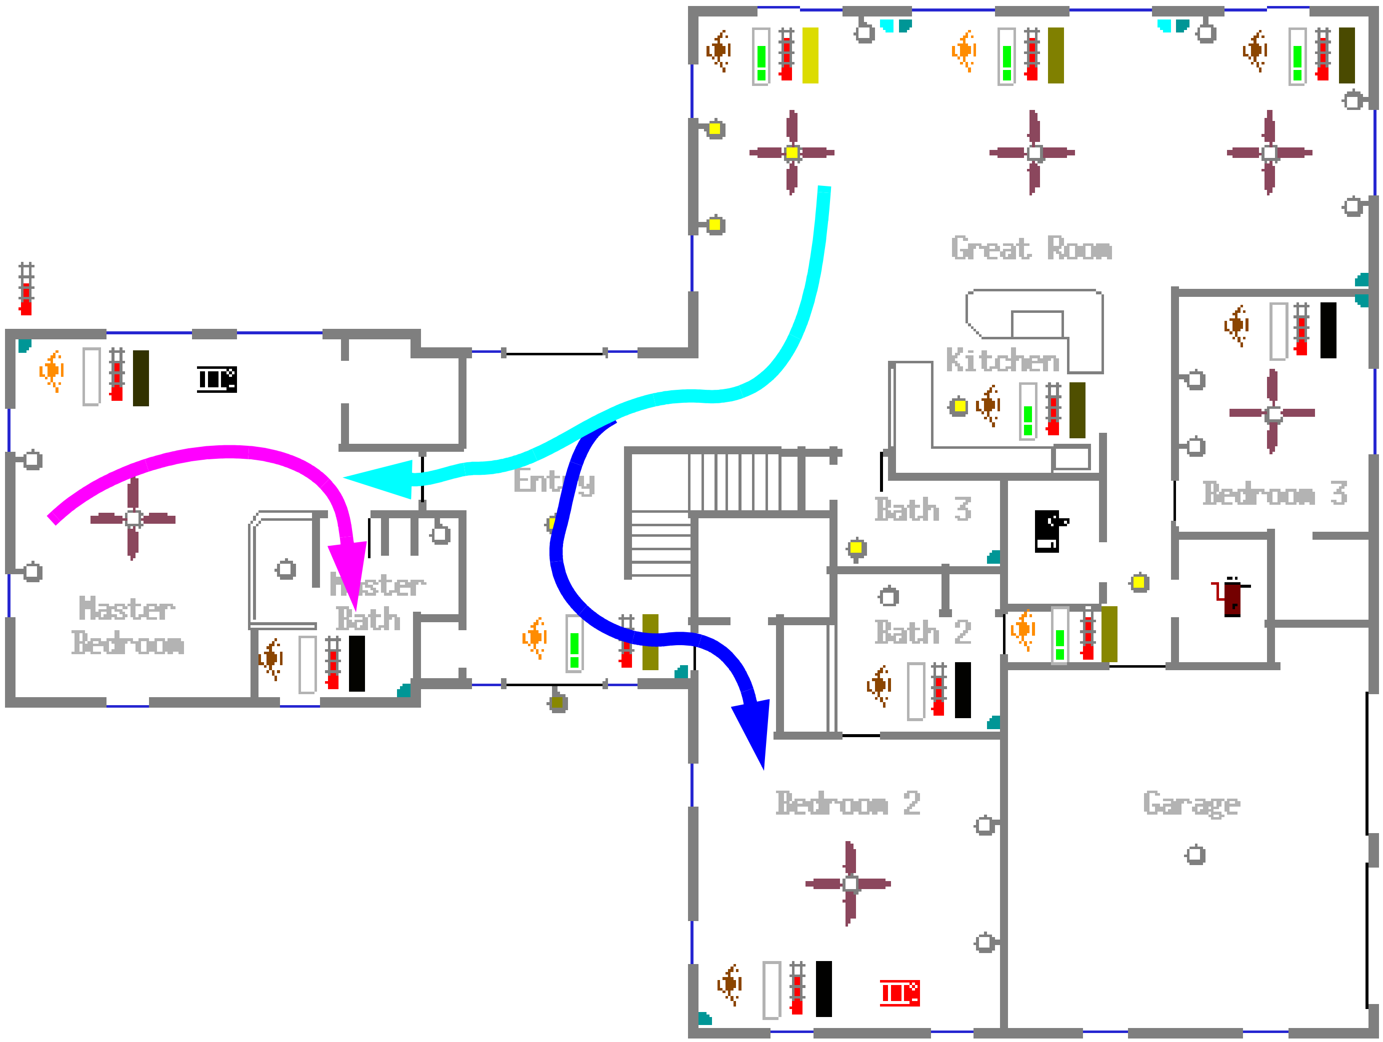

Lighting controller costs

Energy cost

7.2 cents per kW-hr

Discomfort cost

1 cent per device whose level is manually adjusted

Anticipator miss cost

.1 cent per device that was off and should have been on

Anticipator false alarm cost

.1 cent per device that was turned on

Results

- **about three months of data collection**
- **events logged only from 19:00 – 06:59**

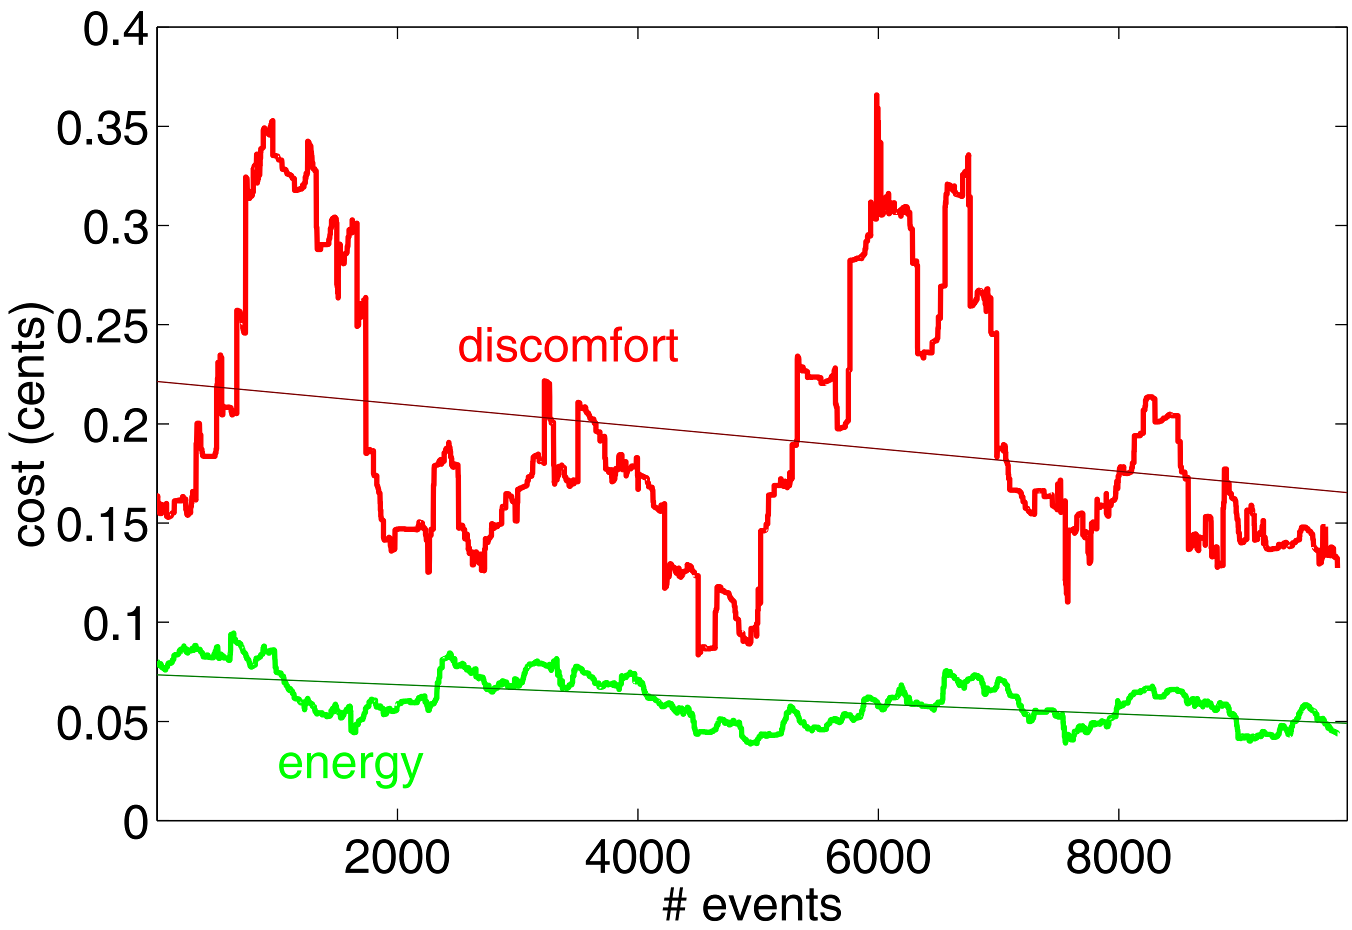

Air temperature control

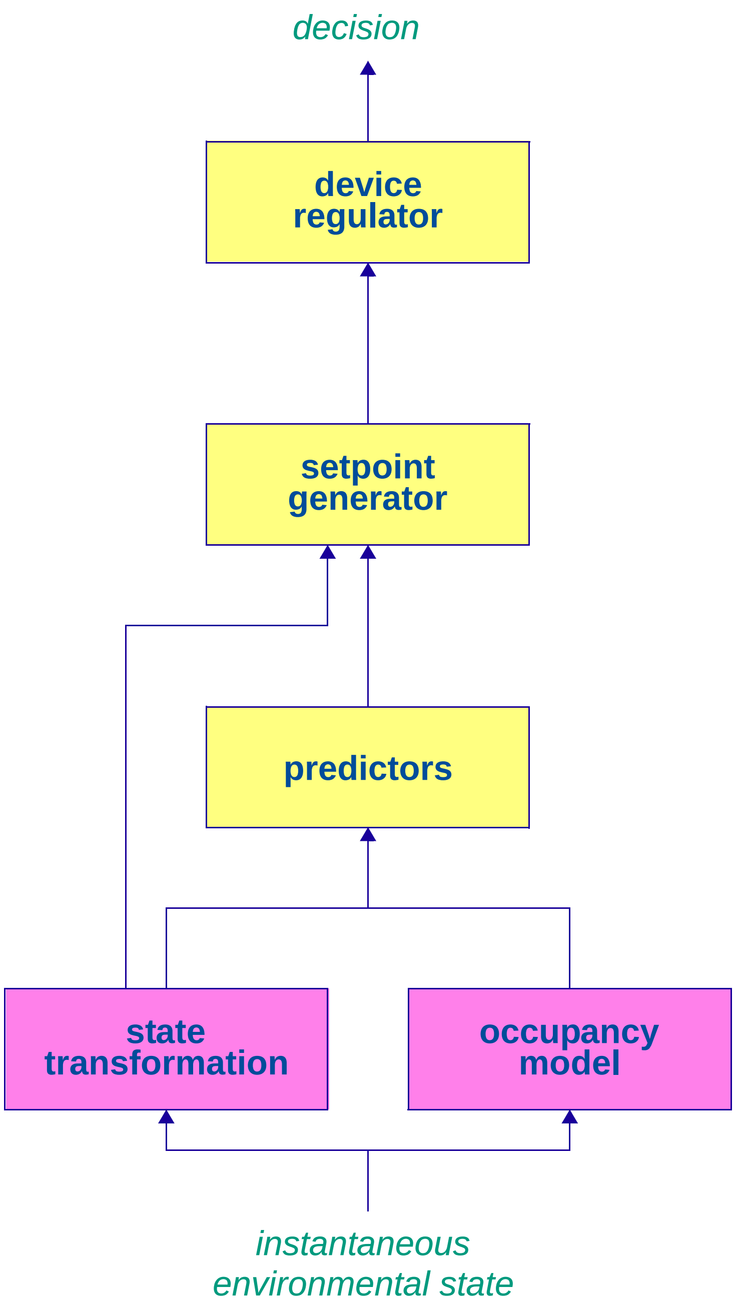

performs first decision in sequence

predictive optimal controller

Searches over a fixed horizon of κ decisions, δ minutes apart, for a decision sequence that minimizes expected cost.

predicts future occupancy status of house

Based on: time of day, day of week, average proportion of time home occupied in the 10, 20, and 30 minutes from present time of day on the previous 3 days and on the same day of the week during the past 4 weeks, proportion of time home was occupied during the past 60, 180, and 360 minutes

reports house occupancy

Misery cost

expected misery cost for decision sequence **u**

occupancy status

indoor air temperature

$$\overline{m}_{\mathbf{u}}(\mathbf{x}_t) = \overline{m}_{\mathbf{u}}[o(t), h_{\mathbf{u}}(t)]$$

$$= p\{o(t) = 0\}\, m[0, h_{\mathbf{u}}(t)] + p\{o(t) = 1\}\, m[1, h_{\mathbf{u}}(t)]$$

$$m(o, h) = o\, \rho\sigma \frac{\delta}{24 \times 60} \frac{\max(0, |\lambda - h| - \varepsilon)^2}{25}$$

salary ($/hr)

productivity loss (hr)

} *economic loss model*

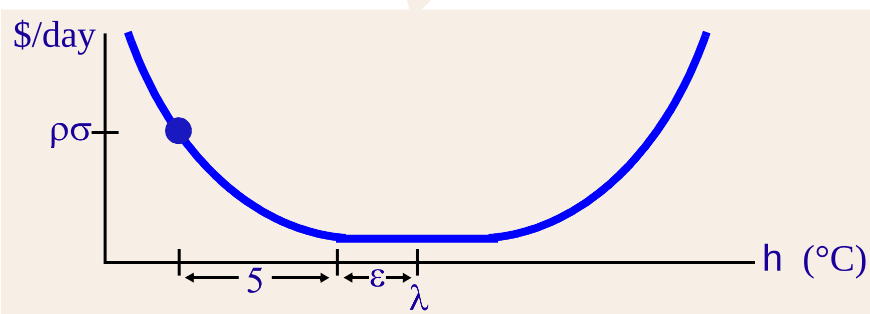

To estimate misery, must predict future *house occupancy* and *indoor temperature*.

Simulation methodology

Simulated environment

- thermal and comfort cost models are exact
- outdoor temperature, *g*, constant 0°C

Occupancy data

- real data collected from neural net house over an 8 month period

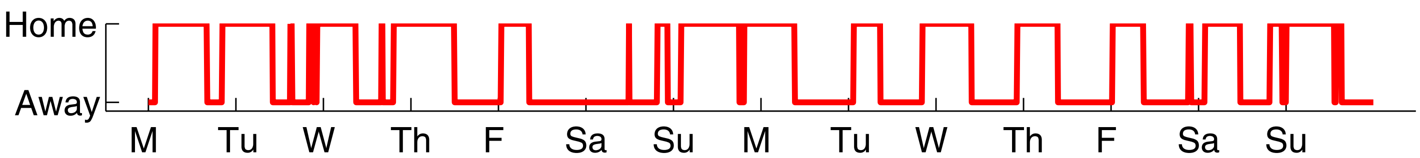

- artificial data, manipulating regularity of occupant schedule

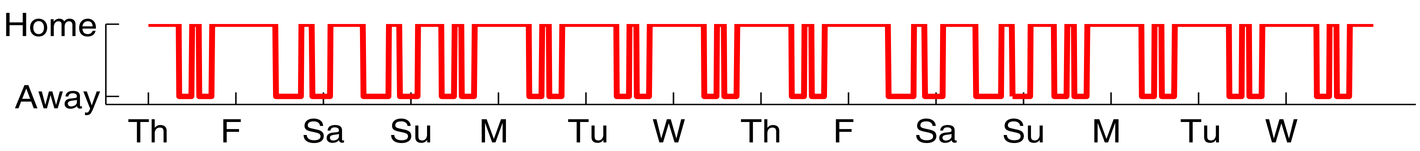

variability index = 0

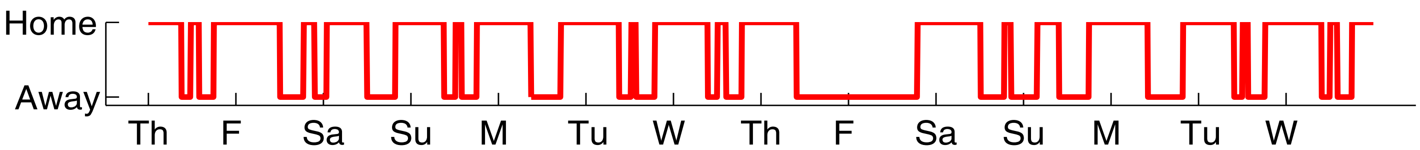

variability index = 1

Alternative heating policies

- **Constant Temperature Policy**

  setpoint = 22.5°C

- **Occupancy Triggered Policy**

  setpoint = 18°C if house empty
22.5°C if house occupied

- **Setback Thermostat Policy**

  setpoint = 18°C half hour before mean morning departure time for day of week
22.5°C half hour before mean evening return time

Each policy produces a setpoint at each time step.

Furnace turns on if actual temperature lower than setpoint.

Comparison of control policies using artificial occupancy data

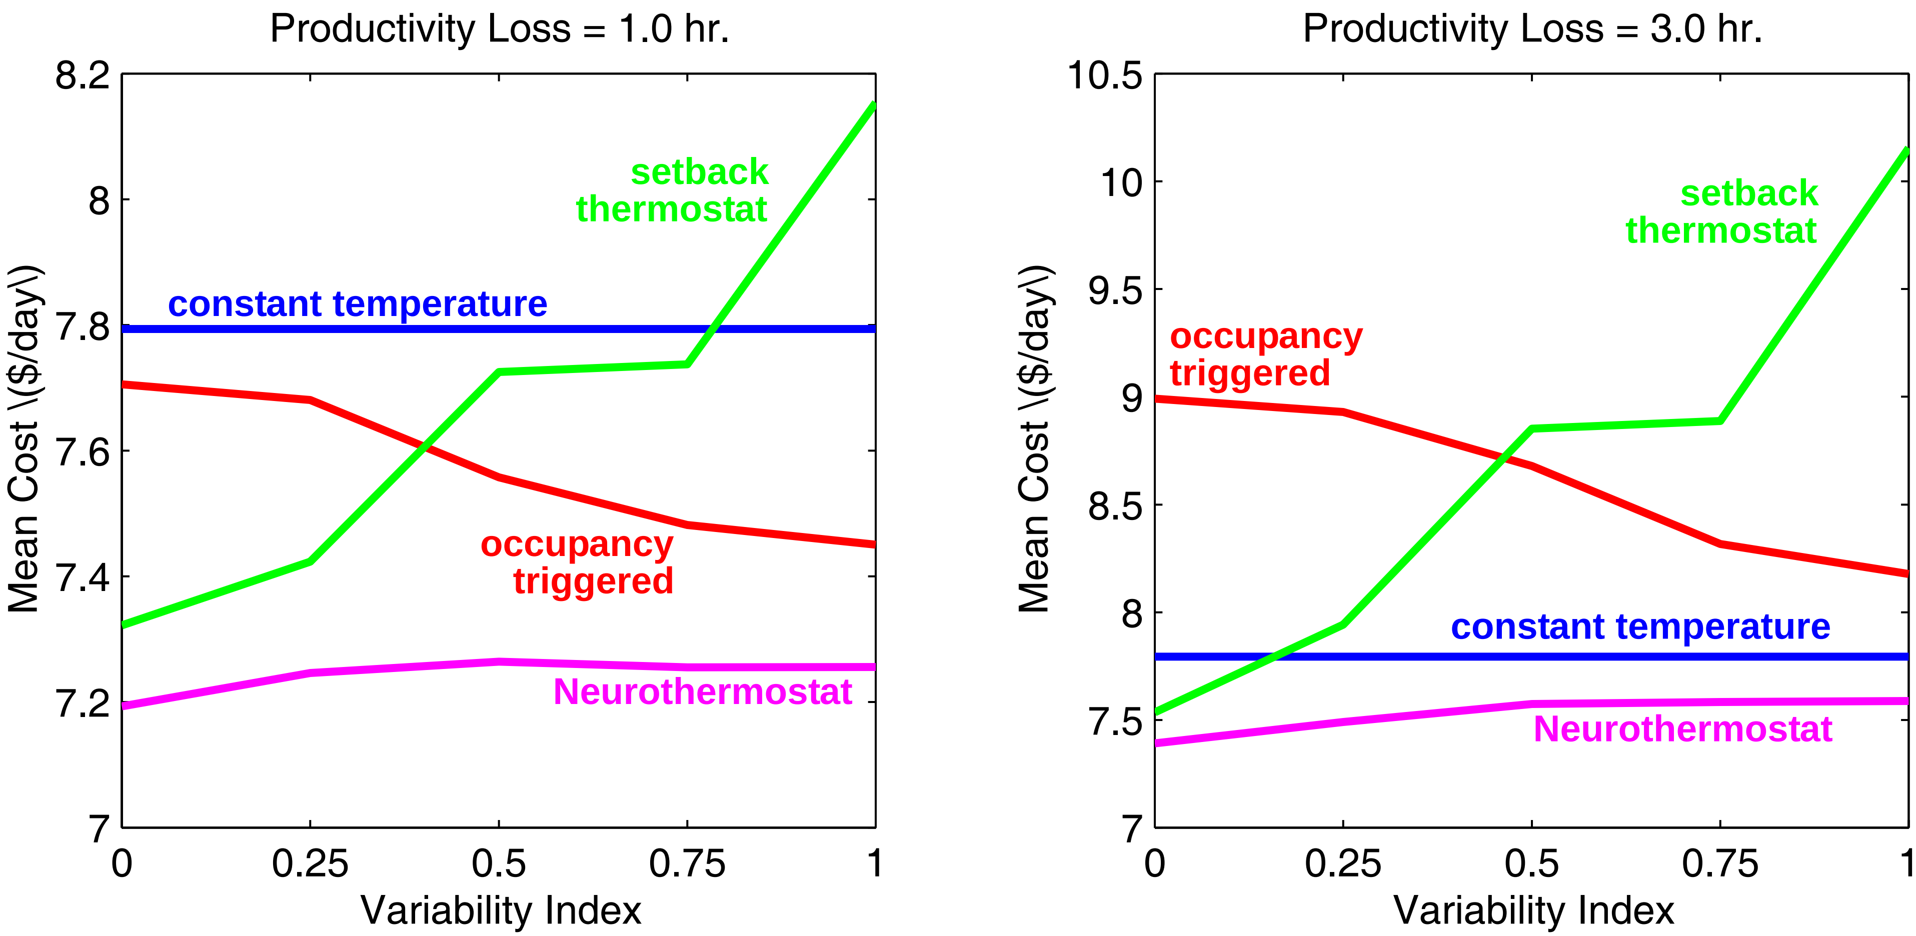

Comparison of control policies using real occupancy data

	Mean Daily Cost: productivity loss $\rho = 1$	Mean Daily Cost: productivity loss $\rho = 3$
Neurothermostat	$6.77	$7.05
constant temperature	$7.85	$7.85
occupancy triggered	$7.49	$8.66
setback thermostat	$8.12	$9.74

Sample Performance

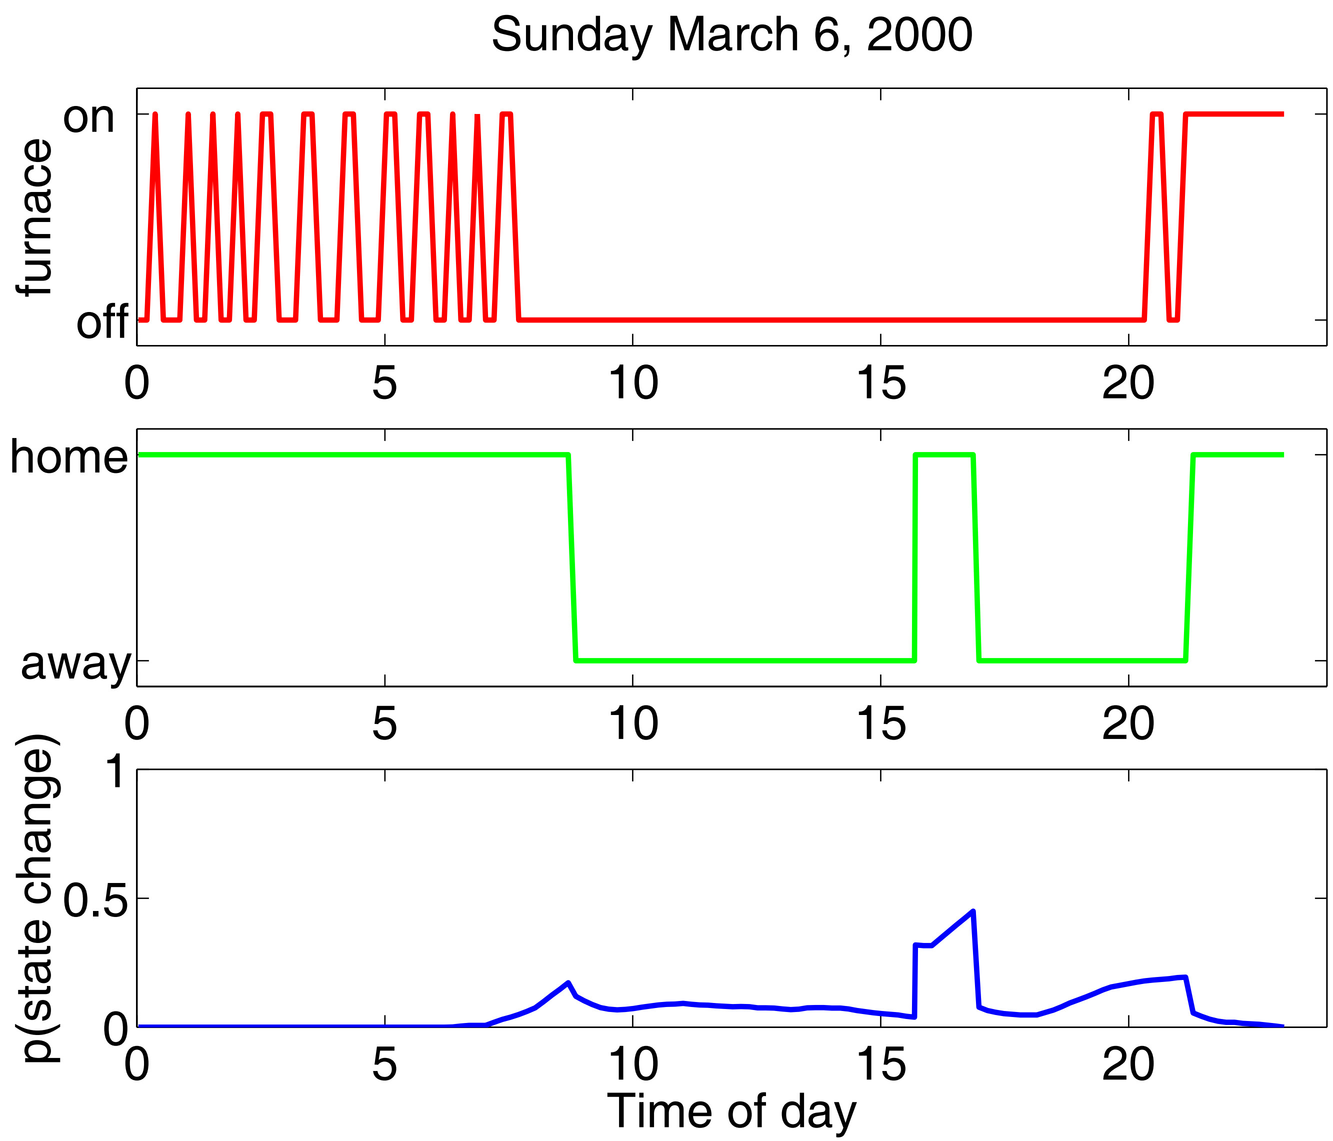

Relation Between Prediction and Temperature

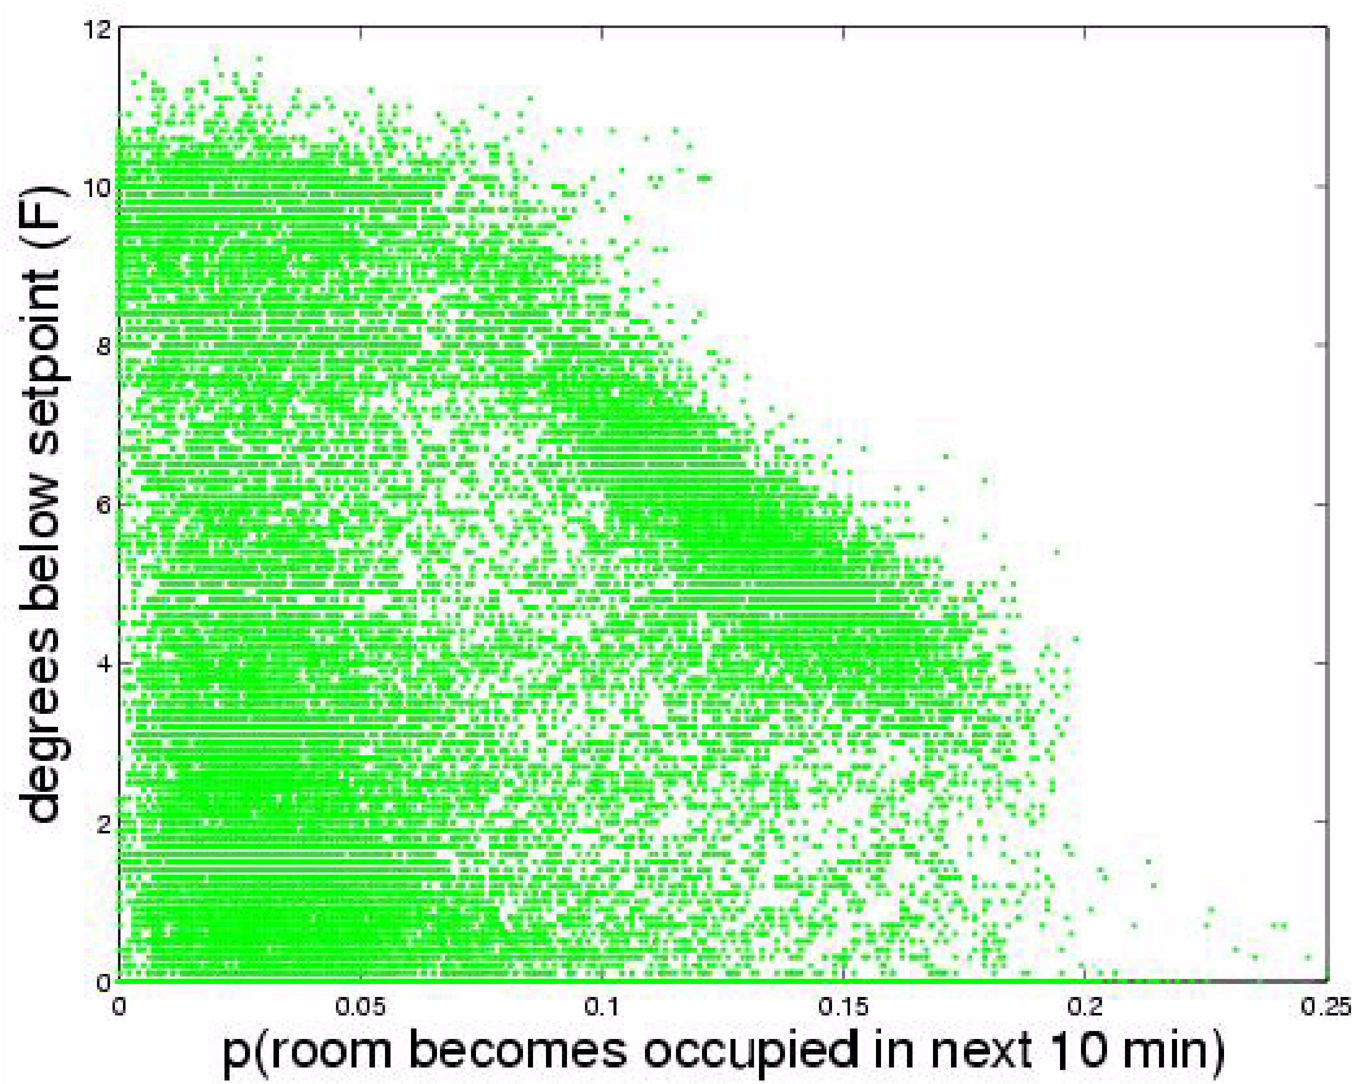

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Daily experience with ACHE was critical for evaluation.

- not another Media Lab demo
- Forced us to solve problems
- e.g., How do we design ACHE to work well out of the box?
- e.g., How much data history is relevant for prediction?

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Daily experience with ACHE was critical for evaluation.

Hide the interface.

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Daily experience with ACHE was critical for evaluation.

Hide the interface.

The more sensors the better.

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Daily experience with ACHE was critical for evaluation.

Hide the interface.

The more sensors the better.

Value of explicit activity classification

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Daily experience with ACHE was critical for evaluation.

Hide the interface.

The more sensors the better.

Value of explicit activity classification

Value in providing inhabitants with information to make informed decisions.

- e.g., consequences of turning up thermostat
- e.g., bathroom sensor

Lessons

Statistical regularities in inhabitant behavior can be exploited to save energy.

Even as a prototype system, ACHE was on balance a win.

Daily experience with ACHE was critical for evaluation.

Hide the interface.

The more sensors the better.

Value of explicit activity classification

Value in providing inhabitants with information to make informed decisions.

Mutual adaptation

Reinforcement learning

Dynamic programming can be used to perform optimization

- requires models of environment and cost function
- computing expectation may be very expensive

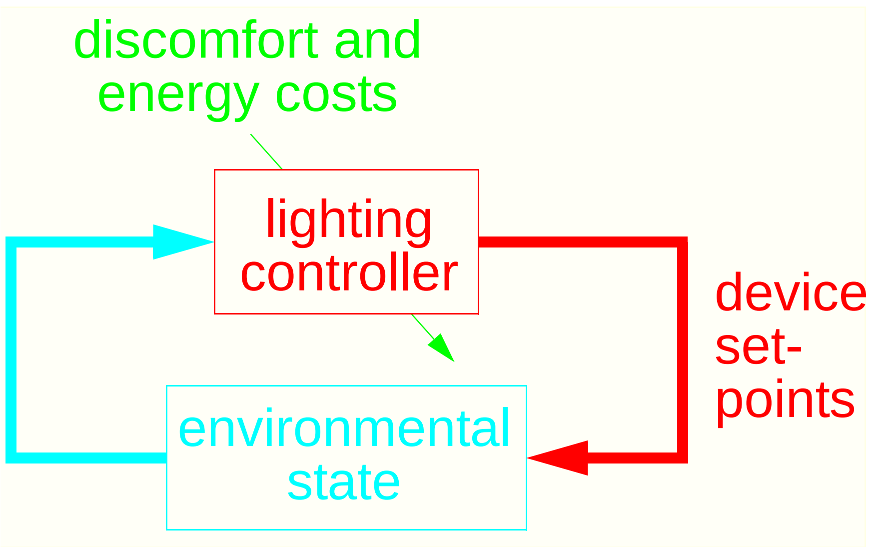

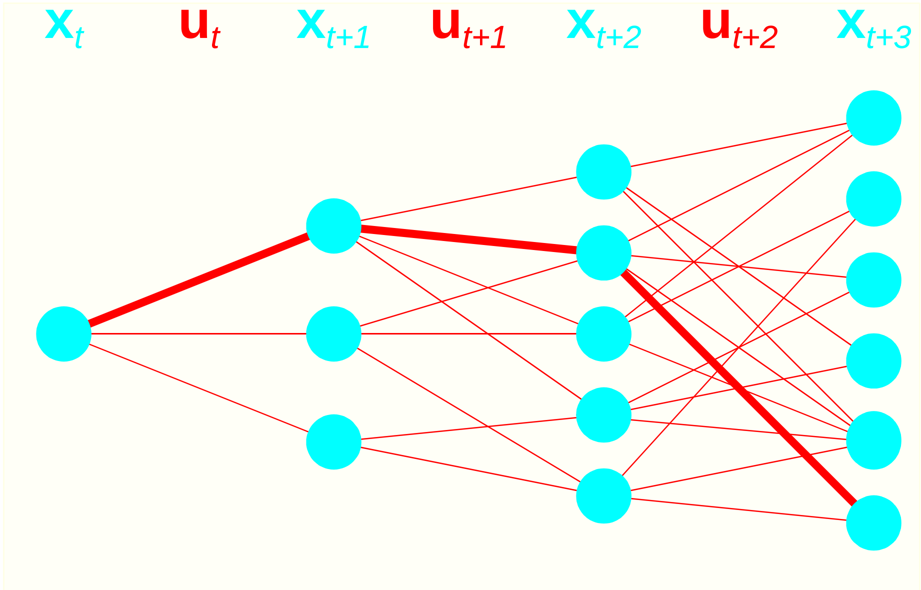

Reinforcement learning is a stochastic form of dynamic programming that samples trajectories in state space.

Q learning

(Watkins, 1989; Watkins & Dayan, 1992)

Q(x,u): If action u is taken in state x, what is the minimum cost we can expect to obtain?

Policy based on Q values:

exploration rate

$$\pi(\mathbf{x}_t) = \begin{cases} \text{argmin}_{\mathbf{u}} Q(\mathbf{x}_t, \mathbf{u}_t) & \text{with probability } (1-\theta) \\ \text{random} & \text{with probability } \theta \end{cases}$$

Incremental update rule for Q values:

$$Q(\mathbf{x}_t, \mathbf{u}_t) \leftarrow (1-\alpha) Q(\mathbf{x}_t, \mathbf{u}_t) + \alpha \max_{\hat{\mathbf{u}}} [c_t + \lambda Q(\mathbf{x}_{t+1}, \hat{\mathbf{u}})]$$

learning rate

discount factor

Given fully observable state, infinite exploration, etc.,

guaranteed to converge on optimal policy.

Decisions have no long term consequences

Effect of decision completely undone by subsequent decision.