CCAI Teaser: *OfficeLearn*, an OpenAI gym environment for office energy demand response

Oct 28, 2020

What's happening here?

Curtailment: current strategy when demand of wind or solar is greater than generation, to shut it off.

CAISO 2020 had days when 30% of solar was curtailed.

Can we abstract away the question of *how* humans respond to price, and make an AI learn it from interaction?

Can we design an RL agent that translates grid prices to human optimized DR prices?

Part 1:

OfficeLearn: An OpenAI Gym Environment

We implement the following flow of actions

-all at once
-hour/hour

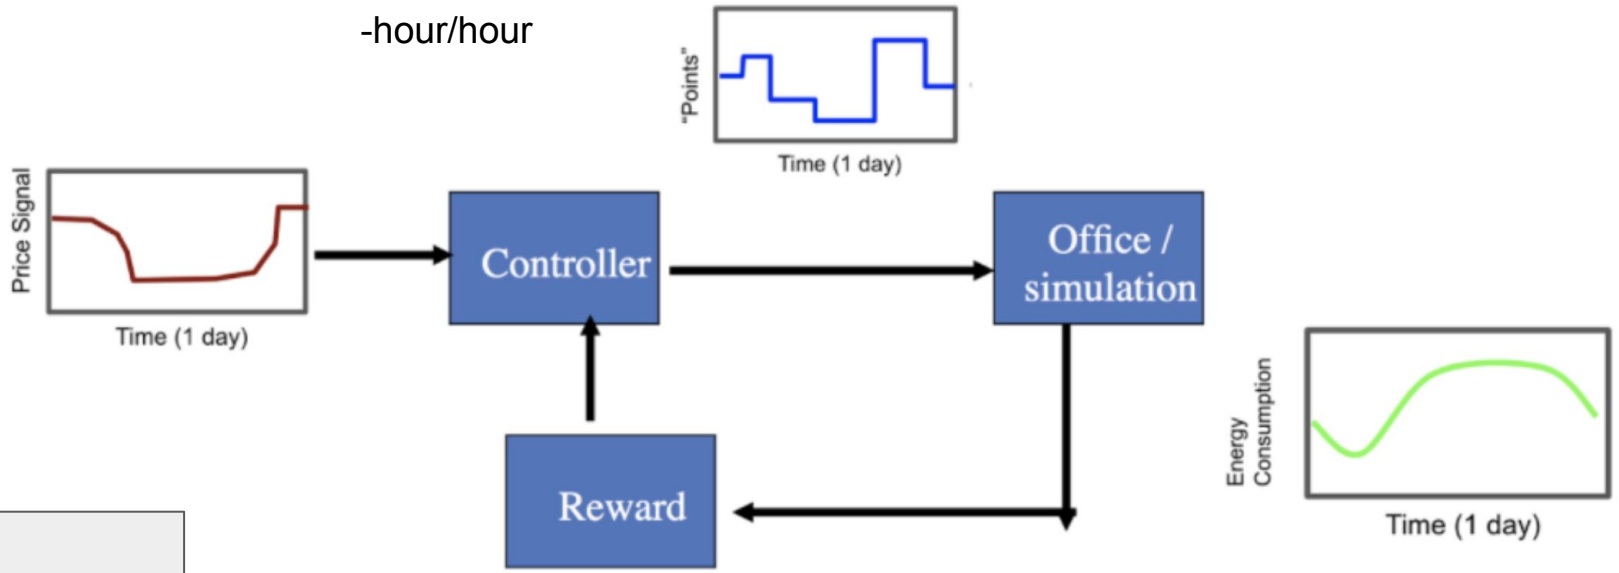

OfficeLearn

OfficeLearn State Space setup

State-space:

**Daily - [grid_price_hour_1,...,
grid_price_hour_10,
yesterday_energy_hour_1,...,
yesterday_energy_hour_10]**

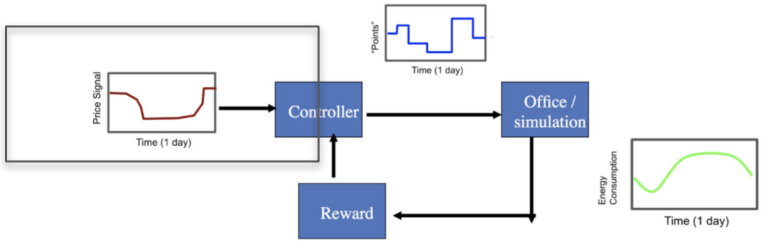

OfficeLearn

OfficeLearn State Space setup

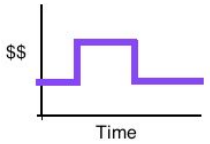

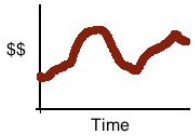

	Time Of Use (pricing)	Real Time Pricing
One Day of prices	(most simple)	
Year of prices		(most complex)

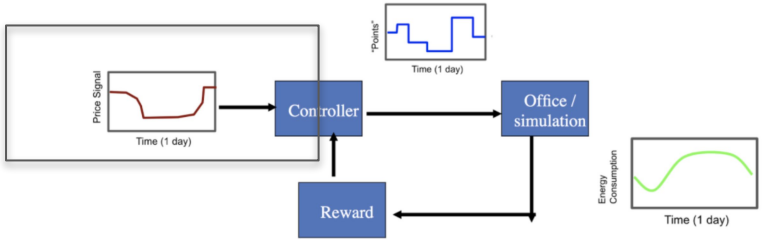

OfficeLearn

OfficeLearn Action Space setup

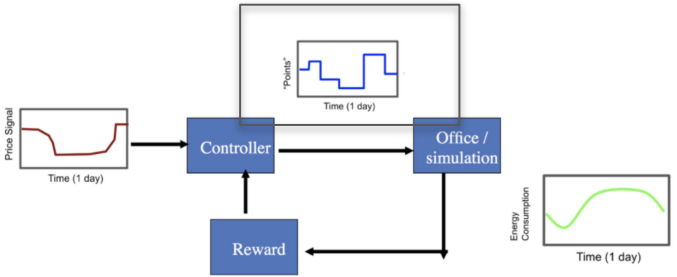

Action space:

[a_1, a_2, …, a_10] with *a*'s in (0, 10)

OfficeLearn

OfficeLearn Office Simulation

Simulation of Office Person

Curtail_and_Shift_person

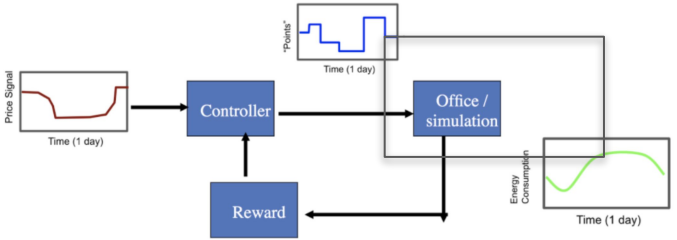

OfficeLearn

OfficeLearn: Reward

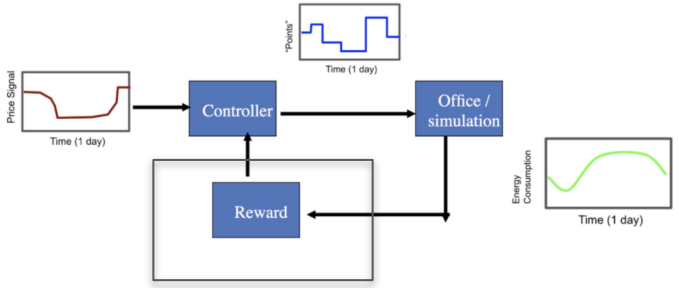

Log cost regularized

d= actual demand
*d**=ideal demand
g= grid prices

$$R(d) = -d^T g - \lambda(\sum d < 10 * (.5 * b_{max})$$

OfficeLearn

Curtail and Shift