

# Frequency regulation with heat pumps using robust MPC with affine policies

Felix Bünning<sup>\*,\*\*</sup>, Joseph Warrington<sup>\*\*</sup>, Philipp Heer<sup>\*</sup>,  
Roy S. Smith<sup>\*\*</sup>, John Lygeros<sup>\*\*</sup>

<sup>\*</sup>Urban Energy Systems Laboratory, Swiss Federal Laboratories for  
Materials Science and Technology, Empa, Dübendorf, Switzerland

<sup>\*\*</sup>Automatic Control Laboratory, Department of Electrical Engineering  
and Information Technology, ETH Zürich, Switzerland

---

**Abstract:** The increase in the renewable energy sources connected to the electricity grid has resulted in an increased need for frequency regulation. On the demand side, frequency regulation services can be provided by electrified heating/cooling systems exploiting the energy stored in thermal mass of buildings. To provide such services a first principles model of the building is needed, which is often difficult to obtain in practice. This issue can be overcome by using a buffer storage between the heating/cooling source and the building. Here, we present a solution based on robust optimization to offer frequency regulation reserves with such a system comprising a heat pump, a thermal storage in the form of a warm water buffer tank, and heating demand from a building that needs to be served. We mitigate the problem of limited thermal storage by introducing affine policies on uncertain variables. In three experiments with a real heat pump and warm water buffer storage and an emulated heating demand, we demonstrate that the system can indeed offer reserves and can successfully track a regulation signal while meeting the heating demand at all times.

*Keywords:* Frequency reserves, Robust Model Predictive Control, Frequency regulation, Buildings

---

## 1. INTRODUCTION

With an increasing amount of renewable energy sources connected to the electricity grid, the need for frequency regulation increases. Besides regulation on the supply side, for example with fast-reacting gas power plants, there is increasing interest in regulation on the consumer side. This concept is commonly referred to as demand-side management. Buildings equipped with electric heating or cooling systems are candidates for demand-side management: they can influence the frequency in the grid as they run on grid electricity and they have flexibility on the generation of heating or cooling energy due to the thermal inertia of the building (Fischer and Madani, 2017).

Zhang et al. (2017) and Vrettos et al. (2016) investigate control schemes based on Robust Model Predictive Control (MPC) to provide day-ahead reserves for frequency regulation with commercial buildings and electrified heating and cooling systems. In (Vrettos et al., 2018b,a) this concept was adapted and successfully tested on a real case study of a single small building equipped with a chiller. A potential drawback of this approach is the need for a first-principles model of the thermal dynamics of the building. Some authors argue that the costs associated with the development and maintenance of such models might be prohibitive for

wide-spread commercial application of MPC in buildings (Sturzenegger et al., 2016; Jain et al., 2018).

This issue could be overcome by decoupling the thermal demand of the building and the supply of the heating/cooling system with a thermal buffer storage, for example a warm water tank. In this case, the buffer storage gives rise to the flexibility in heating/cooling energy production instead of the building itself. The electrified heater/cooler and storage are modelled with first principles, while the demand of the building can be modelled with any forecasting technique (for example (Bünning et al., 2020)).

Here, we adapt the robust MPC approach of Vrettos et al. (2018b) to a system with a ground source heat pump and a warm water tank, which supply heat for a building. We overcome the problem of the low storage capacity of a buffer compared to the thermal capacity of a building by introducing affine policies, as proposed by Warrington et al. (2012) for reserve operation on power systems. In experiments with a real heat pump and buffer tank and an emulated heating demand, we demonstrate that the heat pump successfully follows a frequency regulation signal and that the water temperature in the buffer tank stays sufficiently warm to serve the building heating demand at all times.

The remainder of the article is structured as follows. In Section 2 we describe the reserve scheme, the device set up under consideration and introduce the control scheme. In

---

\* This research project is financially supported by the Swiss Innovation Agency Innosuisse and is part of the Swiss Competence Center for Energy Research SCCER FEED&D.




Fig. 1. Schematic of the heating system with heat pump, water storage tank and heat exchanger

In Section 3 we present the case study and the particular set up of the controllers. In Section 4 we present and discuss the experimental results. In Section 5 we conclude the article.

## 2. METHODOLOGY

### 2.1 Reserve scheme

In this study we assume a reserve scheme inspired by the US frequency regulation market. At midnight, the reserve provider makes a reserve offer  $r \in \mathbb{R}^{96}$  of symmetric reserves in 15-minute intervals for the next 24 hours to the Transmission System Operator (TSO). During the execution phase, the reserve provider can update their base load  $u_k^0$  every timestep  $k$  (every 15 minutes). It then has to follow the load

$$u_k(\tau) = u_k^0 + w(\tau)r_k, \quad (1)$$

with  $w(\tau) \in [-1, 1]$  denoting the regulation signal which is sent by the TSO every 2 seconds and  $r_k$  denoting the  $k^{\text{th}}$  element of  $r$ . As  $u_k^0$  and  $r_k$  are constant for 15 minutes and  $w(\tau)$  is updated every 2 seconds,  $u_k(\tau)$  also changes every 2 seconds.

### 2.2 Heating system

Figure 1 shows a schematic of the heating system under consideration. The heat pump, depicted on the left, uses ambient heat and electricity  $u(t)$  to generate heat at a higher temperature level  $u_{\text{th}}(t)$ , where  $t$  denotes continuous time. The conversion efficiency is described by the coefficient of performance  $\alpha_{\text{COP}}$ :

$$u_{\text{th}}(t) = \alpha_{\text{COP}} u(t) + e. \quad (2)$$

Here,  $e$  denotes an error caused by the assumption of a constant COP. Instead of modelling dependencies of  $e$  (for example on the ambient temperature) explicitly, they will be modelled by an uncertainty set in the following. Note that equation (2) is valid for any  $u$  and  $u_{\text{th}}$ , thus also for  $u_k(\tau)$  and the corresponding  $u_{k,\text{th}}(\tau)$ . The storage tank in the middle is charged by pumping warm water from the outlet of the heat pump to the top of the tank. It is discharged by pumping water from the top through the heat exchanger, which serves the heating demand of the building  $v(t)$ . The average temperature in the storage tank  $x(t)$  can be described by the differential equation

$$m c_p \frac{dx(t)}{dt} = u_{\text{th}}(t) - v(t) + \delta, \quad (3)$$




Fig. 2. Control scheme

where  $m$  denotes the mass, and  $c_p$  denotes the specific heat capacity of the water, and  $\delta$  denotes the error between the actual and the forecast heating demand of the building. Thermal losses are neglected here.

### 2.3 Control scheme

We use a 3-level control approach, as presented by Vrettos et al. (2018a) and indicated in Figure 2. **Level 1** is a robust optimization scheme which is run once per day at midnight to determine the offered reserves  $r$  in 15 minute intervals,  $r_k$ . **Level 2** is the same optimization scheme which is run at the beginning of each 15 minute interval,  $k$ , in a shrinking horizon MPC fashion to determine the nominal heat pump electrical set points  $u_k^0$  held constant during the entire duration of the 15 minute interval; in this calculation the values of the reserves  $r_j$  for  $j \geq k$  are known. **Level 3** is a feedback controller that controls the heat pump's compressor speed to track  $u_k(\tau) = u_k^0 + w(\tau)r_k$ .

For **Level 1**, equation (3) can be written in discrete linear state space form as

$$x_{k+1} = \tilde{A}x_k + \tilde{B}(u_k - v_k). \quad (4)$$

Following a similar approach to Warrington et al. (2012), by redefining  $x := [x_1, \dots, x_N]^T \in \mathbb{R}^N$ ,  $u := [u_1, \dots, u_N]^T \in \mathbb{R}^N$ ,  $v := [v_1, \dots, v_N]^T \in \mathbb{R}^N$  the state trajectory for horizon  $N$  can be described by

$$x = Ax_0 + B(u - v), \quad (5)$$

where

$$A := \begin{pmatrix} \tilde{A} \\ \tilde{A}^2 \\ \vdots \\ \tilde{A}^N \end{pmatrix} \quad B := \begin{pmatrix} \tilde{B} & 0 & \dots & 0 \\ \tilde{A}\tilde{B} & \tilde{B} & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ \tilde{A}^{N-1}\tilde{B} & \dots & \tilde{A}\tilde{B} & \tilde{B} \end{pmatrix}, \quad (6)$$

and  $x_0$  denotes the initial system state. The robust optimization problem can be formulated as

$$\min_{r, u^0, z} f^{\text{el}^\top} u^0 - f^r{}^\top r + \lambda^\top \epsilon \quad (7a)$$

$$\text{subject to } x = Ax_0 + B(u_{\text{th}} - v + \delta + e), \quad (7b)$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + w \odot r), \quad (7c)$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (7d)$$

$$zU_{\min} \leq u^0 + w \odot r \leq zU_{\max}, \quad (7e)$$

$$z \in \mathbb{Z}_2^N, \quad (7f)$$

$$\epsilon \geq 0, \quad (7g)$$

$$\forall w \in W, \forall \delta \in \Delta, \forall e \in E, \quad (7h)$$



Fig. 3. Experimental set-up with heat pump, water storage tanks and heat exchanger

where  $f^{\text{el}}$  denotes costs for electricity,  $f^r$  denotes payments received for offered reserves,  $X_{\min}$  and  $X_{\max}$  denote lower and upper temperature limits for the storage tank (defined by minimum operating temperature for floor heating and maximum supply temperature of the heat pump),  $\epsilon \in \mathbb{R}^N$  is a slack variable to ensure feasibility and  $\lambda$  the associated cost,  $U_{\min}$  and  $U_{\max}$  denote lower and upper limits on the electrical power of the heat pump,  $z$  is a binary variable that determines whether the heat pump is on or off and  $\odot$  denotes the operator for element-wise multiplication. The constraints have to hold for all uncertainty realizations  $w \in W, \delta \in \Delta, e \in E$ , where the elements of  $w, \delta$  and  $e$  are box-constrained. The robust counterpart of this problem is a Mixed Integer Linear Program (MILP).

For large  $\tilde{B}$ , low  $X_{\max}$  or high  $X_{\min}$  (low storage capacity) and large uncertainty in  $W, \Delta$ , and  $E$ , the offered reserves  $r$  become small or the problem becomes infeasible (if no slack variables are used) because the uncertainty on  $x$  builds up over the prediction horizon, making it difficult to guarantee constraint (7d). To alleviate this problem, we introduce a causal policy that allows the system to react to uncertainties that will have been revealed at the time the decision is implemented, though they are still unknown at the time the optimization is performed. As optimizing over the set of all possible policies is intractable, we focus on affine policies, as discussed in (Warrington et al., 2012). In the case of uncertainties stemming from the regulation signal  $w$  for example, equation (7c) can be changed to

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + w \odot r + D_w w), \quad (8)$$

with  $D_w \in \mathbb{R}^{N \times N}$  being a strictly lower triangular matrix:

$$D_w := \begin{pmatrix} 0 & 0 & \cdots & 0 \\ [D_w]_{2,0} & 0 & \ddots & 0 \\ \vdots & \ddots & \ddots & 0 \\ [D_w]_{N,0} & \cdots & [D_w]_{N,N-1} & 0 \end{pmatrix}. \quad (9)$$

The entries of  $D_w$  become decision variables in the optimization problem. Policies on uncertain variables  $\delta$  and  $e$  can be defined accordingly and can be merged into one matrix, as  $\delta$  and  $e$  appear in the same way in (7b).

To further reduce uncertainty, limits on the integral of  $w(t)$  can be formulated by analysing historical regulation signals (see (Vrettos et al., 2018a)). The uncertainty set can therefore be shrunk to  $\bar{W} \subset W$  for constraint (7c), but not for constraint (7e), as the instantaneous electrical

power of the heat pump still needs to be within operational limits.

The resulting optimization problem is

$$\min_{r, u^0, z, D_w, D_{\delta, e}, \epsilon} f^{\text{el}\top} u^0 - f^r\top r + \lambda\top \epsilon \quad (10a)$$

$$\text{subject to } x = Ax_0 + B(u_{\text{th}} - v + \delta + e), \quad (10b)$$

$$u_{\text{th}} = \alpha_{\text{COP}}(u^0 + \bar{w} \odot r + D_w \bar{w} + D_{\delta, e}(\delta + e)), \quad (10c)$$

$$X_{\min} - \epsilon \leq x \leq X_{\max} + \epsilon, \quad (10d)$$

$$zU_{\min} \leq u^0 + w \odot r + D_w \bar{w} + D_{\delta, e}(\delta + e) \leq zU_{\max}, \quad (10e)$$

$$z \in \mathbb{Z}_2^N, \quad (10f)$$

$$\epsilon \geq 0, \quad (10g)$$

$$[D_w]_{i,j} = 0 \quad \forall j \geq i, \quad (10h)$$

$$[D_{\delta, e}]_{i,j} = 0 \quad \forall j \geq i, \quad (10i)$$

$$\forall w \in W, \forall \bar{w} \in \bar{W}, \forall \delta \in \Delta, \forall e \in E. \quad (10j)$$

The constraint regarding lower and upper heat pump capacity limit, now (10e), of course has to be adapted to ensure feasibility under the chosen policies. For this,  $[D_w]_{i,j}$  and  $[D_{\delta, e}]_{i,j}$  have to be zero whenever the on/off condition  $z_k$  is zero too.

Controller **Level 2** is a shrinking horizon MPC scheme that also uses optimization problem (10). However,  $r$  is now fixed (and with it all elements of  $z$  for which the corresponding element in  $r$  is non-zero). The purpose of Level 2 is to potentially change  $u^0$  depending on changed initial conditions  $x_0$  or updated forecasts on  $v$ .

Controller **Level 3** is a PI controller for tracking equation (1). The control input is the set point for the relative compressor speed of the heat pump  $n_{\text{set}}$ . The measured output is the electrical consumption of the heat pump. Clamping is used for anti-windup when the heat pump operates at speed limits. Moreover, as heat pumps often have ramping constraints, integration is stopped if the difference between  $n_{\text{set}}$  and the measured compressor speed exceeds a threshold.

### 3. CASE STUDY

We test the reserve scheme in experiments on a system comprising a real compressor heat pump and water tank. We conducted three different experiments with an emulated heating demand, a horizon of 2 hours and the regulation signal RegD by PJM.<sup>1</sup>

#### 3.1 Devices and heat demand emulation

Figure 3 shows the devices used. The heat pump is the two-compressor model WP-WW-2NES 20.F4-2-1-S-P100 produced by Viessmann with a maximum thermal capacity of 100 kW. It is fed by ground source heat from a borehole

<sup>1</sup> RegD is the faster regulation signal (compared to RegA) which is provided by PJM, a transmission system operator at the east coast of the USA.




Fig. 4. Emulated heating demand

field. We only use the first compressor stage. The storage consists of two 1100 litre Matica water tanks connected in series.

The right picture shows the heat exchanger that is used for heating demand emulation. It has a maximum thermal capacity of 80 kW. The rejected heat is dumped back into the borehole field. To avoid feedback effects, the experiment time is limited to 2 hours.<sup>2</sup> The heat demand profile that was used during experiment 1 is shown in Figure 4. It is an example of the real measured (and discretized) demand from the NEST building at Empa, Switzerland, which shows typical variations of demand in time. For experiments 2 and 3, the demand was scaled by factors 1.3 and 1.8 respectively.

### 3.2 Controller set up

The parameters that are used for controller levels 1 and 2 are shown in Table 1.  $\alpha_{COP}$ ,  $U_{min}$ ,  $U_{max}$  were determined in preliminary experiments with the heat pump, as well as  $E$ , which is a product of the range of COP and the maximum capacity.  $W$  is a property of the regulation signal used (RegD by PJM, shown in Figure 5). The same holds for  $\bar{W}$ , which was determined by Vrettos et al. (2018a) through analysis of historical RegD signals.  $\tilde{A}$  and  $\tilde{B}$  were determined from the physical properties of the storage tank under the assumption of negligible thermal losses.  $\Delta$  is an approximation of the tracking accuracy of the heating demand emulation.  $x_0$  is initialized by taking a weighted average of six temperature measurements from different water layers in the storage tanks.

For the feedback controller in level 3, the proportional gain  $k_p = 2.0$  and integral gain  $k_i = 0.4$  were determined by pre-tuning on a first-order representation of the heat pump in Simulink<sup>®</sup> and manual adjustments after transferring the controller to the actual plant. The controller output limits of  $n_{set}$  (relative compressor speed) are 20% and 50%<sup>3</sup>. The maximum allowed difference between  $n_{set}$  and the measured speed was set to 2% before anti-windup activates. The sampling time is 500 milliseconds.

The optimization problems for level 1 and 2 were formulated using Matlab<sup>®</sup> with YALMIP (Löfberg, 2012), which automatically derives robust counterparts, and solved using CPLEX (v.12.9). Level 3 was implemented in Python 3. Communication between Matlab and Python was set up through shared csv files. The communication with sensors and actuators of the devices is done via a Python OPC-UA client.

<sup>2</sup> For longer experiments, the temperature of the ground around the boreholes rises such that the supply temperature for the heat pump is affected. As a result, the COP rises.

<sup>3</sup> The heat pump is a two stage model where the second compressor is started if the relative speed exceeds 50%.

Table 1. Parameters for controller level 1 and 2

general	heat pump	storage
$N = 8$ ,	$\alpha_{COP} = 3.45$ ,	$\tilde{A} = 1$ ,
$\lambda = 5$ ,	$E = [-3, 3] kW$ ,	$\tilde{B} = 0.0978 \frac{K}{kW}$ ,
$f^{el} = 1$ ,	$W = [-1, 1]$ ,	$X_{min} = 28^\circ C$ ,
$f^r = 1.2$	$\bar{W} = [-0.25, 0.25]$ ,	$X_{max} = 38^\circ C$ ,
	$U_{min} = 8.2 kW$ ,	$\Delta = [-2.5, 2.5] kW$
	$U_{max} = 12.8 kW$	

## 4. RESULTS AND DISCUSSION

### 4.1 Experiments

Figure 5 shows results of the first experiment. The first graph shows the possible range of the electrical power of the heat pump depending on the set point  $u^0$  and the offered reserves  $r$  in dotted red. The orange line shows the combination of these two variables with the regulation signal (depicted in the third graph), as described in equation (1), which defines the set point for the electrical power of the heat pump. The blue line shows the actual electrical power of the heat pump.

The first thing to note is that level 1 offered reserves of varying size in periods 1 (2.3 kW), 2 (0.8 kW), 3 (2.3 kW), 7 (2.3 kW) and 8 (2.3 kW). The cost function of level 1 is 16% lower compared to the solution of problem (10) without affine policies. The sum of offered reserves is 21% higher. During the reserve periods, the regulation signal is tracked generally very well. In period 4, the heat pump operates at a constant set point, in periods 5 and 6 it is switched off. At the end of period 6, it can be noted that the heat pump is switched on five minutes before the next regulation period starts. This is done as the heat pump takes several minutes before being fully operational after starting. It ensures that the regulation signal can be tracked at the beginning of the next period.

The third graph shows the tracking error in detail during periods where reserves were offered. It can be seen that the error becomes large where high gradients in the set point are present for an extended period of time (e.g middle of period 1, beginning of period 2, middle of period 7). This is due to the compressor's ramping limit. However, the tracking performance is more than sufficient to operate as a reserves provider for PJM. PJM benchmarks reserve providers with an hourly performance score for *accuracy*, *delay* and *precision*. Averaged, these scores give rise to a *composite* score. For the presented experiment, the individual scores for hour 1 and hour 2 are 0.9628, 0.9999, 0.8318 and 0.9656, 1.0000, 0.8369 respectively. The composite scores are 0.9315 and 0.9341 respectively. The minimum requirement is a composite score of 0.75. (For more details on the scoring, please refer to (PJM, 2019)).

Besides being able to track the regulation signal, the control approach needs to ensure that the heating demand can be met at all times. The last graph of Figure 5 shows the average tank temperature in black, the temperature constraints in dotted black and the measurements at different heights of the storage tanks transparent in the background. It can be seen that the average temperature is between the constraints and all individual temperatures are between  $36^\circ C$  and  $28^\circ C$  at all times, which is sufficient for meeting the heating demand.

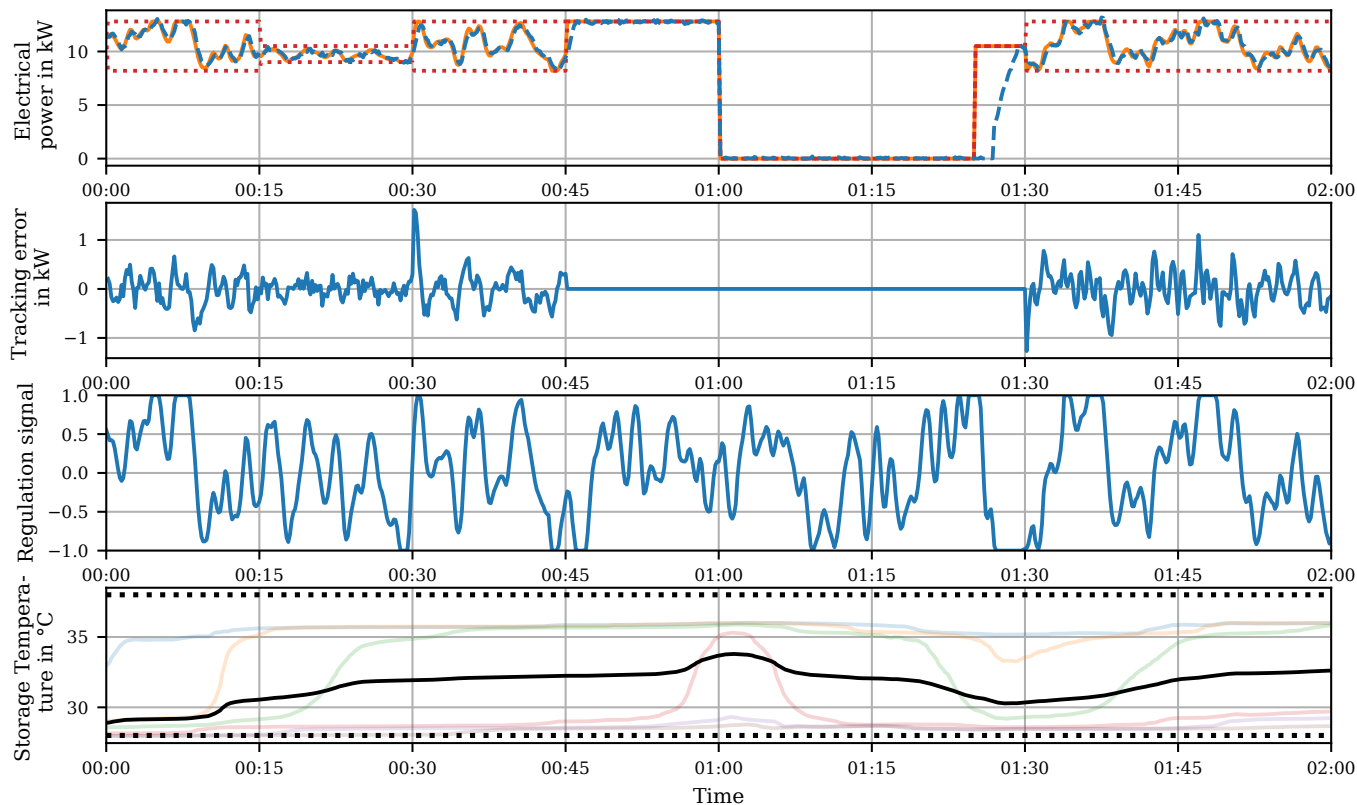


Fig. 5. Results of first experiment. First plot: possible set points according to offered reserves in dotted red, set point for electrical capacity in orange, actual electrical capacity in blue. Second plot: tracking error. Third plot: regulation signal. Fourth plot: temperatures of layers in water tank in different colours.

As discussed in Section 3, the experiment was repeated with different heating demands (scaled by 1.3 and 1.8). The results for tracking performance and storage temperature are shown in Figures 6 and 7. The offered reserves are  $[0,0,1.4,2.3,1.4,2.3,2.3,2.1]$  kW and  $[2.3,0.1,0.7,0.7,0.7,2.0,2.3,2.3]$  kW for experiments 2 and 3 respectively. Compared to the level 1 solution without affine policies, the cost functions decrease by 3% and 13% respectively. The sum of reserves is slightly decreased by 2% in experiment 2 but increased by 102% in experiment 3 compared to when no recourse is used. The approach with affine policies would further benefit from longer prediction horizons, as the build up of uncertainties in the case where no affine policies are used, is not as prominent for  $N=8$  (compared to  $N=96$ ).

The tracking performance is satisfactory in both experiments. Moreover, the average storage temperature stays within limits in both experiments. The performance scores of all three experiments are shown in Table 2. The lower precision performance in hour 1 of the third experiment can be explained by the small value of offered reserves (0.1 kW) in the second reserve period. This leads to a high relative tracking error. The problem could be omitted by imposing a lower constraint on  $r$ . However, as the scores are well above the limit of 0.75, this would be optional. Generally, the good scores could also be used to tune the feedback controller in level 3 to be less aggressive. While this would lower the performance scores, it would lessen the mechanical stress on the compressor of the heat pump.

Table 2. PJM performance scores for all experiments

score:	accuracy		delay		precision		composite	
hour:	1	2	1	2	1	2	1	2
exp 1	0.96	0.97	1.00	1.00	0.83	0.84	<b>0.93</b>	<b>0.93</b>
exp 2	0.89	0.97	1.00	1.00	0.80	0.81	<b>0.90</b>	<b>0.93</b>
exp 3	0.80	0.96	1.00	1.00	0.69	0.82	<b>0.83</b>	<b>0.93</b>

## 5. CONCLUSION

By using a buffer storage between an electrified heating/cooling source and a building, frequency regulation reserves can be provided without the need for a first-principles model of the building. In this study, we have shown a control approach based on robust optimization to control such a system. We have introduced affine policies on uncertain variables to overcome the problem of low storage capacity of a buffer tank compared to the thermal mass of a building. In three experiments with a real heat pump, real water storage and an emulated heating demand, we have demonstrated that the approach can offer reserves, track a regulation signal with sufficient accuracy and meet the energy demand at all times. Ongoing research focusses on combining the approach with demand forecasting to conduct experiments with a real heating demand of a building and a full prediction horizon of one day.

## ACKNOWLEDGEMENTS

We would like to thank Reto Fricker, Ralf Knechtle and Sascha Stoller for sharing their expertise and help with the implementation. Moreover, We would like to thank Kristina Orehounig and Viktor Dorer as well as




Fig. 6. Results of second experiment

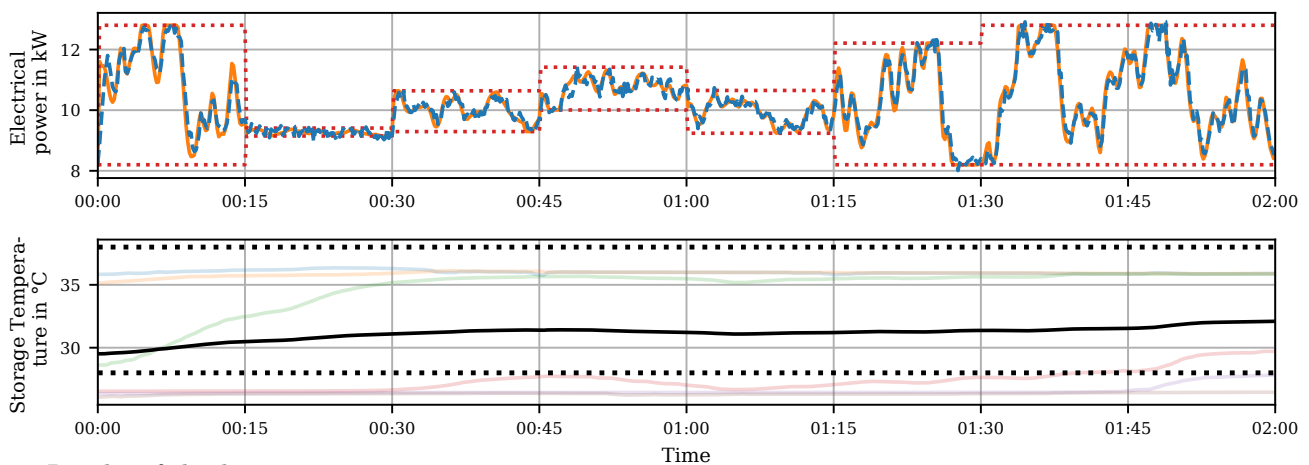


Fig. 7. Results of third experiment

Ahmed Aboudonia, Annika Eichler, Benjamin Flamm, Marc Hohmann, Mathias Hudoba de Badyn, Andrea Iannelli, Mohammad Khosravi, Anil Parsi and Suli Zou for fruitful discussions.

This research project is financially supported by the Swiss Innovation Agency Innosuisse and is part of the Swiss Competence Center for Energy Research SCCER FEED&D.

#### REFERENCES

- Bünning, F., Heer, P., Smith, R.S., and Lygeros, J. (2020). Improved day ahead heating demand forecasting by online correction methods. *Energy and Buildings*, 211, 109821.
- Fischer, D. and Madani, H. (2017). On heat pumps in smart grids: A review. *Renewable and Sustainable Energy Reviews*, 70, 342–357.
- Jain, A., Smarra, F., Behl, M., and Mangharam, R. (2018). Data-Driven Model Predictive Control with Regression Trees—An Application to Building Energy Management. *ACM Transactions on Cyber-Physical Systems*, 2(1), 1–21.
- Löfberg, J. (2012). Automatic robust convex programming. *Optimization Methods and Software*, 27(1), 115–129.
- PJM (2019). PJM Manual 12: Balancing Operations. Technical report.
- Sturzenegger, D., Gyalistras, D., Morari, M., and Smith, R.S. (2016). Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost–Benefit Analysis. *IEEE Transactions on Control Systems Technology*, 24(1), 1–12.
- Vrettos, E., Kara, E.C., MacDonald, J., Andersson, G., and Callaway, D.S. (2018a). Experimental Demonstration of Frequency Regulation by Commercial Buildings—Part I: Modeling and Hierarchical Control Design. *IEEE Transactions on Smart Grid*, 9(4), 3213–3223.
- Vrettos, E., Kara, E.C., MacDonald, J., Andersson, G., and Callaway, D.S. (2018b). Experimental Demonstration of Frequency Regulation by Commercial Buildings—Part II: Results and Performance Evaluation. *IEEE Transactions on Smart Grid*, 9(4), 3224–3234.
- Vrettos, E., Oldewurtel, F., and Andersson, G. (2016). Robust Energy-Constrained Frequency Reserves from Aggregations of Commercial Buildings. *IEEE Transactions on Power Systems*, 31(6), 4272–4285.
- Warrington, J., Goulart, P.J., Mariethoz, S., and Morari, M. (2012). Robust reserve operation in power systems using affine policies. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 1111–1117. IEEE.
- Zhang, X., Kamgarpour, M., Georghiou, A., Goulart, P., and Lygeros, J. (2017). Robust optimal control with adjustable uncertainty sets. *Automatica*, 75, 249–259.