Making an air-source heat pump smart-grid ready

Making an air-source heat pump smart-grid ready

S Thorsteinsson[1], H Cai[2], J D Bendtsen[1], P Heer[2] and J Vivian[2]

[1] Section for Automation and Control, Aalborg University, Denmark, [2] Urban Energy Systems, Swiss Federal Laboratories for Material Science and Technology, Switzerland

E-mail: sith@es.aau.dk

Abstract. In this paper we discuss the differences between the industrial and academic definition of smart-grid readiness of heat pumps and investigate possibilities to bridge this gap such that existing heat pumps, to a higher degree, can participate in a smart-grid or flexibility setting. A weather compensated heat pump has been equipped to enable heating reference tracking by overriding the outdoor air temperature measurement. The reference tracking has been performed using a model predictive controller and a PI-controller. The results show that it is possible to track a heating reference but also that operating the heat pump according to smart-grid objectives has consequences, such as reduced COP and increased defrosting rate. The model predictive and PI-controller showed little difference in performance, meaning that even simple rule-based controllers can be used for tracking. Lastly, the paper provides suggestions for signals that should be available if heat pumps in the future should operate closer to the researchers' expectations of smart-grid readiness.

1. Introduction

In recent years, the heat pump (HP) sales in the European Union have increased significantly from 0.8 million units sold in 2008 to 1.6 million in 2020 [1]. As the number of HPs in the grid increases a number of challenges with respect to grid stability arise, but the HP does not only pose risks to the grid, using the heat capacity in buildings to provide flexibility and other services that supports the grid is an opportunity which is currently being explored. To fully utilize the HPs capacity to support the grid it needs to encompass a certain level of *smartness*, entailing that the controller can receive signals about the situation and respond appropriately.

The definition of *smart* varies depending on the context. From the perspective of the grid operator a smart-grid ready (SGR) HPs' main task is to support grid stability. This view is best reflected in the second version of the smart-grid ready label published by the Bundesverband Wärmepumpen e.V. (bwp) in 2020. Here a SGR labeled HP is demanded to be able to operate in four modes (forced shut-down, normal, intensified heating, forced start-up) [2]. Reviewing academic literature, the definition for a SGR HP goes further. In the literature (e.g. [3–5]) a Model Predictive Control (MPC) algorithm uses electric power consumption P_{HP} as decision variable and relates it to heat using an expression for coefficient of performance (COP). Other decision variables are the compressor speed [6] and the forward temperature [7]. In the literature, the HP is used to support high-level objectives such as reducing operation costs, CO_2-emission or primary energy consumption. In all the references provided the objective is to minimize energy costs using an electricity price signal.

Given the differences of definitions we pose the question: *Is it possible to bridge from the industrial definition of smart-grid ready to the academic one by overwriting the ambient*

temperature input and thereby make the HP follow a heat reference as it is often required in the scientific literature?

In this paper we show a proof of concept that a weather compensated HP can actually be operated according to a heat reference by replacing the outdoor temperature sensor with a control signal. The outdoor temperature is the chosen control signal since it can be replaced without accessing any inner parts of the HP. The heat reference used in this paper is generated by a supervisory controller aimed at minimizing operation costs for space heating. The work was funded by the EUDP project OPSYS 2.0 (Case no.: 64018-0581).

2. Method

This section provides an overview of the control problem and describes the MPC controller.

2.1. Control problem

The control objective is to make the hourly accumulated heat production, Q_{HP}, track an hourly heat reference, Q_{ref}, by replacing the outdoor temperature reading, T_{a}, which is used by the heating curve to determine the supply temperature, with a control signal, $\hat{T}_{\mathrm{a}}$, as shown in Fig. 1. The challenge here is to choose the replacing outdoor temperature value and the flow of the heating system, q, such that the HP acts as if it is designed to track a heat reference. The electric power, P_{HP}, and heat flow, $\dot{Q}_{\mathrm{HP}}$, are included in the figure since they are the actual signal used by the controllers. Note, that in the readings for $\dot{Q}_{\mathrm{HP}}$ and P_{HP} there are artifacts from domestic hot water (DHW) production and defrosting (DF) meaning that it is a requirement that the controller can handle these events. DHW events are recognised as $P_{\mathrm{HP}} > 0$ without heat production ($\dot{Q}_{\mathrm{HP}} = 0$) and DF events are recognised as sharp downwards slopes on P_{HP} and $\dot{Q}_{\mathrm{HP}}$ during heating. In this paper these events have priority over the controller, which is operated in a passive mode during their occurrences. The control architecture, shown in Fig. 2,

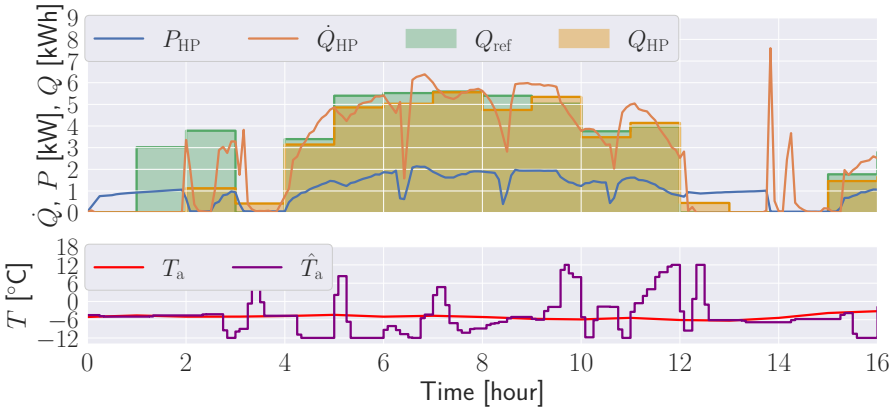

Figure 1. Example of the control input, $\hat{T}_{\mathrm{a}}$, electricity consumption, P_{HP} and heat production, Q_{ref}, against the heat reference, Q_{ref}, for a 16-hour sample period.

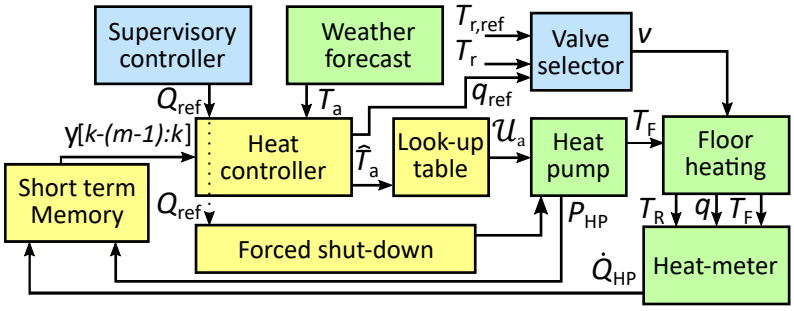

Figure 2. The control diagram with system and upstream controllers.

consists of three components: the *Heat controller*, which provides $\hat{T}_{\mathrm{a}}$ and the flow reference, q_{ref}, a *look-up table* which translate from $\hat{T}_{\mathrm{a}}$ to a voltage, $\mathcal{U}_{\mathrm{a}}$, which mimics the temperature reading,

and a *Short-term memory* which holds the M latest measurement points. The *Heat controller* is either a simple PI-controller, with limits and anti-windup, or MPC. The coefficients of the PI-controller are seen in Table 1. The *Supervisory controller* is the upstream controller which provides the heat reference and the *Valve selector* determines the flow. Both blue modules are described in [8] which describes the *Supervisory controller*.

The controller performance is evaluated using the root mean square error (RMSE) and normalized RMSE (nRMSE) of the tracking error, $e = Q_{\mathrm{HP}} - Q_{\mathrm{ref}}$. Only data-points where the HP was heating at least 50 minutes of the hour are included, since DHW production otherwise creates too much noise in the results. The nRMSE is normalized using $Q_{\mathrm{HP,max}}$-$Q_{\mathrm{HP,min}}$.

2.2. Model predictive controller

This section presents the MPC used to track Q_{ref} provided by the *Supervisory controller*. Note that this is a low-level controller in the hierarchy, therefore it does not operate according to high-level objectives such as minimizing costs or CO_2-emissions. The model seeks to emulate the internal controller dynamics of the HP and not the dynamics of the refrigeration-cycle.

The model used for the MPC algorithm is a discrete full state feedback ($y[j] = x[j]$) linear state space system, $x[j+1] = \mathbf{A}x[j] + \mathbf{B}u[j] + \mathbf{E}d[j]$, with P_{HP} and $\dot{Q}_{\mathrm{HP}}$ being states. The control inputs are $\hat{T}_{\mathrm{a}}$ and q_{ref} with T_{a} being a measured disturbance.

The method chosen for the model identification is Dynamic Mode Decomposition with Control (DMDc), which is a black-box method that provides a linear discrete state-space model [9]. DMDc is chosen for the following three reasons: firstly, a discrete linear model is ideal for MPC, second, since information on the internal controllers and states are unavailable, a black-box method based on input-output data is preferred and, third, the method naturally incorporates multiple time-series of input-output data. This is relevant because data-sets where $\hat{T}_{\mathrm{a}} \neq T_{\mathrm{a}}$ are needed since they are both inputs to the system. Using multiple short data-sets spread in time disturbs the occupants less, than performing a single long excited time series measurement. The DMDc method is well described by Proctor et al [9].

The optimization problem seeks to deliver the heat requested by the supervisory controller with a minimum rate of change in electric consumption, P_{HP}, in order to reduce volatility in heat supply and obtain a smooth transition between hourly demanded heating loads.

$$\min_{u_{\mathrm{c}}} J(u_{\mathrm{c}}) = \min_{u_{\mathrm{c}}} \|\Delta P_{\mathrm{HP}}\|_2^2 + c_s s^{\mathrm{T}} s \tag{1a}$$

s.t.

$$\Delta P_{\mathrm{HP}}[j] = P_{\mathrm{HP}}[j] - P_{\mathrm{HP}}[j-1] \qquad P_{\mathrm{HP}}[j-1] = P_{\mathrm{HP},-1} \qquad j \in \{0,\dots,n-1\} \tag{1b}$$

$$\hat{Q}_{\mathrm{ref}}[i] = \begin{cases} \sum_{j=iR_Q}^{iR_Q+M_i} \frac{\Delta T_{\mathrm{S}}}{\Delta T_{\mathrm{L}}} \dot{Q}_{\mathrm{HP}}[j] + s_Q[i] & Q_{\mathrm{ref}}[i] > 0 \\ 0 & Q_{\mathrm{ref}}[i] = 0 \end{cases} \qquad i \in \{0,\dots,N-1\} \tag{1c}$$

$$x[j+1] = \mathbf{A}x[j] + \mathbf{B}u[j] + \mathbf{E}d[j] \qquad x[0] = x(k) \tag{1d}$$

$$u_{\mathrm{c}}[\ell] = u\left[\ell R_{\mathrm{c}}\right] = u\left[\ell R_{\mathrm{c}} + 1\right] = \cdots = u\left[\ell R_{\mathrm{c}} + (R_{\mathrm{c}} - 1)\right] \qquad \ell \in \left\{0,\dots,\frac{n}{R_{\mathrm{c}}} - 1\right\} \tag{1e}$$

$$x_{\min} - s_{\min}[j] \leq x[j] \leq x_{\max} + s_{\max}[j], \;\; u_{\min} \leq u[j] \leq u_{\max}, \;\; \Delta T_{\min} \leq \Delta \hat{T}_{\mathrm{a}}[j] \leq \Delta T_{\max} \tag{1f}$$

$$x[j] = \begin{bmatrix} \dot{Q}_{\mathrm{HP}}[j] & P_{\mathrm{HP}}[j] \end{bmatrix}^{\mathrm{T}} \qquad u[j] = \begin{bmatrix} \hat{T}_{\mathrm{a}}[j] & q[j] \end{bmatrix}^{\mathrm{T}} \qquad d[j] = \begin{bmatrix} T_{\mathrm{a}}[j] \end{bmatrix}^{\mathrm{T}} \tag{1g}$$

The cost-function consists of a term which minimizes the rate of change of P_{HP} [kW] and a slack term. Constraint (1b) calculates the rate of change of P_{HP}. Constraint (1c) forces the HP to deliver $\dot{Q}_{\mathrm{HP}}$ [kW] as requested by $\hat{Q}_{\mathrm{ref}}$ [kW h] for hour i with the left side of the equation summing the energy contributions over each hour. Constraint (1d) defines the dynamics with initialization of state and (1e) defines a zero-order-hold on the control signal in the intermediate simulation steps between control periods. The simulation period is ΔT_{S}, control period is ΔT_{c} and the heat reference period is ΔT_{L}. The sample ratios are: $R_Q = \frac{\Delta T_{\mathrm{L}}}{\Delta T_{\mathrm{S}}} \in \mathbb{N}$, $R_{\mathrm{c}} = \frac{\Delta T_{\mathrm{c}}}{\Delta T_{\mathrm{S}}} \in \mathbb{N}$.

Note, it is required that the ratios are natural numbers. The constraints (1f) limits the state x, and control inputs, u, and limits rate of change of the input, $\Delta\hat{T}_\mathrm{a}$. $\Delta\hat{T}_\mathrm{a}$ is calculated according to (1b). The expressions in (1g) present the content of x, u and d vectors. Since the control period is 15 minutes and the reference period is 1 hour the control window is not always aligned with full hours creating incomplete control hours in the beginning and end of the reference vector. This misalignment is handled by the partial functions in (2).

$$M_i = \begin{cases} ((R_Q - 1) - r & i = 0 \\ R_Q - 1 & 0 < i < N - 1 \\ r - 1 & i = N - 1 \end{cases}, \; \hat{Q}_\mathrm{ref}[i] = \begin{cases} \max(Q_\mathrm{ref}[0] - Q_\mathrm{del}, 0) & i = 0 \\ Q_\mathrm{ref}[i] & 0 < i < N - 1 \\ Q_\mathrm{ref}[N-1]\left(1 - \frac{R_Q - r}{R_Q}\right) & i = N - 1 \end{cases} \tag{2}$$

The reference $Q_\mathrm{ref}[0]$ is the heat that remains to be delivered in the current hour and $N = \left\lfloor \frac{n}{R_Q} \right\rfloor$ is the length of the reference vector with n being the length of the simulation vector and r the number of simulation steps into the hour.

2.3. Experiment description

The evaluation test spanning the period from 2022-11-07 to 2023-03-07 was carried out in an inhabited single-familyhouse. The test period contains four sub-periods; two periods with PI-control and two with MPC. The results from the tracking experiment is compared to a benchmark data-set from the period 2021-11-01 to 2022-03-01 where the HP was operated using the controller provided by the manufacturer. Table 1 shows the periods for operation together with coefficients, ranges and weather indicators for each period. The range for $\hat{T}_\mathrm{a}$ and its change between consecutive steps is shown in the fifth and sixth column of Table 1, respectively. The average outdoor temperature for each period is given as $\overline{T}_\mathrm{a}$ and minimum and maximum as T_min and T_max, respectively. The column $\overline{E}_\mathrm{PV}$ is the average daily electricity production from the solar panel, which is here used as an indicator of solar radiation intensity. Note that PI1 and MPC1 are more aggressively tuned towards tracking performance than PI2 and MPC2. The update periods are $\Delta T_\mathrm{L} = 60\mathrm{min}$, $\Delta T_\mathrm{c} = 15\mathrm{min}$ and $\Delta T_\mathrm{S} = 5\mathrm{min}$. The limits of the HP are: $Q_\mathrm{HP,min} = 1.5\mathrm{kWh}$ and $Q_\mathrm{HP,max} = 4.5\mathrm{kWh}$. The readers are referred to [8] for more details.

Table 1. Controller and experiment indicators.

Ctrl	Period	P	I	$\hat{T}_\mathrm{a}$	$\Delta\hat{T}_\mathrm{a}$	$\overline{T}_\mathrm{a}$	T_min	T_max	$\overline{E}_\mathrm{PV}$
MPC1	08/12-22 - 18/01-23	-	-	[-12, 12]	-	2.6	-5.9	8.9	4.8
MPC2	18/01-23 - 15/02-23	-	-	[-12, 12]	[-5, 10]	2.3	-2.2	8.4	7.3
PI1	07/11-22 - 25/11-22	5.3e-03	2.2e-06	[-15, 7]	-	5.0	-0.1	12.9	7.6
PI2	15/02-23 - 06/03-23	3.2e-03	7.4e-07	[-12, 6]	-	4.6	1.1	8.4	26.2

3. Results

RMSE and nRMSE of the heat reference signal's tracking are shown in Table 2. Here it can be seen that PI1 and MPC1 have the best performance. For the case with PI1, it is most likely due to a more aggressive tuning of the parameters as well as a broader input range, compared to PI2. For MPC1, the reason is that there is no limit to the rate of change of $\hat{T}_\mathrm{a}$. The corresponding distribution of the error between $Q_\mathrm{ref}[i]$ and $Q_\mathrm{HP}[i]$ is shown in Fig. 3. The distributions are similar, as in both cases MPC tends to overshoot less than PI. A visual example of the tracking is given in Fig. 1. Fig. 4 presents the frequency of defrosting cycles for each type of controller as a function of the outdoor air temperature. Here it is clear that the MPC algorithm causes more defrost events than the benchmark controller. Where the benchmark and MPC shows a

linear trend, results of the PI-controller are more scattered which is likely caused by the fewer data-points. It is therefore difficult to conclude on the actual defrosting rate of the PI-controller, but it seems lower than the MPC and higher than the benchmark. The higher DF frequency is likely caused by Q_{ref} demanding higher heating loads. Fig. 5 presents the relation between hourly accumulated electricity consumption, E_{HP}, and the COP for that hour. Here it is also clear that the HP has been operated at higher heating loads since the controller data points (scattered dots and triangles) move further out the x-axis than the benchmark data (contours). Furthermore, the COP drops as expected at higher loads and lower temperatures. The plot also reveals that, unsurprisingly, hours with DF events have a significantly lower COP.

Table 2. Root mean square error achieved during heat reference tracking.

Ctrl	Color	Period	#DP	RMSE	nRMSE
MPC1	●	08/12-22 - 15/01-23	337	0.774	0.172
MPC2	●	15/01-23 - 15/02-23	315	1.024	0.228
PI1	●	07/11-22 - 25/11-22	73	0.784	0.174
PI2	●	15/02-23 - 06/03-23	87	1.163	0.259

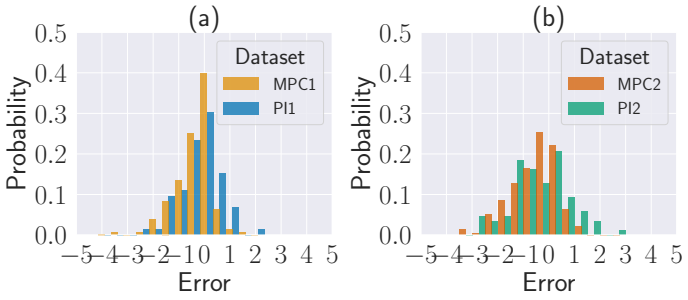

Figure 3. Sub-figure (a) and (b) shows the error distribution of the tracking.

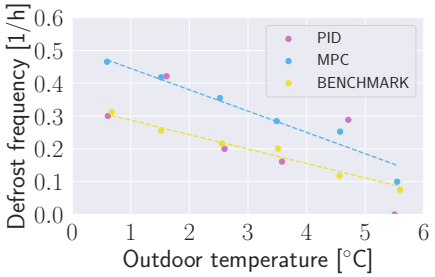

Figure 4. Shows the defrost frequency as a function of controller and outdoor temperature.

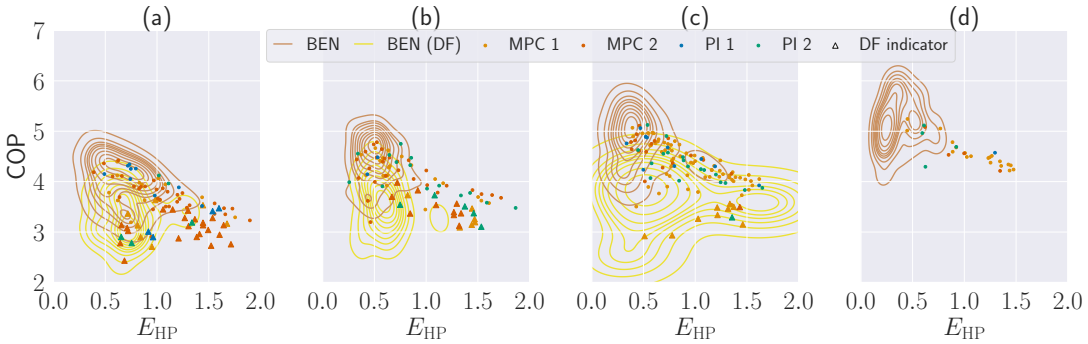

Figure 5. Hourly COP with respect to the hourly accumulated electricity consumption, E_{HP}, for four outdoor temperature intervals: (a) T_{a}: [1.0,2.0), (b) T_{a}: [3.0,4.0), (c) T_{a}: [5.0,6.0) and (d) T_{a}: [7.0,8.0).

4. Discussion

Results have shown that a HP can be operated according to a heat reference, but they also reveal certain efficiency drawbacks on COP and DF frequency which need to be taken into account.

The high-level controller was able to reduce space heating costs by 7.8% despite the reduction in efficiency of the HP presented here [8]. This is possible since the focus shifts from maximizing COP to high-level objectives such as reducing costs or CO_2-emissions which are factors closer to the concerns of society. Here it is important to stretch that the *Heating controller* only aims at delivering the requested heat, not at maximizing efficiency. Therefore, efficiency losses of the HP when operated according to high-level objectives are intrinsic and a proper analysis of the benefits of such operation are required. In this case the prices were volatile, allowing the *Supervisory controller* to benefit thereof. However, it is not certain that all situations provide the proper conditions for covering the drawbacks on efficiency. Although in this project, signals such as P_{HP} and $\dot{Q}_{\text{HP}}$ were measured using externally installed equipment, it is our belief that making these signals available by default can be a promising path forward for the sector. Further, we have noticed that signals such as state of charge of the domestic hot water tank, measured outdoor air temperature and information on defrost state could be beneficial for any controller which seeks to exploit the grid supportive capabilities of the HP. Further, having reference tracking build into the HP is far more efficient than providing a retrofit solution.

5. Conclusion

In this paper we showed that a weather compensated HP can be adapted to track a heat reference by overwriting the outdoor air temperature signal, thus enabling it perform according to the requirements provided in literature on MPC control of buildings. Although this improves flexibility capabilities of the HP and adds an extra layer on top of the smart-grid ready label's definition, the HP efficiency can be significantly impaired when operated according to high level objectives. Future work needs to find a trade-off between the benefit of dynamic operation as presented in the paper and the decrease of HP efficiency.

References

[1] Thomas Novak 2021 *The REHVA European HVAC Journal* **58** 40–43 URL https://www.rehva.eu/rehva-journal/chapter/european-heat-pump-market

[2] Bundesverband Wärmepumpen eV 2020 Regulations for the "SG ready" label for electrically driven heating and hot water heat pumps and interface-compatible system components

[3] Gonzato S, Chimento J, O'Dwyer E, Bustos-Turu G, Acha S and Shah N 2019 *Energy and Buildings* **202** 109421 ISSN 0378-7788 URL https://www.sciencedirect.com/science/article/pii/S0378778819307042

[4] Rastegarpour S, Scattolini R and Ferrarini L 2021 *Journal of Process Control* **99** 69–78 ISSN 0959-1524 URL https://www.sciencedirect.com/science/article/pii/S0959152421000068

[5] Ferrarini L, Rastegarpour S and Caseri L 2020 *IEEE Robotics and Automation Letters* **5** 5363–5369 ISSN 2377-3766 conference Name: IEEE Robotics and Automation Letters

[6] Lee Z, Gupta K, Kircher K J and Zhang K M 2019 *Energy and Buildings* **198** 75–83 ISSN 03787788 URL https://linkinghub.elsevier.com/retrieve/pii/S0378778819302609

[7] Tahersima F, Stoustrup J, Rasmussen H and Meybodi S A 2012 *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)* pp 7583–7588 iSSN: 0743-1546

[8] Thorsteinsson S, Kalaee A A S, Vogler-Finck P, Stærmose H L, Katic I and Bendtsen J D 2023 Long-term experimental study of price responsive predictive control in a real occupied single-family house with heat pump arXiv:2303.16289 [cs, eess]

[9] Proctor J L, Brunton S L and Kutz J N 2016 *SIAM Journal on Applied Dynamical Systems* **15** 142–161 publisher: Society for Industrial and Applied Mathematics URL https://epubs.siam.org/doi/10.1137/15M1013857