Field Demonstration of Predictive Heat Pump Water Heater Control

Levi D. REYES PREMER[1,2*], Leo SEMMELMANN[3], Elias N. PERGANTIS[1,2], Echkard A. GROLL[1,2], Davide ZIVIANI[1,2], Kevin J. KIRCHER[1,2],

[1] Ray W. Herrick Laboratories, School of Mechanical Engineering, Purdue University, West Lafayette, IN 47907-2099, USA

[2] Center for High Performance Buildings, Ray W. Herrick Laboratories, Purdue University West Lafayette, IN 47907-2099, USA

[3] Karlsruher Institut für Technologie (KIT) Kaiserstraße 12, 76131 Karlsruhe, Germany

[*] Corresponding Author: lreyespr@purdue.edu

ABSTRACT

Heat pump water heaters (HPWHs) could significantly reduce energy costs and greenhouse gas emissions from water heating, the second largest energy use in residential buildings. Today, most HPWHs use electric resistance heating elements to maintain comfortable water temperatures even during large water draws. Unfortunately, heating elements significantly decrease energy efficiency, and their current and voltage requirements may necessitate costly electrical work in older homes. This paper develops and field-tests a model predictive control (MPC) system that enables a HPWH with no heating elements to maintain comfort at high efficiency. By contrast to most prior experimental studies on water heater MPC, which often use perfectly-forecasted water draws in controlled laboratory settings, this paper reports field tests from a real home with three full-time occupants. The occupants' water draws are forecasted using a machine learning model and a scalable training methodology. This paper also presents occupant feedback on thermal comfort, as well as an Internet of Things infrastructure that enables real-time data acquisition and control. In the MPC formulation, the energy savings were 11% with the same thermal comfort as the manufacturer's constant set-point control. An adjusted MPC formulation substantially improved thermal comfort while modestly increasing energy costs.

1. INTRODUCTION

In 2023, residential buildings caused 18% of United States primary energy consumption (EIA, 2023). Improving residential energy efficiency is therefore crucial to climate goals, such as the 80% emission reduction from 2022 levels by 2035 that the International Energy Agency estimates is necessary from developed nations (IEA, 2023). Electrifying appliances is a necessary step in this direction for sectors such as water heating, the second-largest energy end-use in the residential sector (EIA, 2020). Heat pump water heaters (HPWHs) – a system that uses electricity to move heat from ambient conditions into water tanks – typically use two to three times less energy than electric resistance water heaters (Priyadarshan et al., 2024). HPWHs could therefore play a significant role in achieving emission reduction targets. With increasing federal and state incentives, the adoption of heat pump water heaters is projected to grow significantly in the coming years (IEA, 2022).

Although HPWHs are much more efficient than water heaters that use natural gas electric resistance heating, a HPWH's energy use and thermal comfort performance depend on its control strategy. Today, most HPWH control strategies maintain a constant water temperature set-point with a lower dead band to determine when to initiate heating. This is a robust strategy for comfort when there is a sufficient heat transfer rate from the system; however, because heat pumps in HPWHs are limited by their heat transfer rate, these systems often use one or more electric resistance heating elements when the heat pump's heat capacity is not sufficient to meet thermal comfort. This operation mode, therefore, decreases the overall efficiency of the HPWH as the resistance element has a maximum coefficient of performance (COP) of 1 while the heat pump can reach a COP of up to 5 (Sparn et al., 2014). The constant set-point operation also results in the water heater staying thermally charged at all non-hot water draw periods since there is no insight into hot water usage. This results in additional heat losses through the tank walls.

Recent research has explored improving the control strategy in HPWHs through variations of model-predictive control

(MPC). There have been simulation studies performed to minimize operational costs through economic MPC (dela Rosa et al., 2023), minimize carbon emissions (Mande et al., 2022), and to explore data-driven water heater control (Heidari et al., 2020). These studies provide important control designs and simulation evaluations; however, simulated control performance tends to differ from experiment and field test results and also lacks implementation difficulties (Drgoňa et al., 2020). There are a set of experimental studies on HPWHs that have been performed as well. These studies have aimed to enhance thermal comfort and reduce costs (Jin et al., 2014), and mitigate peak loads (Starke et al., 2020) using a two-node or single-node water heater model. All of these studies are solely considering simulated water usage data, some of which have data from real homes. While these approaches can reasonably replicate specific water usage patterns in households, they may not fully capture the dynamic nature of household behaviors. For instance, unpredictable events like showering at unusual times are challenging to replicate and may not be ideal for experimental data due to the complexities they introduce for controllers (Khabbazi et al., 2024). These controllers are also developed for the water heater in hybrid mode where the electric resistance elements can be used. There has been an experimental effort to reduce operational costs in a HPWH using only the heat pump; however, this work was also conducted in a laboratory environment with simulated water usage (Baumann et al., 2023). Additionally, in these workings, there is little to no mention of the water forecast method and performance. There has been work to improve electric water heater (EWH) performance through predictive control using a Prophet machine learning model for the forecast trained from 77 EWH's over 150 days (Shen et al., 2021); however, the suitability of this model for a single-family home has not been demonstrated. Training the model on data from a single-family home with just one water heater presents significantly fewer data points, raising a concern about the model's effectiveness in this context.

This paper uses MPC to reduce costs and maintain comfort with a hot water forecast using a Prophet model (Shen et al., 2021). The Prophet model is trained on a month's worth of data on one HPWH, representing a more realistic training method for single-family homes. The controller developed aims to improve the control strategies for the HPWH in pure heat pump operation without an electric resistance heating element to ideally maintain higher efficiencies while providing acceptable thermal comfort. This controller is implemented in the DC Nanogrid House at Purdue University for 7 days. This field test provides valuable insight into unique challenges presented by real occupant behavior, human feedback, and limitations in sensing equipment and IoT communication (Pergantis, Premer, et al., 2024). This work has impacts beyond improving HPWH control by enabling more dynamic load shifting in a home.

2. TEST SITE

The field demonstrations in this paper took place in the DC Nanogrid House, pictured in Fig. 1. The DC Nanogrid House is a 208 m^2, two-story, 1920s-era detached single-family home near Purdue University's campus in West Lafayette, Indiana, USA. The home has three occupants who live year-round. The hot water in the house is produced via a Rheem 50-gallon heat pump water heater equipped with two electric resistance heating elements supplied by AC power. The water heater is also equipped with two thermistors on the upper and lower parts of the tank used in the default controls. The water heater has five different operational modes: electric mode, high demand mode, hybrid mode, power savings mode, and heat pump mode. The water heater has a uniform energy factor (UEF) of 3.75. The UEF is the industry standard for measuring water heater efficiency and for most HPWHs, is evaluated on the performance of the HPWH in hybrid mode (Star, 2023). It is running in hybrid mode because this mode enables the electric resistance elements when water temperatures become uncomfortable. Running the water heater in heat pump mode saves the occupants energy costs; however, at the expense of occupants being more susceptible to discomfort conditions. This work enables a predictive controller on the water heater in heat pump mode to evaluate the cost reduction and its ability to meet comfort needs.

One of the efforts of this work is to keep the scalability of the control solution in mind. Minimal sensing equipment was added to ensure the water heater's observability. Through the API connection with the water heater, set-point temperatures and mode changes can be updated to the water heater. Due to the older model of the water heater, there are no sensor readings available through the API connection, such as the thermistors on the water heater. Thermocouples were therefore placed next to the manufacturers' thermistors to mirror the measurements that would be available to the manufacturer and the newer models of the water heater through API access. The only additional sensing equipment that is not used by the manufacturer is a flow sensor on the hot water line, a thermocouple on the inlet water line to the tank, and a data acquisition system (DAQ), that could be replaced by a raspberry pi. The flow sensor enables hot water usage data that can then be used to predict future patterns. The thermocouple on the inlet line is used to measure the heat loss from entering water. The DAQ pulls these measurements and the data is then fed by a raspberry pi to

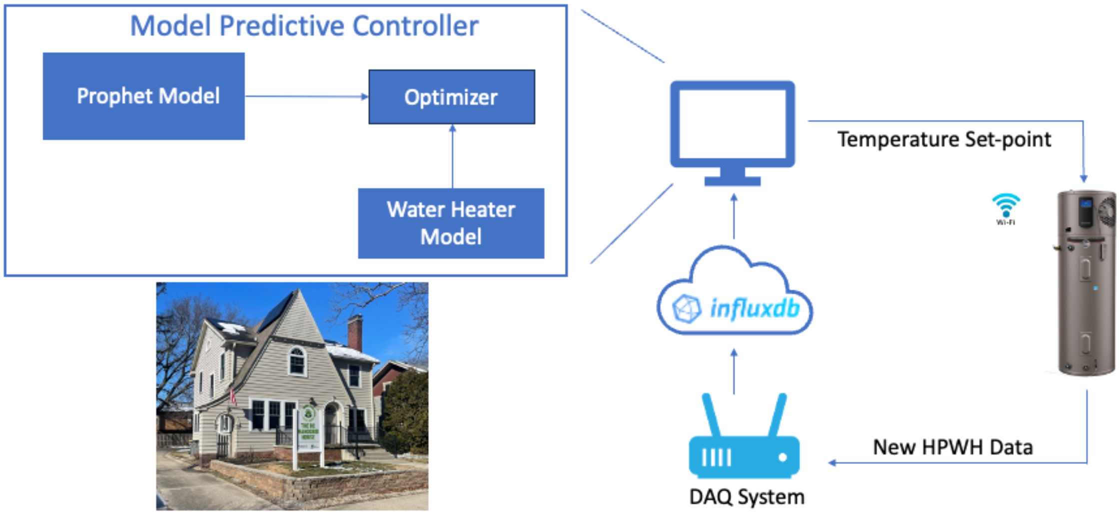

Figure 1: Presented is a diagram of the IoT infrastructure in the DC Nanogrid House at Purdue University. There is a Yokogawa DAQ system reading the sensor data and this data is then fed to Influxdb for cloud storage. A computer in the home runs the MPC controller and the Python scripts to push the set-point adjustment to the water heater through the API connection. The Prophet model provides the hot water forecast to the optimizer that calculates the optimal set-point.

the cloud database, Influxdb ("InfluxDB 2.0", n.d.), using Python scripts (Pergantis, Priyadarshan, et al., 2024). The controller is then operated on a computer in the home that also pushes the set-point adjustments to the water heater Fig. 1. The architecture of the IoT involved for this work could be minimized to a single raspberry pi that acts as the DAQ system, holds the minimum necessary data, and runs the MPC controller to update the water heater set-point.

3. MODELING

Initially, a single-node model was developed for the water heater, where the tank was modeled as a weighted average between the upper and lower tank temperatures. The tank-temperature was modeled as 75% of the upper tank temperature and 25 % of the lower tank temperature used in Starke et al., 2020. For the majority of the operation, this model form resembled the tank temperature reasonably well; however, since the set-point temperature was set to be the optimized set-point temperatures, during hot water use periods, the optimization problem would think it would need to heat the tank with the lower tank temperature dropping from cold inlet water. In reality, the upper tank temperature would drop insignificantly or much slower than modeled due to stratification over the height of the tank.

To better resemble the stratification of the water inside the tank, a two-node model was developed. There are two ways to develop a two-node model. The first is to control the height of the stratification and assume the temperatures are held constant (Diao et al., 2012). Another way to model the two nodes is to hold the height of the stratification layer constant and control the temperatures (Jin et al., 2014). Since the control parameter for a water heater is the set-point temperature, a two-node model where the height of the stratification layer is constant is used. This is because of the difficulty of translating a change in height of a stratification layer to a set-point temperature to send to the water heater. There are then two energy balance equations to describe the two-node model with a stratification layer resembling a resistance layer (R_{sl}) in Equation (1 & 2).

$$zC\dot{T}_H = \frac{T_L - T_H}{R_{sl}} + \frac{z(\theta - T_H)}{R_\theta} + \dot{m}c_p(T_L - T_H) + \lambda * q \tag{1}$$

$$(1-z)C\dot{T}_L = \frac{T_H - T_L}{R_{sl}} + \frac{(1-z)(\theta - T_H)}{R_\theta} + \dot{m}c_p(T_{in} - T_L) + (1-\lambda) * q \tag{2}$$

The left side of the Equation (1) and Equation (2) represent the change in capacitance (C) of the tank due to a temperature change ($\dot{T}$) for their respective portion of the tank. The resistance value (R_θ) to the ambient temperature surrounding the tank (θ) are used to calculate the heat loss through the tank wall. The heat transfer due to the flow across the boundaries is a function of the mass flow rate of water ($\dot{m}$), specific heat capacity of the water (c_p) and the temperature difference between the inlet and outlet water. The heat input into the tank from the condenser is represented by q and is split between the two nodes. To minimize unnecessary variables, z and λ are used to represent the fraction of the height of the top node and the fraction of the heat being injected to the top part of the tank from the condenser coil. These parameters are tuned through comparing the model to the measured temperatures on the upper and lower part of the tank. To most closely resemble the real water heaters operation, 0.6 and 0.66 for z and λ are chosen for the parameters respectively. Converting Equation (1 & 2) to a continuous-time state space equation results in:

$$\begin{bmatrix} \dot{T}_H \\ \dot{T}_L \end{bmatrix} = \tilde{A} * \begin{bmatrix} T_H \\ T_L \end{bmatrix} + \tilde{B} * q + \tilde{w} \tag{3}$$

Where $\tilde{A}$, $\tilde{B}$, and $\tilde{w}$ represent the thermodynamics dependent on the state variables T_H and T_L, the input variable q and the disturbance variables respectively.

4. FORECASTING

Based on (Shen et al., 2021), a study that builds a Model Predictive Controller for hot water scheduling of a multi-family housing complex, we apply a Prophet forecast model for the prediction of the upcoming water draws over the next day in 5-minute time steps, yielding 288 forecasted values. The Prophet model is an open-source tool that treats time-series forecasting problems as curve-fitting problems, with a consideration of seasonality and trends (Taylor & Letham, 2018).

In Table 1, we compare the forecasting quality of the Prophet algorithm against a naive recency-based forecast, that simply predicts the day-before values for the upcoming 288 time steps, based on the Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE), which are frequently used metrics to evaluate time-series forecasts (vom Scheidt et al., 2020). For the comparison, we conduct forecasts for the period from 2024-03-10 and 2024-03-17, with model retraining every 5 minutes. We use training data going back to 2024-02-14. For both metrics, the Prophet model outperforms the Naive model. However, we note that the performance difference is insignificant for the Mean Absolute Error.

Table 1: The forecast evaluation results. The Naive forecast is close in the MAE performance metric; however, the prophet model far outperforms in the RMSE performance metric.

	Naive	Prophet
RMSE	0.206	0.141
MAE	0.059	0.058

In Figure 2, we depict an exemplary forecast of the Prophet and Naive forecast during the MPC validation period, comparing the forecast results with the actual values. We can observe that the Prophet forecast predicts probable shower events in the morning with a rather smooth curve, which can be effectively used to schedule pre-heating. We note that the Prophet algorithm might also predict negative values due to its setup, which we normalize to zero before applying the prediction for the MPC. The negative values in the Prophet model are also filtered to be zero when being used in the MPC controller.

5. CONTROLLER DESIGN

The MPC controller is a linear convex optimization problem using the default solver for CVX (Grant and Boyd, 2014). The objective of the optimization problem is to minimize the cost of operating the water heater while trying to maintain the temperature at 46.1 °C. This temperature preference (T_{pref}) is to maintain thermal charge if there are any unexpected water draws. A one-norm formulation was used (Borrelli et al., 2016) to minimize the total cost of operation over the prediction horizon and the cost was set to be the electricity price in West Lafayette, \$0.15 per kWh. To try to maintain

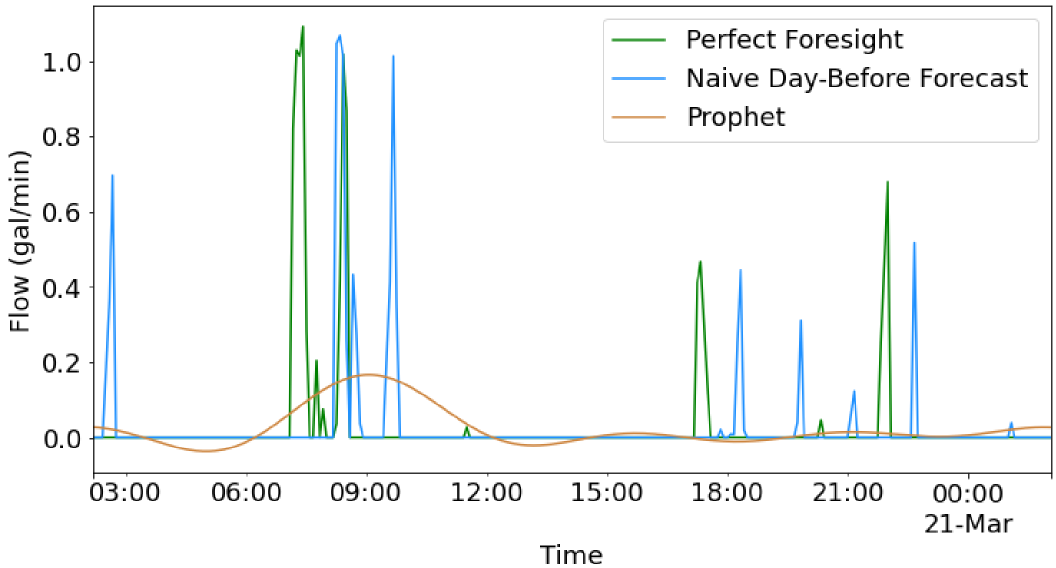

Figure 2: Exemplary Prophet and Naive forecast. The Prophet model has a more wave-like form capturing the overall trend of the historical hot water usages. The Naive forecast has much higher peaks and is a better replication of real hot water usage; however, does not consider the historical hot water usage trends in the home.

46.1 °C outlet water, a one-norm was also used with the price being the same as the cost of electricity. This is considered a soft constraint that provides a preference in the optimization to keep the upper tank temperature at this temperature seen in Equation (4e) (Rawlings et al., 2017). With the continuous energy balance model being discretized with 5-minute time-steps Equation (4b), the MPC optimization problem results in the following formulation:

Objectives: $$\min J = \Delta t * \left(\sum(\pi_{elec} * P) + \sum |T_{pref} - T_H|\right) \tag{4a}$$

State: $$\begin{bmatrix} \dot{T}_H \\ \dot{T}_L \end{bmatrix} = A * \begin{bmatrix} T_H \\ T_L \end{bmatrix} + B * q + w \tag{4b}$$

Initial states: $$T_H = T_{\text{initial,sensor,H}} \; ; T_L = T_{\text{initial,sensor,L}} \tag{4c}$$

Expressions: $$P = \frac{q}{COP} \tag{4d}$$

Constraints: $$0 \leq q \leq P_{\max} * COP \tag{4e}$$

The constraint is the limitation of the modeled heat allowed by the system. These limitations are related to the capacity of the heat pump, a function of the COP which is displayed in Equation (4d & 4e). The COP is assumed to remain constant to preserve linearity, and any deviations to the real COP will prompt adjustments to the set-points through the inherent feedback loop in the MPC operation. There are no constraints set on the water heater temperature as this would cause infeasibilities when the temperature drops far enough below the minimum set point where the heat pump can't heat to the minimum set point with its limited capacity. The set points are therefore post processed to be at least 43.3°C and to be whole numbers in Fahrenheit.

For the forecast of the hot water usage, the Prophet model is used because of its more accurate RMSE and distribution of likely water usage periods. While the Naive forecast more accurately captures the peaks of a water usage event, the MPC operation would be highly sensitive to the timing of these peaks, which are often inaccurate. The Prophet model offers the MPC controller more reliable estimates of likely water usage periods, allowing for better planning.

6. RESULTS

A four-day field test was conducted, during which the MPC controller operated the water heater under a constant electricity price. This operation was compared to a baseline test conducted the previous week, from Friday to Tuesday, using the default constant set-point control logic for the water heater. The set-point was set to 48.89°C (120°F) as this is the suggest water temperature set-point for most water heaters (HUD, 2023). The baseline controller and the MPC controller experienced very similar water usages as seen in Figure 3, with the total liters of hot water amounting to 700.3 and 681.4 respectively.

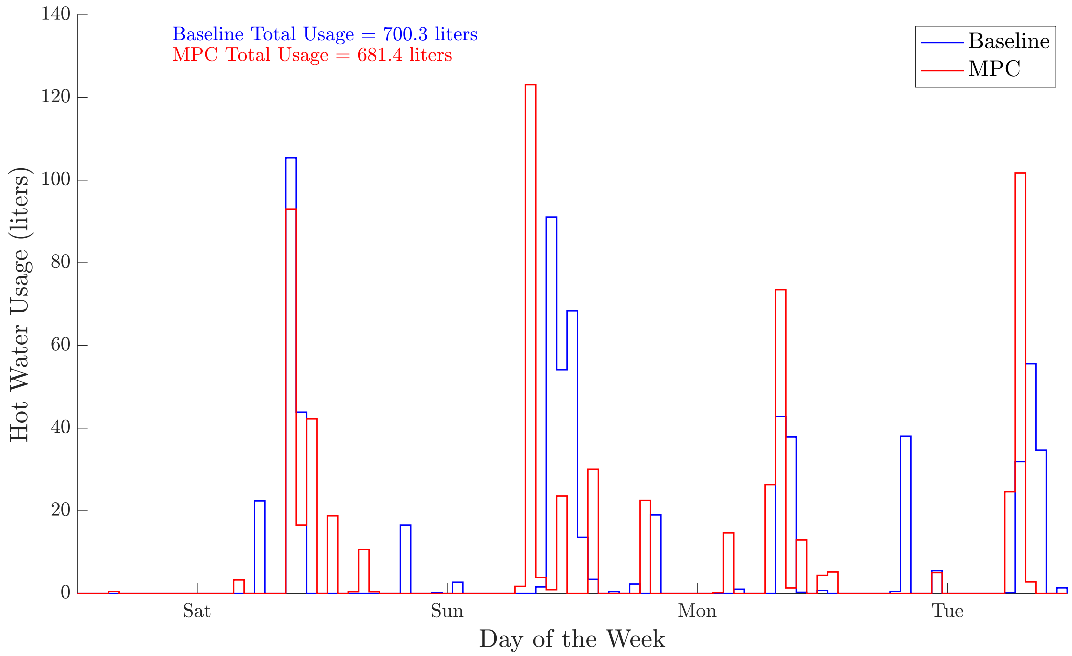

Figure 3: The blue represents the hot water usage during the MPC operation and the red under the baseline controller. The total hot water usages during the field test periods are 700.3 liters for the baseline operation and 681.4 liters for the MPC. The time-step in this figure is an hour to provide a less scattered visualization. The hot water usage occurs at similar times and at similar magnitudes, therefore making them reasonable comparison periods.

During shower periods confirmed by the occupants, the minimum temperatures felt were approximately 33.2°C and 31.1°C for the MPC controller and the constant set-point control respectively. Both of these temperatures are below the human body temperature, therefore indicating uncomfortable temperatures (Tadakatsu Ohnaka and Watanabe, 1994). These colder temperatures occurred twice out of the twelve showers taken for both controllers. This occurred when the three occupants showered consecutively with only a couple of minutes in between showers. These colder temperatures were felt by the occupant that took the third shower causing them to end the shower early. When only two showers were taken consecutively, or a single shower was taken, the temperatures stayed within a comfortable regime where there were no occupant complaints. The overall minimum temperatures were 29.7°C and 28.7°C for the MPC and baseline controller which are likely due to hot water draws after the shower, the delay of water temperatures mixing and the delay of the inlet flow where cold water still enters the tank after hot water usage has ended for a short period of time.

During the operation with the MPC controller, the set-point temperature closely stayed around the 46.1°C preference until dropping to the minimum set-point during high hot water use periods as seen in Figure 4. This is in part due to the objective function that was chosen where the upper tank temperature had the same penalization for increasing above 46.1°C as it did for dropping below. Additionally, because of the dead band in the default control for the temperature tracking the set-point in both operations, the MPC would sometimes increase the set-point thinking the heat pump would initiate when it wouldn't.

The MPC controller did raise the set-point temperature slower than the baseline after a water use period. This was due

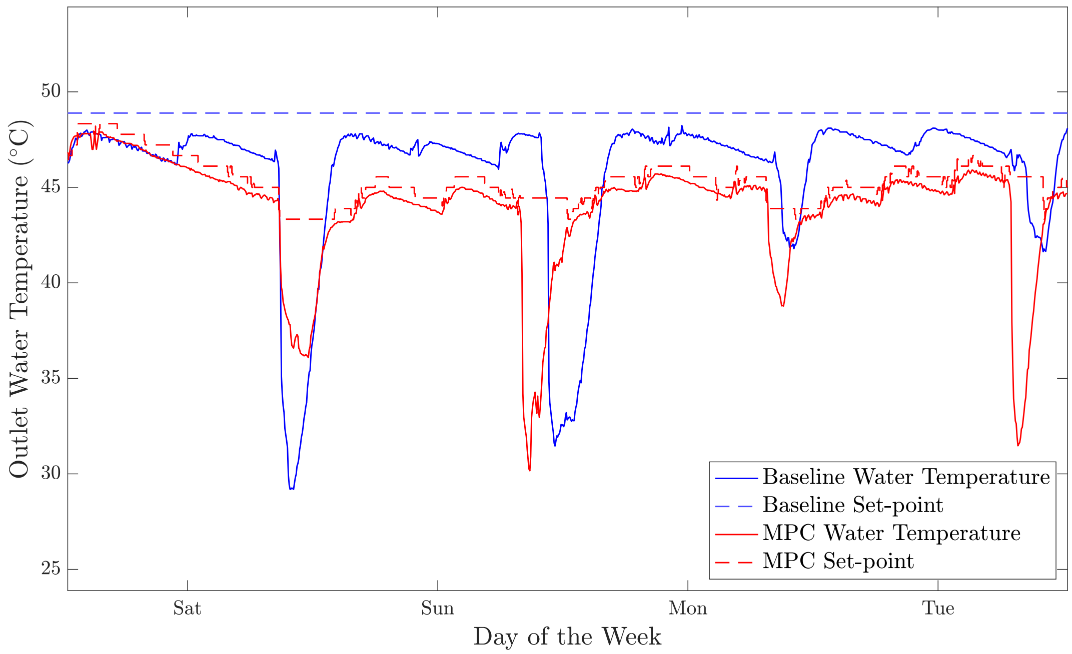

Figure 4: The blue represents the MPC operation and the red represents the baseline control where there is a constant set-point. The solid line is the outlet water temperature which is assumed to be the upper tank temperature. The dashed lines represent the set-point temperatures for the corresponding controllers.

to thermodynamic understanding of the MPC controller it found the most efficient solution was to heat the water back up at a slower rate to minimize the heat loss to the ambient conditions through the tank walls. Including this logic in a rule based controller would be difficult because the rate at which the tank temperature is heated back to its preferred temperature is dependent on the next expected water draw.

For the four days, the constant set-point controller and the MPC controller had the same thermal comfort; however, the MPC controller saved approximately 11 % in energy and cost, seen in Fig. 5. This was a result of the lower set-point temperatures in MPC operation and the minimization of heat losses through a slower ramp in temperature set-points. This correlates to roughly $20 in annual savings with a $0.15 per kWh price, as in West Lafayette, to upwards of $50-$60 in more expensive flat rate structures such as California for this size of a water heater.

An additional three days of testing was performed with the inclusion of a tracking model between the upper node temperature and the set-point in the MPC controller to enable better temperature tracking. The performance of this MPC formulation is summarized in Table 2 in comparison to the original MPC formulation and the baseline controller.

Table 2: A summary of the results including the testing performed with an updated MPC formulation. The first version of the MPC controller (MPC 1) had a lower operating cost than the adjusted MPC controller (MPC 2); however, MPC 2 had much better thermal comfort performance

	Baseline Control	MPC 1	MPC 2
Testing Duration (Days)	4	4	3
Discomfort Rate Out of Total Showers	2/12	2/12	0/8
Average Upper Tank Temperature (°C)	42.5	42.4	48.5
Shower Minimum Outlet Tank Temperature (°C)	31.1	33.2	43.2
Overall Minimum Tank Temperature (°C)	28.7	29.7	40.4
Energy Usage Per Liter ($\frac{kWh}{liter}$)	0.0453	0.0406	0.0554

With the adjusted MPC controller maintaining a higher set-point and better tracking, the thermal comfort was enhanced

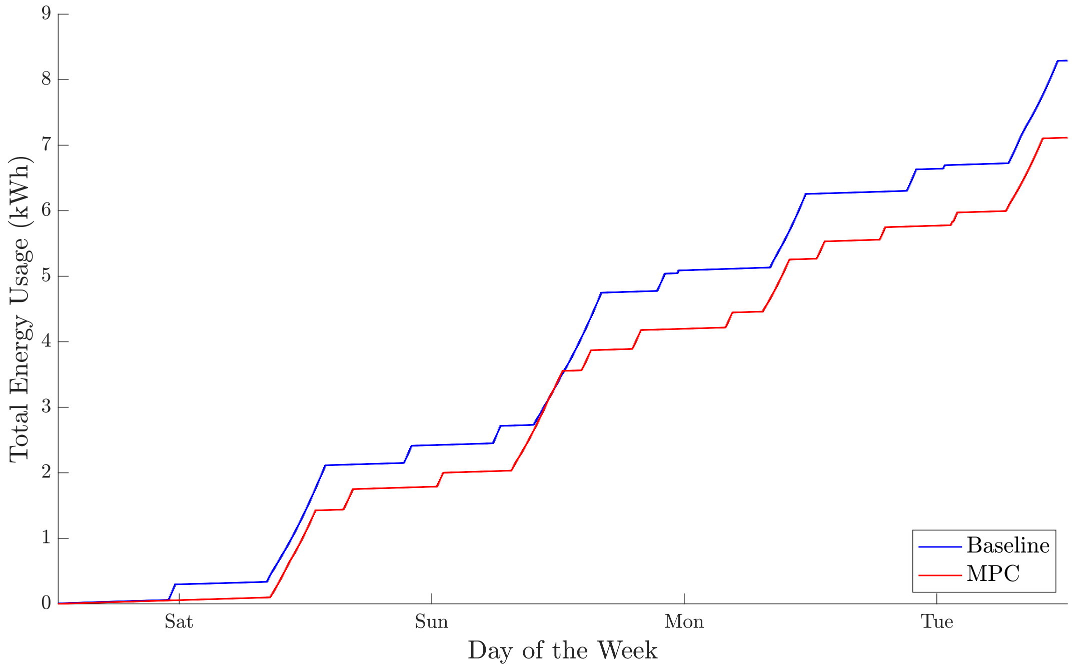

Figure 5: The blue represents the total energy usage under MPC operation and the red under baseline control operation. The MPC uses saves approximately 11% energy due to its lower set-points and its minimization of heat loss to ambient conditions.

significantly. Out of the eight showers taken, all of them were comfortable. This increase in thermal comfort did come at an increased cost of energy per liter though. No additional tuning of the objective function was performed either; therefore allowing for future improvement in energy savings.

7. DISCUSSION

An advanced controller for a heat pump water heater may not yield the highest savings in isolation; however, given the increasing adoption of dynamic rate structures, a controller capable of identifying optimal heating periods could prove advantageous. Further work will be completed to compare the comfort and energy performance against other control modes in the HPWH to provide a holistic evaluation. Moreover, the challenge of determining the optimal operation times for various devices within a home is a complex one, a challenge that smart home systems are currently tackling. This controller facilitates the integration of water heaters, including more sophisticated heat pump based systems, into smart home systems without resorting to blind control. Integrating water flow measurements and inlet water temperature readings into the tank is essential for this purpose. This can be achieved through additional sensors or by employing accurate state estimators that leverage the inherent temperature sensors on the upper and lower parts of many water heaters. In this study, we minimized sensing equipment and utilized convex linear programming for the MPC controller, enabling the use of cost-effective micro-controllers. This approach contrasts with other algorithmic solutions, such as mixed-integer programming, which require higher computational power (Linderoth & Savelsbergh, 1999).

8. CONCLUSION

This paper investigated the enhancement of a Heat Pump Water Heater (HPWH) by utilizing only the heat generated from the heat pump without electrical resistance elements. The findings elucidate the limitations of representative current market-available HPWH when operated in heat pump mode. The MPC controller, while maintaining the same comfort levels as the default heat-pump control mode, demonstrated an 11% energy and cost savings under a flat rate structure. Furthermore, through the incorporation of a tracking model and accurate water forecasting, the MPC controller exhibited superior thermal comfort maintenance at a higher cost than the baseline heat pump mode.

Future research directions could involve fine tuning the MPC controller's objective function to optimize thermal comfort at minimal costs and investigate how different rate structures impact water heater operational costs and the performance disparities between various water heater control modes. Additionally, a comparative analysis of water forecasting models, such as a random forests model, with a realistic training methodology in a single-family home and at smaller time steps, could provide valuable insights. Moreover, developing state estimators to replace the water flow meter and thermocouple for measuring inlet water temperature could reduce additional sensing equipment costs. The experimental setup at the DC Nanogrid House with the MPC controller on the water heater, developed in this work enables the future investigations mentioned, offering valuable insights and guiding further research directions.

NOMENCLATURE

T	water temperature	(°C)
$\dot{T}$	change in water temperature	(°C)
R	resistance value	$(\frac{°C}{kW})$
C	thermal capacitance of the tank	$(\frac{kW}{°C})$
$\dot{m}$	water mass flow rate	$(\frac{kg}{s})$
c_p	specific heat capacity of water	$(\frac{kJ}{kg°C})$
q	heat transfer supplied from the condenser	(kW)
P	power from water heater	(kW)
θ	ambient air temperature	(°C)
λ	top node height as a fraction of tank height	
z	top node heat injection relative to total heat supplied	
$\tilde{A}$	continuous time state matrix	
$\tilde{B}$	continuous time input matrix	
$\tilde{w}$	continuous time disturbance matrix	
A	discrete time state matrix	
B	discrete time input matrix	
w	discrete time disturbance matrix	
COP	coefficient of performance	
UEF	uniform energy factor	

Subscript

H	upper tank water
L	lower tank water
in	inlet water
pref	preference
sl	stratification layer
θ	ambient air

REFERENCES

Baumann, C., Huber, G., Alavanja, J., Preißinger, M., & Kepplinger, P. (2023). Experimental validation of a state-of-the-art model predictive control approach for demand side management with a hot water heat pump. *Energy and Buildings*, *285*, 112923.

Borrelli, F., Bemporad, A., & Morari, M. (2016). *Predictive control for linear and hybrid systems* (1st). Cambridge University Press.

dela Rosa, L., Mande, C., & Ellis, M. J. (2023). Supervisory multi-objective economic model predictive control for heat pump water heaters for cost and carbon optimization.

Diao, R., Lu, S., Elizondo, M., Mayhorn, E., Zhang, Y., & Samaan, N. (2012). Electric water heater modeling and control strategies for demand response. *Power and Energy Society General Meeting, 2012 IEEE*, 1–8.

Drgoňa, J., Arroyo, J., Cupeiro Figueroa, I., Blum, D., Arendt, K., Kim, D., Ollé, E. P., Oravec, J., Wetter, M., Vrabie, D. L., & Helsen, L. (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control*, *50*, 190–232.

EIA. (2020). Residential energy consumption survey (recs).

EIA. (2023). Total energy-annual energy review.

Grant, M., & Boyd, S. (2014, March). CVX: Matlab software for disciplined convex programming, version 2.1.

Heidari, A., Olsen, N., Mermod, P., Alahi, A., & Khovalyg, D. (2020). Adaptive hot water production based on supervised learning. *Sustainable Cities and Society*, *66*.

HUD. (2023). Operations and maintenance: Domestic hot water.

IEA. (2022). The future of heat pumps.

IEA. (2023). Net zero roadmap: A global pathway to keep the 1.5 °c goal in reach.

InfluxDB 2.0 [Accessed: 03/02/2024]. (n.d.).

Jin, X., Maguire, J., & Christensen, D. (2014). Model predictive control of heat pump water heaters for energy efficiency.

Khabbazi, A. J., Pergantis, E. N., Premer, L. D. R., Lee, A. H., Jie Ma, H. L., Henze, G. P., & Kircher, K. J. (2024). What have we learned from field demonstrations of advanced commercial hvac control? *8th International High Performance Buildings Conference at Purdue*.

Linderoth, J. T., & Savelsbergh, M. W. P. (1999). A computational study of search strategies for mixed integer programming. *INFORMS Journal on Computing*, *11*(2), 173–187.

Mande, C., Aboud, A., dela Rosa, L., et al. (2022). Timing is everything: Optimizing load flexibility of heat pump water heaters for cost, comfort, and carbon emissions. *UC Davis*.

Pergantis, E. N., Premer, L. D. R., Lee, A. H., Priyadarshan, Liu, H., Groll, E., Ziviani, D., & Kircher, K. J. (2024). Protecting residential electrical infrastructure through advanced control: The first field results. *8th International High Performance Buildings Conference at Purdue*.

Pergantis, E. N., Priyadarshan, Theeb, N. A., Dhillon, P., Ore, J. P., Ziviani, D., Groll, E. A., & Kircher, K. J. (2024). Field demonstration of predictive heating control for an all-electric house in a cold climate. *Applied Energy*, *360*, 122820.

Priyadarshan, Pergantis, E. N., Crozier, C., Baker, K., & Kircher, K. J. (2024). Edgie: A simulation test-bed for investigating the impacts of building and vehicle electrification on distribution grids.

Rawlings, J., Mayne, D., & Diehl, M. (2017). *Model predictive control: Theory, computation, and design*. Nob Hill Publishing.

Shen, G., Lee, Z. E., Amadeh, A., & Zhang, K. M. (2021). A data-driven electric water heater scheduling and control system. *Energy and Buildings*, *242*, 110924.

Sparn, B., Hudon, K., & Christensen, D. (2014). Laboratory performance evaluation of residential integrated heat pump water heaters.

Star, E. (2023). Water heater key product criteria.

Starke, M., Munk, J., Zandi, H., Kuruganti, T., Buckberry, H., Hall, J., & Leverette, J. (2020). Real-time mpc for residential building water heater systems to support the electric grid, 1–5.

Tadakatsu Ohnaka, Y. T., & Watanabe, Y. (1994). The effects of variation in body temperature on the preferred water temperature and flow rate during showering. *Ergonomics*, *37*, 541–546.

Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, *72*(1), 37–45.

vom Scheidt, F., Medinová, H., Ludwig, N., Richter, B., Staudt, P., & Weinhardt, C. (2020). Data analytics in the electricity sector–a quantitative and qualitative literature review. *Energy and AI*, *1*, 100009.

ACKNOWLEDGMENT

The Center of High Performance Buildings (CHPB) at Purdue University and the National Science Foundation (NSF) supported this work. Additionally, this work was made possible by donations from Rheem Manufacturing and the occupants in the DC Nanogrid House.