6th International Conference on Smart Energy Systems
6-7 October 2020
#SESAAU2020

Reinforcement Learning based HVAC Optimization in Factories

- Debmalya Biswas

Powered by

HVAC Energy Optimization

- As part of our sustainability efforts, we are exploring ways to optimize the HVAC (Heating, Ventilation and Air Conditioning) controls in our factories.

- This is a complex problem as it requires computing an optimal state taking into account multiple variable factors, e.g. the occupancy in a building zone, temperature requirements of operating machines, air flow dynamics within the building, external weather conditions, etc.

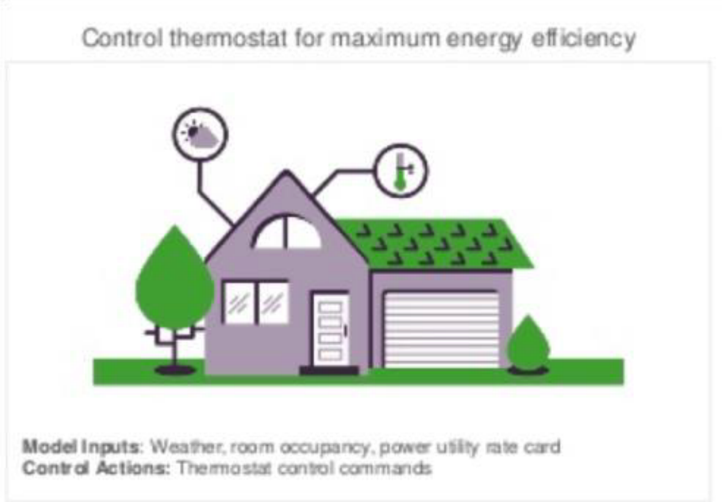

HVAC Functioning

- The primary goal of the HVAC units is to keep the temperature and (relative) humidity within the prescribed manufacturing tolerance ranges.

- By controlling 4 Output valves: Cooling, Heating, Re-heating and Humidifier.

- This needs to be balanced with energy savings and CO_2 emission reductions to offset the environmental impact of running them.

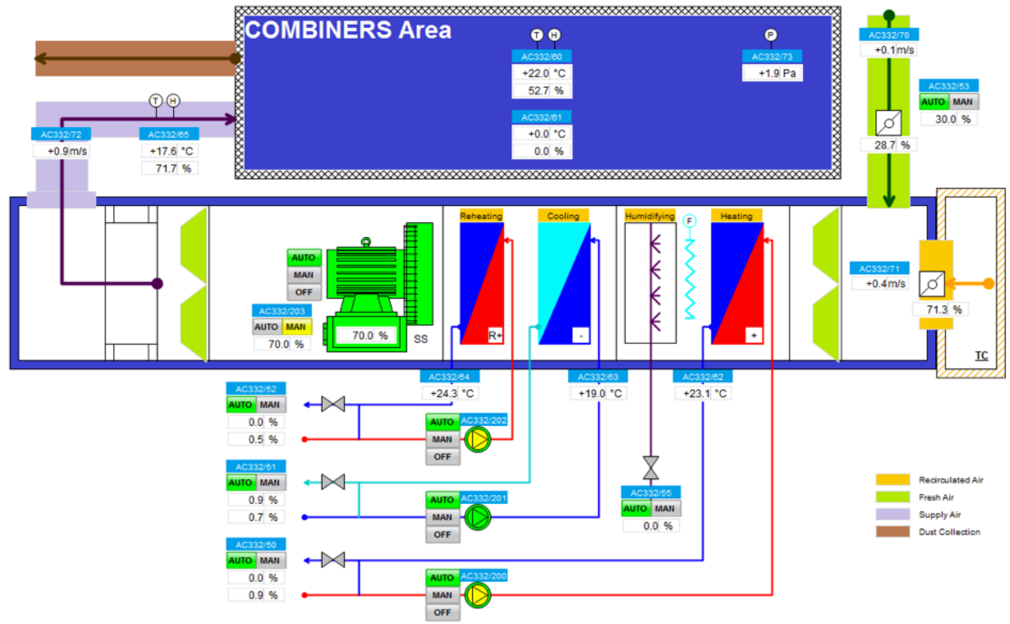

Reinforcement Learning (RL) Basics

- RL refers to a branch of AI/ML, which are targeted towards goal-oriented problems.
- RL algorithms are able to achieve complex goals by maximizing a reward function over many steps, e.g. the points won in a game over many steps.
- The reward function works similar to incentivizing a child with candy and spankings, such that the algorithm is penalized when it takes a wrong decision and rewarded when it takes a right one – this is reinforcement.

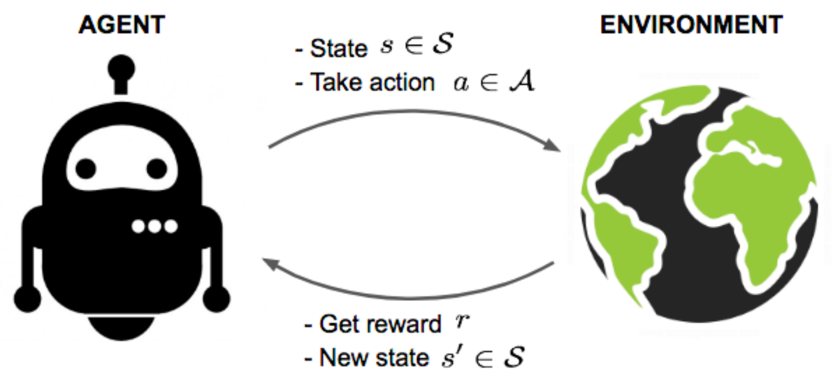

RL Formulation

- At any point in time, a factory zone is in a state characterized by the temperature and (relative) humidity values observed inside and outside the zone.
- The game environment in this case corresponds to the temperature and humidity tolerance levels, which basically mandate that the zone temperature and humidity values should be within the range: 19–25 degrees and 45–55%.
- The set of available actions in this case are the Cooling, Heating, Re-heating and Humidifier valve opening percentages (%).

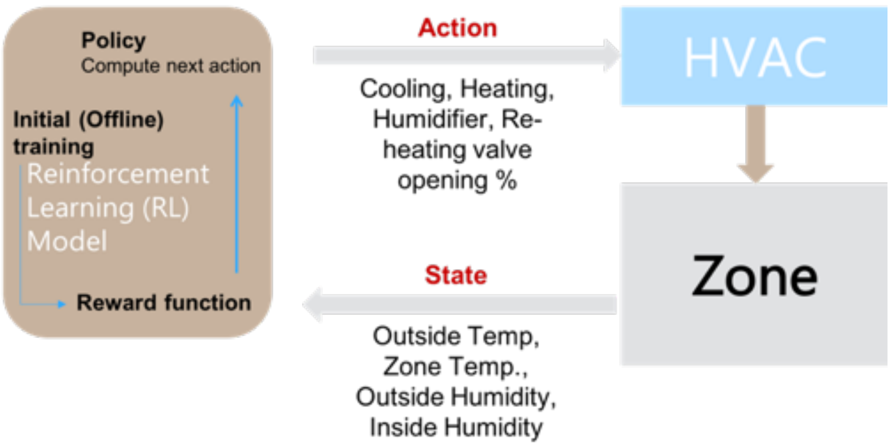

RL Formulation

- At any point in time, a factory zone is in a state characterized by the temperature and (relative) humidity values observed inside and outside the zone.

- The game environment in this case corresponds to the temperature and humidity tolerance levels, which basically mandate that the zone temperature and humidity values should be within the range: 19–25 degrees and 45–55%.

- The set of available actions in this case are the Cooling, Heating, Re-heating and Humidifier valve opening percentages (%).

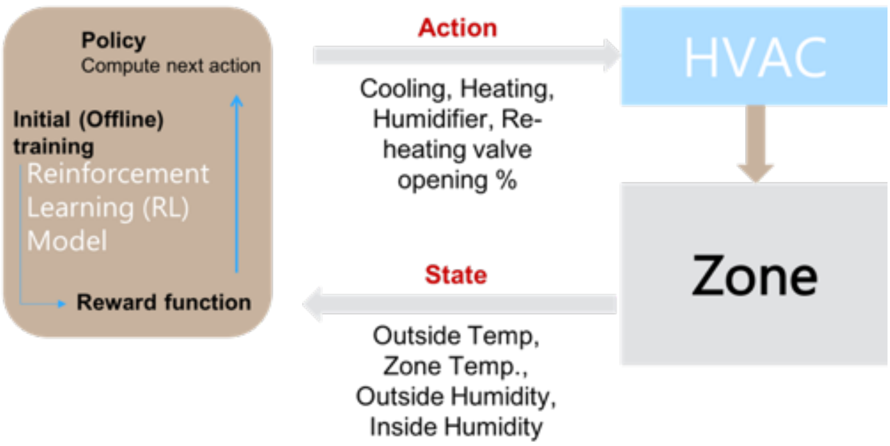

RL Formulation – Reward Function

- The Reward Function assigns a reward to each action based on the following three parameters:

$$Reward(a) = (w_1 \times SC) - (w_2 \times EC) - (w_3 \times TV)$$

- A control strategy is to decide on the weightage of the three parameters: Setpoint Closeness (SC), Energy Cost (EC), Tolerance Violation (TV).
    - The Energy Cost is captured in terms of electricity consumption and CO_2 emission.

- Setpoint Closeness encourages a "business friendly" policy where the RL model attempts to keep the zone temperature as close as possible to the temperature / humidity setpoints, implicitly reducing the risk of violations, but at a higher Energy Cost.

We opt for a "balanced" control policy which maximizes Energy Savings and Setpoint Closeness, while minimizing the risk of Tolerance Violations.

Initial RL Implementation

Temperature / (Rel.) Humidity	> 22.0 degrees	< 22.0 degrees
> 50%	Cool and Re-heat	Heat, Cool and Re-heat
< 50%	Cool and Humidify	Heat and Humidify

The reward value is computed as a measure of the 'effectiveness' of the previous (output) valve openings.

- Time *t*: Input - (Indoor Temp., Indoor Humidity); Output – (Cooling, Heating, Re-heating, Humidifier) valve opening%
- Time *t+1*: Input - (Indoor Temp$_t$, Indoor Humidity$_t$);
- Compare *t* and *t+1* Input/Output values to adapt the Output valve opening% at *t+1*.

```
#invoked every 1 min
def rl_hvac
    (it, ih, pit, pih, pHV, pCV, pUV, pRV):

    #initialize to previous values
    oHV = pHV
    oCV = pCV
    oUV = pUV
    oRV = pRV
    #Heat and Humidify
    if (it < 22.0) and (ih < 50.0):
        if (it < pit):
            inc = hiW + abs((it - pit)/0.1)
            oHV += inc
        if (ih < p_ih):
            inc = uiW + abs((ih - pih)/0.1)
            oUV += inc
        oRV = 0.0
        oCV = 0.0
    # Cool and Re-heat
    if (it > 22.0) and (ih > 50.0):
```

Tipping Point

- The stochastic RL algorithm always starts opening the valves at 0.0%,

- the temperature and (or) humidity deviation from the setpoint then keeps increasing, until the valve opening percentage reaches the tipping point, after which the deviation starts decreasing again until it becomes 0.

- Given this behavior, if we knew the Cooling tipping point beforehand to be 22.3 degrees, we could have opened the Cooling valve earlier to the tipping point - leading to a lower energy cost.

- The caveat here is that the tipping point needs to be estimated properly for all the valves, otherwise opening a valve to more than the tipping point percentage might actually lead to a higher energy cost.

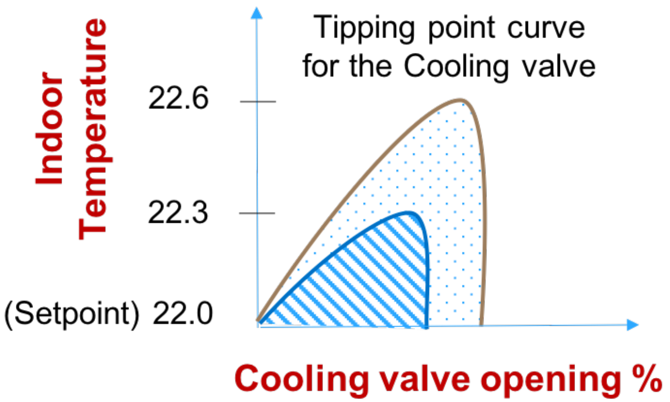

RL Offline Training

1. Take the data generated by the initial RL algorithm

2. apply a filter to extract the data corresponding to the ‘tipping point’ values

3. which are then used as training data,

4. to develop a model to predict the ‘tipping point’ of the valves for each state of the factory zone.

```
#Heat and Humidify
if (it < 22.0) and (ih < 50.0):
    if (oHV == 0.0):
        oHV = h_model.predict(it, ih)
    elif (it < pit):
        inc = hiW + abs((it - pit)/0.1)
        oHV += inc
    if (oUV == 0.0):
        oUV = u_model.predict(it, ih)
    elif (ih < p_ih):
        inc = uiW + abs((ih - pih)/0.1)
        oUV += inc
```

The (offline) trained models are then embedded in our initial RL algorithm.

Benchmarking

Comparing the average valve opening percentages (over 1 week), we can see that all the RL based valve opening percentages are lower; ranging from 10% savings for the Heating valve to almost 45% savings for the Re-heating valve - leading to 25% savings on average.

PID

	Room Temperature ValueY	Room Humidity ValueY	Heating Valve ValueY	Cooling Valve ValueY	Humidifier Valve ValueY	Reheat Valve ValueY
mean	22.027662	50.037363	9.372753	10.498550	30.770681	20.468397
std	0.066739	0.922815	20.184126	14.498391	29.895784	31.726365
min	21.743055	46.518517	0.839120	0.607639	0.000000	0.491898
25%	21.990740	49.479164	1.012731	0.839120	0.000000	0.665509
50%	22.019676	50.086803	1.215278	5.049190	29.300000	1.157407
75%	22.063078	50.665508	4.166667	17.100695	49.960000	34.143517

RL

	Room Temperature ValueY	Room Humidity ValueY	Heating Valve ValueY	Cooling Valve ValueY	Humidifier Valve ValueY	Reheat Valve ValueY
mean	22.009329	49.529951	8.446135	7.832123	25.714461	11.139734
std	0.152894	1.419767	12.799755	9.436951	23.573753	21.848293
min	21.151621	44.531250	0.839120	0.578704	0.000000	0.491898
25%	21.931133	48.726852	0.925926	0.781250	2.157418	0.578704
50%	22.019676	49.469906	1.591435	3.327546	24.203690	0.954861
75%	22.092014	50.173611	11.400463	12.413195	38.818600	9.143518

Conclusions

- Considered the problem of HVAC energy optimization in factories, which has the potential of making a significant environmental impact in terms of energy savings and reduction in CO_2 emissions.
- Outlined a RL based HVAC controller that is able to learn and adapt to real-life factory settings, without the need for any offline training.
- Provided benchmarking results that show the potential to save upto 25% in energy efficiency.

Future work

- Note that we have considered energy savings as proportional to the valve opening percentages (the lower the better). In reality, the energy consumption and CO_2 emissions of the different valves may not be proportional - depending on the underlying mechanism.
- We leave this as future work to adapt the RL logic to accommodate the energy consumption and CO_2 emissions aspects.

Thanks for your attention

Any Questions?

Paper: (link)

Contact: Debmalya Biswas, Linkedin

Powered by