Optimizing Cost and Carbon Footprint: Laboratory Testing of Model Predictive Control for Smart Management of Heat Pump Water Heaters

Caton Mande, UC Davis Western Cooling Efficiency Center
Fred Meyers, UC Davis Western Cooling Efficiency Center
Loren dela Rosa, UC Davis Chemical Engineering
Matt Ellis, UC Davis Chemical Engineering

ABSTRACT

Heat pump water heaters (HPWHs) are an efficient way to heat domestic hot water, but their performance can be further improved with advanced control strategies. This paper presents a laboratory study comparing two control strategies on a 65-gallon HPWH: setpoint-tracking rule-based control (RBC) and economic model predictive control (MPC). The RBC strategy is a simple, robust, and reactive approach that does not consider energy prices or greenhouse gas (GHG) emissions when it turns the heat pump on or off. In contrast to RBC, MPC is a more sophisticated and proactive control strategy that uses a mathematical model of the HPWH, the electricity tariff, and forecasts for exogenous inputs (like weather and marginal GHG emissions) to optimize the operation of the HPWH to minimize energy costs and GHG emissions without compromising user comfort.

This study evaluates a cloud based MPC against manufacturer's RBC based on cost, CO2 emissions, and peak runtime. The MPC controlled the HPWH by sending new setpoints through an application programming interface. Testing was conducted inside an environmental chamber to vary the ambient air temperature, in addition to the hot water demand profiles. Results over a 24-hour period demonstrate MPC's advantages: potential cost reduction of up to 29 percent, CO2 emission reduction of up to 61 percent, and peak runtime reduction of up to 79 percent compared to RBC. These findings highlight MPC's potential to significantly improve energy efficiency, reduce emissions, and enhance user comfort in HPWH operations.

Introduction

Water heating constitutes a significant portion of U.S. residential energy consumption, accounting for 18% of energy usage (EIA 2023) and 15% of residential greenhouse gas (GHG) emissions (Leung 2018). As renewable energy sources (RES) gain traction in the electric grid, transitioning to highly efficient electric water heaters is key for decarbonizing homes and cutting utility costs for end-users (Bailey 2020). Heat pump water heaters (HPWHs) present a promising solution, with energy efficiency two to three times greater than that of electric resistance water heaters (DOE 2018). However, widespread adoption of HPWHs raises concerns about increased demand on the grid (Brooks 2021). Additionally, effective decarbonization through electrification requires coordinating HPWH loads with the intermittent availability of RES. This can be achieved by implementing time-of-use (TOU) rates and marginal grid GHG emissions rate signals. TOU rates encourage customers to shift energy use during off-peak periods when electricity costs are lower and cleaner energy sources are more prevalent (Faruqui et al. 2019). Marginal grid GHG emissions rates indicate the environmental impact of additional grid demand, motivating customers to adjust energy use to periods of lower emissions (WattTime 2022).

Currently, most HPWHs are controlled by rule-based strategies that track a temperature setpoint (Lissa et al. 2021), regardless of electricity cost or marginal grid GHG emissions (Pean et al. 2019). Economic model predictive control (MPC) stands out as an advanced control strategy leveraging dynamic system modeling to predict future system behavior and compute optimal control actions minimizing an economic cost function (e.g., Ellis et al. 2014). MPC's predictive capability makes it ideal for providing load flexibility for HPWHs. By using a dynamic thermal model of the HPWH tank and forecasts of exogenous inputs, like hot water draws, MPC can predict tank temperature changes and determine optimal preheating strategies. Additionally, MPC can incorporate TOU rates and GHG emissions rate signals into its cost function, ensuring the HPWH operating strategy minimizes both electricity cost and GHG emissions while maintaining user comfort.

This paper presents laboratory tests validating the effectiveness of MPC in optimizing HPWH operation to reduce electricity costs and GHG emissions. In this study, the MPC functioned as a supervisory controller, sending water temperature setpoint schedules to the original equipment manufacturers (OEM) rule-based control (RBC) via an application programming interface (API), as detailed by dela Rosa et al. (2023). The MPC used a blended cost function to co-optimize the reduction of electricity costs and GHG emissions. Additionally, to maintain model linearity, two versions of the HPWH tank model were used to investigate the MPC performance impacts when tank temperatures were combined using arithmetic and geometric means. The performance of the MPC was compared to the existing RBC without MPC to evaluate its advantages based on electricity cost, GHG emissions, peak electricity cost runtime, and peak GHG emissions runtime.

Technical Approach

Laboratory Test Apparatus

The laboratory test apparatus was built to measure the performance of a 65-gallon HPWH (59-gallon nominal) in a laboratory setting. The plumbing for the apparatus was designed to follow real-world installation requirements and accommodate needs for measurement and verification (M&V). Figure 1 shows how water flows through the apparatus. First, the cold water from the lab's domestic water system passes through a coarse filter (25 μm) before entering the expansion tank. After the expansion tank, the cold water is drawn into the HPWH tank and the mixing valve (if the HPWH tank outlet is hotter than the mixing valve setpoint). After leaving the HPWH tank, the hot water flows through five feet of straight piping length to allow for proper flow development before entering the flow meter. This piping length was insulated to minimize temperature loss of the additional piping section (<1 °C loss). After the flow meter, the water goes into the hot inlet of the mixing valve. The outlet of the mixing valve is connected to a flow controller that simulates hot water draws from the various household fixtures.

The hot water flow controller was designed to accurately simulate flow ranging between 0.01 GPM to 10.27 GPM and operates using a binary-inspired additive flow to combine multiple calibrated flows to achieve the hot water draw profile's target flow at a frequency of 1 Hz. By design, the flow controller splits the hot water from the mixing valve outlet into an array of twelve flow channels of varying diameter. Each channel has a solenoid valve and a manually adjustable needle valve. The needle valves were adjusted to calibrate the specific flows shown in Figure 3, using flow by mass. For each channel, flow was drawn until 10 pounds of water were collected, the process was repeated until five consecutive tests were within 3% of the target flow rate. Based on this design and calibration process, the flow controller can simulate flows up to 3 GPM with an error <1.5% and for higher flows the error increases to <5% because of small pressure changes in the manifold.

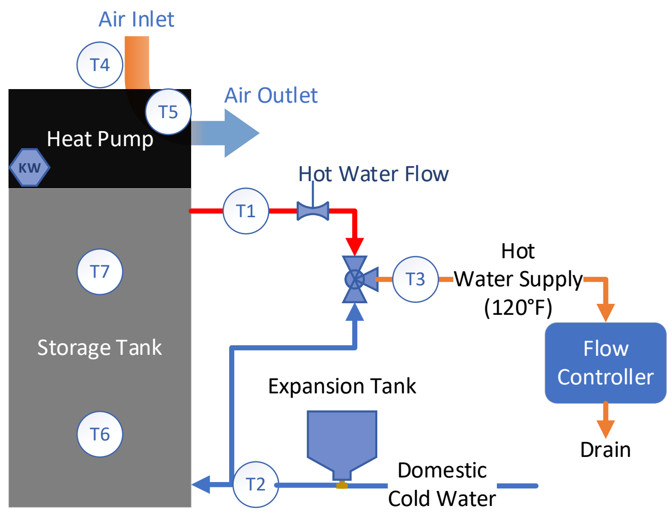

	Name	Source
T1	Hot Water Temperature	Lab M&V
T2	Cold Water Temperature	Lab M&V
T3	Mixing Valve Outlet Temperature	Lab M&V
T4	Air Temperature (4x)	Lab M&V
T5	Air Temperature	Manufacturer API
T6	Lower Tank Temperature	Manufacturer API
T7	Upper Tank Temperature	Manufacturer API
KW	HPWH Power/Energy Use	Lab M&V, Manufacturer API

Figure 1: Schematic of laboratory test apparatus and M&V layout

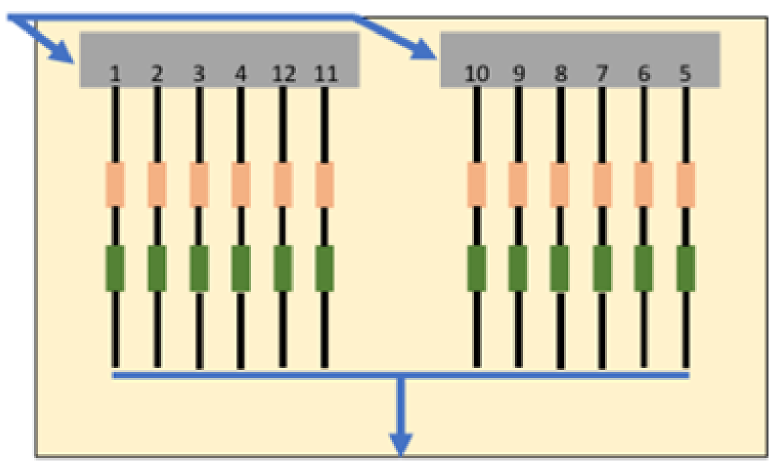

Channel	GPM	Channel	GPM
1	2	7	.32
2	2	8	.16
3	2	9	.08
4	2	10	.04
5	1	11	.02
6	.64	12	.01

Figure 2: Flow controller schematic and design flow rates for each channel

Table 1: Calibration results for each of the 12 channels of the flow controller

Valve	Desired Flow	Measured Flow	Percent Difference	Valve	Desired Flow	Measured Flow	Percent Difference
1	2	2.007	.35	7	0.32	0.322	0.769
2	2	1.98	1.15	8	0.16	.1624	1.51
3	2	1.99	0.35	9	0.08	0.0811	1.38
4	2	2.02	0.92	10	0.04	0.0403	0.64
5	1	0.990	0.988	11	0.02	0.0202	1
6	0.64	0.639	0.199	12	0.01	0.0985	1.5

HPWH Tank Initialization

An experiment initialization procedure was developed to ensure the water temperature in the tank would be consistent at the start of each test. This was required for repeatability and to allow results to be more comparable to simulation because tank temperature stratification happens as the HPWH sits unused. The initialization procedure was developed through trial and error, which identified that drawing approximately 9 gallons at a rate of 1 GPM, would trigger the HPWH to start heating the tank and return it to a repeatable temperature condition.

Testing and Data Acquisition

The laboratory test apparatus used NI LabVIEW to run experiments and capture sensor data. For each 24-hour experiment, the desired hot water flow profile for the flow controller would be inputted, the tank initialization procedure would begin, and the experiment would start once the HPWH shutoff at the end of the initialization procedure. Once the flow profile completed, the test would end unless the HPWH was actively heating the tank. In that scenario, the test would end once the HPWH shut off. All testing data was recorded at one of two sampling frequencies: 1Hz for the lab M&V sensors, and 0.2Hz for the data from the manufacturer's API. All test data was saved to a CSV for post processing. The exact sensor type and corresponding accuracy are shown in Table 2.

Table 2: Sensors used for lab testing HPWH. The tag corresponds to the locations shown in Figure 2

Tag	Name	Type/Source	Accuracy
T1	Hot water temperature	3-wire in-line RTD	±0.15 °C
T2	Cold water temperature	3-wire in-line RTD	±0.15 °C
T3	Mixing valve outlet temperature	3-wire in-line RTD	±0.15 °C
	Hot water flow	Turbine flow meter	±1.5% of reading
T4	Ambient temperature #1	3-wire air RTD	±0.15 °C
T4	Ambient temperature #2	3-wire air RTD	±0.15 °C
T4	Ambient temperature #3	3-wire air RTD	±0.15 °C
T4	Ambient temperature #4	3-wire air RTD	±0.15 °C
T5	Ambient temperature	Manufacturer API	-
T6	Lower tank temperature	Manufacturer API	-
T7	Upper tank temperature	Manufacturer API	-
KW	Power/Energy use	Dent PowerScout 3+	±0.2% of reading
	Current tank setpoint temperature	Manufacturer API	-

Cloud-based Supervisory MPC

This paper presents the next step in the control development described by the authors in dela Rosa et al. (2023) and presents results for the cloud based MPC using a blended cost function and two versions of the HPWH tank model.

One of the primary benefits of cloud based approach over a local deployment is its adaptability to retrofit existing cloud connected HPWHs without necessitating alterations to the existing hardware. However, a notable drawback of cloud-based deployment is the potential for communication disruptions or delays. To tackle this challenge, MPC operates as a supervisory controller, issuing temperature setpoints to the local RBC to motivate it to operate the HPWH as if there was a local MPC.

The process for running the MPC and communication with external data sources unfolds as follows. At intervals of 5 minutes, tank temperature data (upper, lower, and setpoint) from the lab HPWH is fetched from the local device, alongside forecasts of exogenous inputs sourced from APIs and forecasting models. This data serves as inputs for the MPC to calculate the optimal sequence of operation for the HPWH over the future 24-hour horizon.

Blended Cost Function

The MPC determines the set of control actions over the 24-hour time horizon that minimizes the blended cost function shown in Equation 1. The blended cost function consists of three terms: the electricity cost, the GHG emissions, and a comfort violation penalty term. Since the electricity rate and the GHG emissions rate are of different orders of magnitude, the emissions rate is standardized. A weighted sum of the electricity rate and the standardized GHG emissions rate is used to blend the two rates. The resulting sum is called the effective rate and is given by:

$$p_{eff,j+k} \coloneqq p_{elec,j+k} + P_{emission}\left(\frac{\tilde{p}_{ghg,j+k} - p_{ghg,ave}}{p_{ghg,stdev}}\right)$$

Equation 1: The effective rate

Where:

$p_{elec,k}$ is the electric tariff at the kth time step in the prediction horizon,
$P_{emission}$ is the weighing coefficient for the standardized GHG emissions term,
$\tilde{p}_{ghg,k}$ is the predicted marginal GHG emissions rate at the kth time step in the prediction horizon,
$p_{ghg,ave}$ is the average GHG emissions rate (computed offline based on historical GHG data),
$p_{ghg,stdev}$ is the standard deviation of the GHG emissions rate (computed offline based on historical GHG data),

At the jth time step of the 24-hour time horizon, the blended cost function is given by:

$$\sum_{k=0}^{N-1} p_{eff,j+k}\tilde{P}_k\Delta t + \omega_T\tilde{T}_{viol,k}$$

Equation 2: The MPC blended cost function

Where:

$\tilde{P}_k$ is the predicted power consumption of the HP at the kth time step in the prediction horizon,
Δt is the sample period,
ω_T is the weighing coefficient for the comfort violations term,
$\tilde{T}_{viol,k}$ is the predicted deviations of the average tank temperature from the comfort bounds at the kth time step in the prediction horizon

The MPC is designed to prioritize comfort over electricity cost and GHG emissions reduction. Thus, $\omega_T = 2$ to heavily penalize comfort violations. The value of $P_{emission}$ is chosen such that the MPC can achieve emissions reduction with minimal trade-off on cost and comfort. The effective rate, for a given residential TOU tariff, with a peak period from 4-9pm, and GHG emissions rate ($p_{ghg,ave} = 674$ and $p_{ghg,stdev} = 391$), is shown in Figure 3 and is used to provide insight into how to choose $P_{emission}$. When $P_{emission} = 0$, the effective rate is equal to the TOU rate. Thus, we would expect MPC to maximize HP runtime during the off-peak periods to minimize the blended cost function value. When $P_{emission} = 0.1$, the effective rate reaches its lowest point from 8 AM until 4 PM (the start of the peak period), representing a period abundant in low-cost, low-GHG electricity. Therefore, we would expect MPC to maximize HP runtime during this interval. Although the effective rate increases at 4 PM with the onset of the peak rate, the effective rate from 4 PM to 6 PM remains lower compared to 6-9pm. In response, the MPC may activate preheating during this period to ensure sufficient hot water in the tank, because the high-cost, low-GHG electricity has a lower effective rate than the period from 6-9pm.

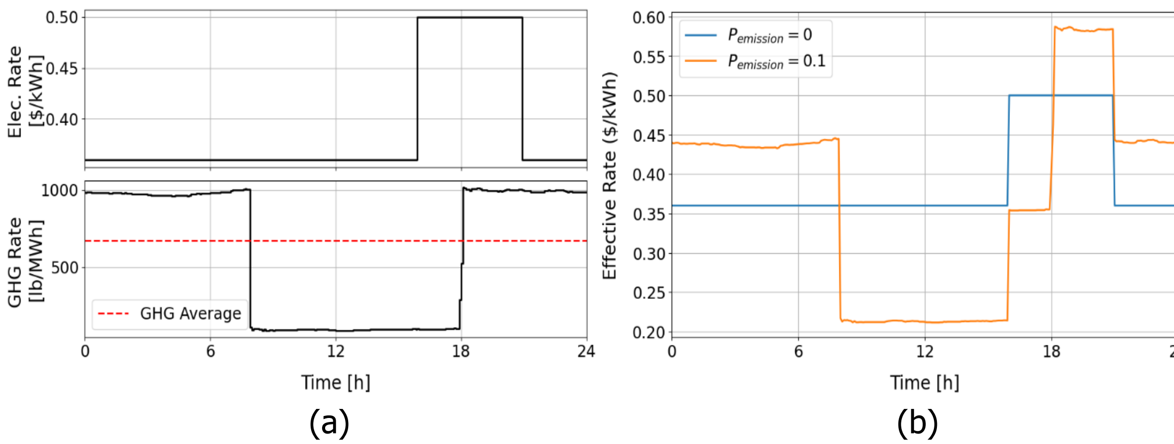

Figure 3: The (a) TOU tariff and GHG emissions rate, (b) effective rate.

Each HPWH laboratory test yields two sets of MPC results. First, $P_{emission} = 0$, where MPC optimizes HP operation based on cost and comfort, and second, $P_{emission} = 0.1$, where MPC optimizes HP operation based on cost, GHG emissions, and comfort. These MPC results are compared with those obtained from the RBC strategy for the same flow profile.

Averaged HPWH Tank Temperature

The 65-gallon HPWH used in laboratory testing has two water temperature sensors inside the storage tank (Figure 4, Left). This configuration leads to dividing the tank into at least two vertical layers in the dynamic thermal model. However, heat conduction between the layers introduces non-linearity into the model. To keep the MPC calculations simple, the two layers are combined into a single average temperature, keeping the model linear and reducing computational complexity. Figure 4 (Middle and Right), shows how the two temperatures were combined using the arithmetic and geometric means. For the geometric mean, the 80/20 split was determined by minimizing the error between open-loop simulations and actual data captured from the laboratory HPWH.

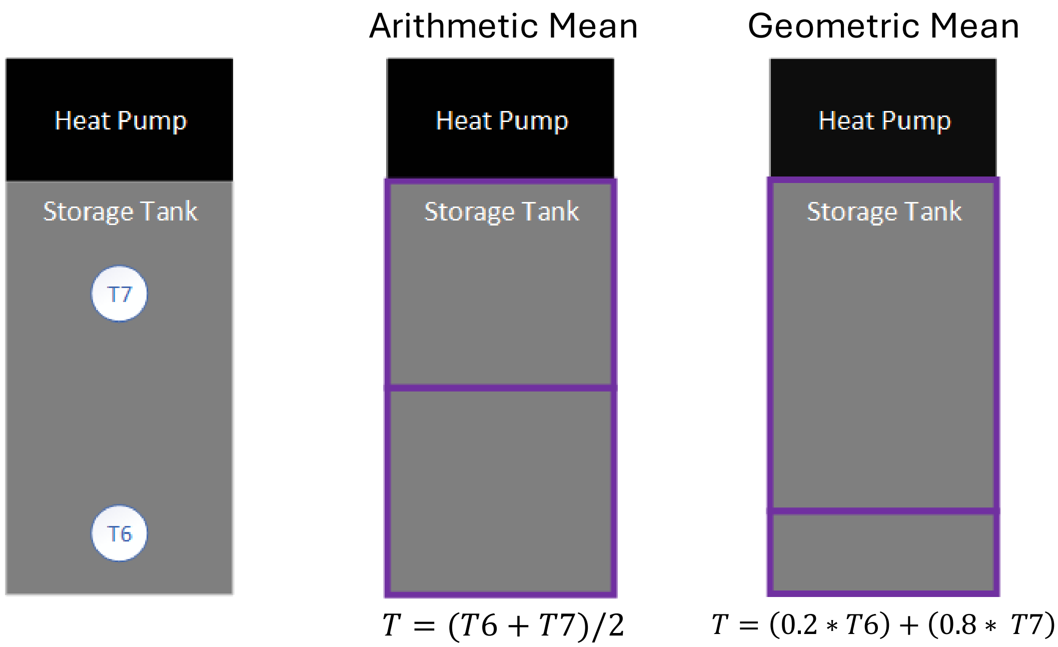

Figure 4: Illustration of how layers were combined into a single average temperature. Left, laboratory HPWH with two temperature sensors. Middle, arithmetic mean method where each layer is given equal weight in the average. Right, geometric mean method with an 80/20 split.

Hot Water Draw Profile

The hot water use profile has a major impact on HPWH performance since the hot water leaving the tank is replaced with water at the domestic supply temperature. The three water draw profiles used in the HPWH laboratory testing were taken from baseline field monitoring of a 65-gallon single-tank HPWHs in the two multifamily housing properties that will be used for field testing the MPC. Each HPWH has a thermostatic mixing valve to keep the water supplied to each fixture at or below 120°F. Figure 5 shows the 24-hour cumulative flow and the flow rate for each of the three profiles. This data represents the volume of hot water flowing out of the HPWH tank, so the actual flow at any fixture could have been higher than what is shown.

Profile #1

The first flow profile was selected based on the temporal and volume relationships between the draws and it represents a "best-case scenario" for RBC because no runtime occurs during the peak TOU period. Over the 24-hour period, the total volume drawn from the HPWH is 31-gallons, which is 53 percent of the nominal tank volume.

Profile #2

Flow profile #2 selected to define a "worst-case scenario" for RBC. For Profile #2, the selection criteria were: 1) identifying days where a field HPWH had significant runtime during the peak period, and 2) there was minimal or no water draws between 3-4pm, to allow MPC time to charge the storage tank. These criteria were not expected to produce the "best-case scenario" for MPC because high RBC runtimes during peak typically correlate with high water usage in that same time window. While MPC benefits from proactive control of the HPWH, it is limited to the storage capacity of the HPWH tank. Over the test period, the total volume drawn from the HPWH is 121-gallons, which is 205 percent of the nominal tank volume.

Test Cases

This work presents and compares three test cases that were run on the lab HPWH for each of the two flow profiles. For each test, the HPWH was configured to operate in "heat pump only" mode and each group of tests used the same hot water draw profile, electricity price tariff, and marginal GHG rate. The electricity price for all cases was based on a residential time-of-use tariff with a peak from 4-9pm every day and the marginal GHG rate represented a sunny California day where renewables were meeting the local distribution demand between 8am – 6pm in California climate zone 12.

Case 1: Baseline (RBC)

In the baseline RBC case, the HWPH operated in "heat pump only" mode with the tank setpoint temperature set to 120°F. The supervisory MPC was deactivated, and the unit ran according to the OEM control logic to maintain the setpoint.

Case 2: Cost only

For the cost only case, the HPWH operated in "heat pump only" mode and $P_{emissions} = 0$ in the blended cost function (Figure 5, b). Thus, the setpoints sent to the HPWH by the MPC were expected to minimize the operating cost for the 24-hour test.

Case 3: Cost and GHG

For cost and GHG case, the HPWH operated in "heat pump only" mode and $P_{emissions} = 0.1$ in the blended cost function For this weighting factor, the effective rate used in the cost function has four price changes over the 24-hour period, which can be described by four labels: low electricity cost and high GHG, low electricity cost and low GHG, high electricity cost and low GHG, and high electricity cost and high GHG (Figure 5, b).

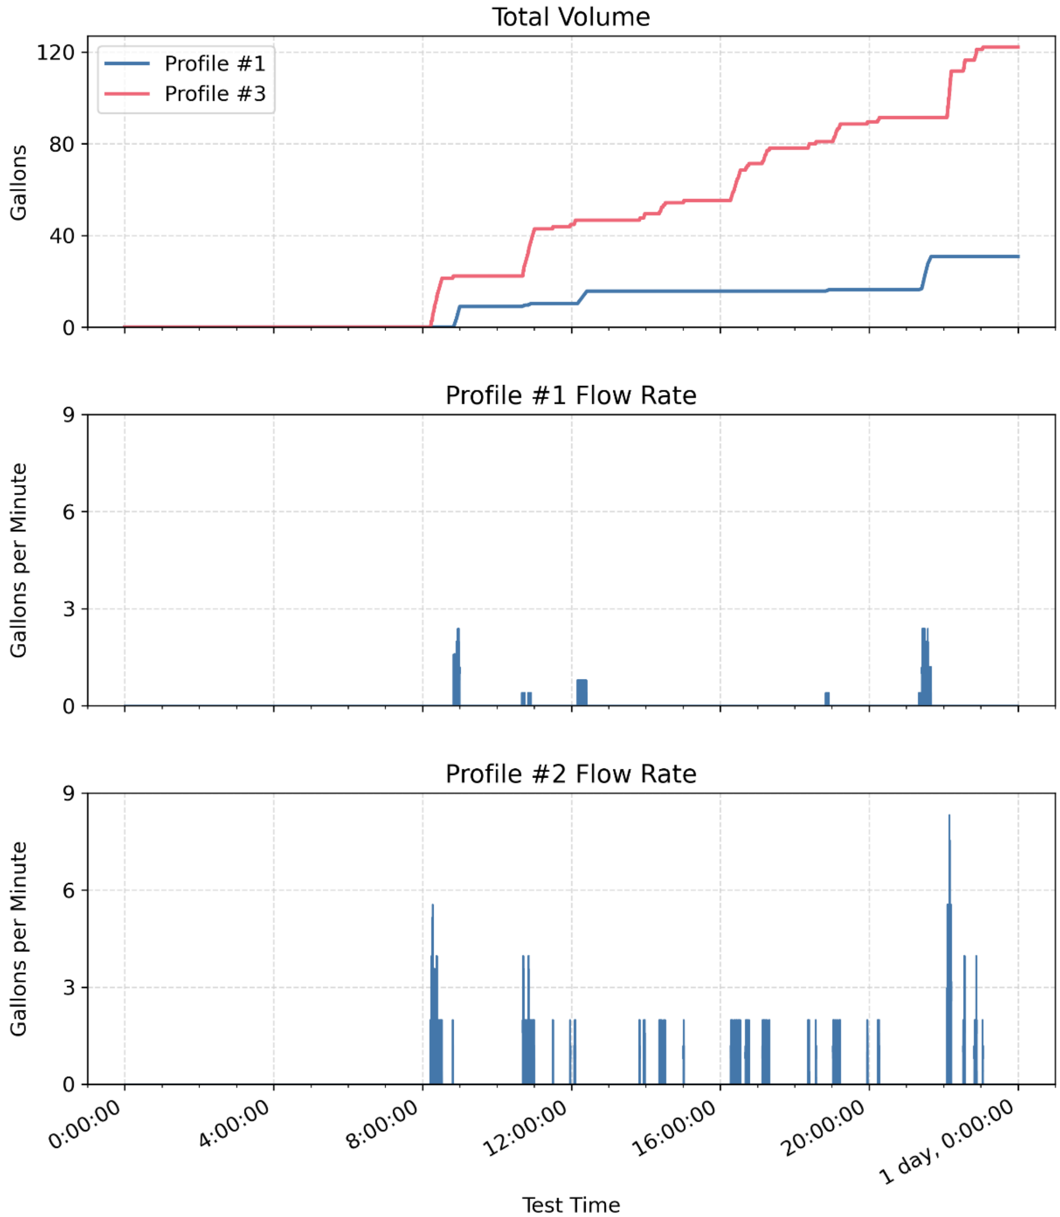

Figure 5: Hot water draw profile, totaling 33-gallons, selected from field data.

Results

Arithmetic Mean

This section presents the results for Profile #1 and #2 when MPC used the arithmetic mean to calculate the average tank temperature.

Profile #1

The results for the three test case scenarios (RBC, MPC cost only, MPC cost and GHG) are shown in Figure 6 The difference in HPWH runtimes demonstrates how the setpoints sent by the supervisory MPC in test case 2 and 3 changed the HPWH operation. Under RBC, the HPWH went through 3 compressor cycles, which primarily happened after large draws occurred. For the cost only case, the HPWH went through 10 compressor cycles, and it also did not run during the peak period. For the final case of cost and GHG, the compressor did 14 cycles during the test, two of which occurred during the peak period. This type of operation in the cost and GHG case was expected because the effective rate is at the second lowest level between 4-6pm, when electricity cost peaks but there is still a high availability of RES. Therefore, comparing the two MPC cases shows the cost and GHG case shifted part of the load from after 9pm to before the GHG emissions peak at 6pm.

Overall, the results show the supervisory MPC performed better than the baseline RBC in terms of cost and CO2 emissions, and the value for $P_{emissions}$ determined how much savings ended up in each bucket. The RBC had a total operating cost of $0.557 and generated 0.771 lbs of CO2. Cost only reduced the operating cost to $0.397 (-29 percent) and the CO2 emissions to 0.514 lbs (-33 percent). For cost and marginal GHG, the operating cost was $0.482 (-14 percent), and the CO2 emissions dropped to 0.301 lbs (-61 percent). The tabulated results for the three cases are shown in Table 3.

Profile #2

The results with Profile #2 for each of the three test cases are shown in Figure 7. The variation in HPWH runtimes explains how the setpoints provided by the supervisory MPC in test cases 2 and 3 altered the HPWH's operation. Under RBC, the HPWH went through 6 compressor cycles and had 167 minutes of runtime during peak electricity cost. For cost only, the HPWH went through 12 compressor cycles and reduce the peak runtime to 129 minutes. For cost and GHG, the compressor did 17 cycles during the test and reduced peak runtime to 75 minutes, which was split between three cycles.

Profile #2 included sufficient tank draws during peak to deplete the take storage potential, so peak runtime was expected for all test cases. For cost only, multiple compressor cycles occurred during peak which shows how the MPC with the blended cost function is evaluating cost, thermal losses, and user comfort. Additionally, in the cost only case, more non-unique solutions to the MPC problem are likely because the TOU rate only changes once during peak, so it can be difficult to find a global minimum, if one exists.

Overall, Profile #2 demonstrates the MPC can outperform the baseline RBC in terms of cost, CO2 emissions, or peak runtime. However, neither MPC case outperformed the RBC across all metrics. The RBC had a total operating cost of $1.266, generated 1.175 lbs of CO2, and had a peak cost runtime of 167 minutes. The cost only case reduced the operating cost to $1.223 (-3 percent), had the same amount of CO2 emissions, 1.614 lbs (0 percent), and reduced peak runtime to 130 minutes (-22 percent). For the cost and GHG case, the operating cost increased a little to $1.270 (+0.3 percent), the CO2 emissions dropped to 1.110 lbs (-6 percent), and the peak runtime was reduced to 75 minutes (-55 percent). The reduction in cost and CO2 savings compared to Profile #1 are driven by reduction in HPWH operating efficiency as the water is heated to a higher temperature. The tabulated results are shown in Table 4.

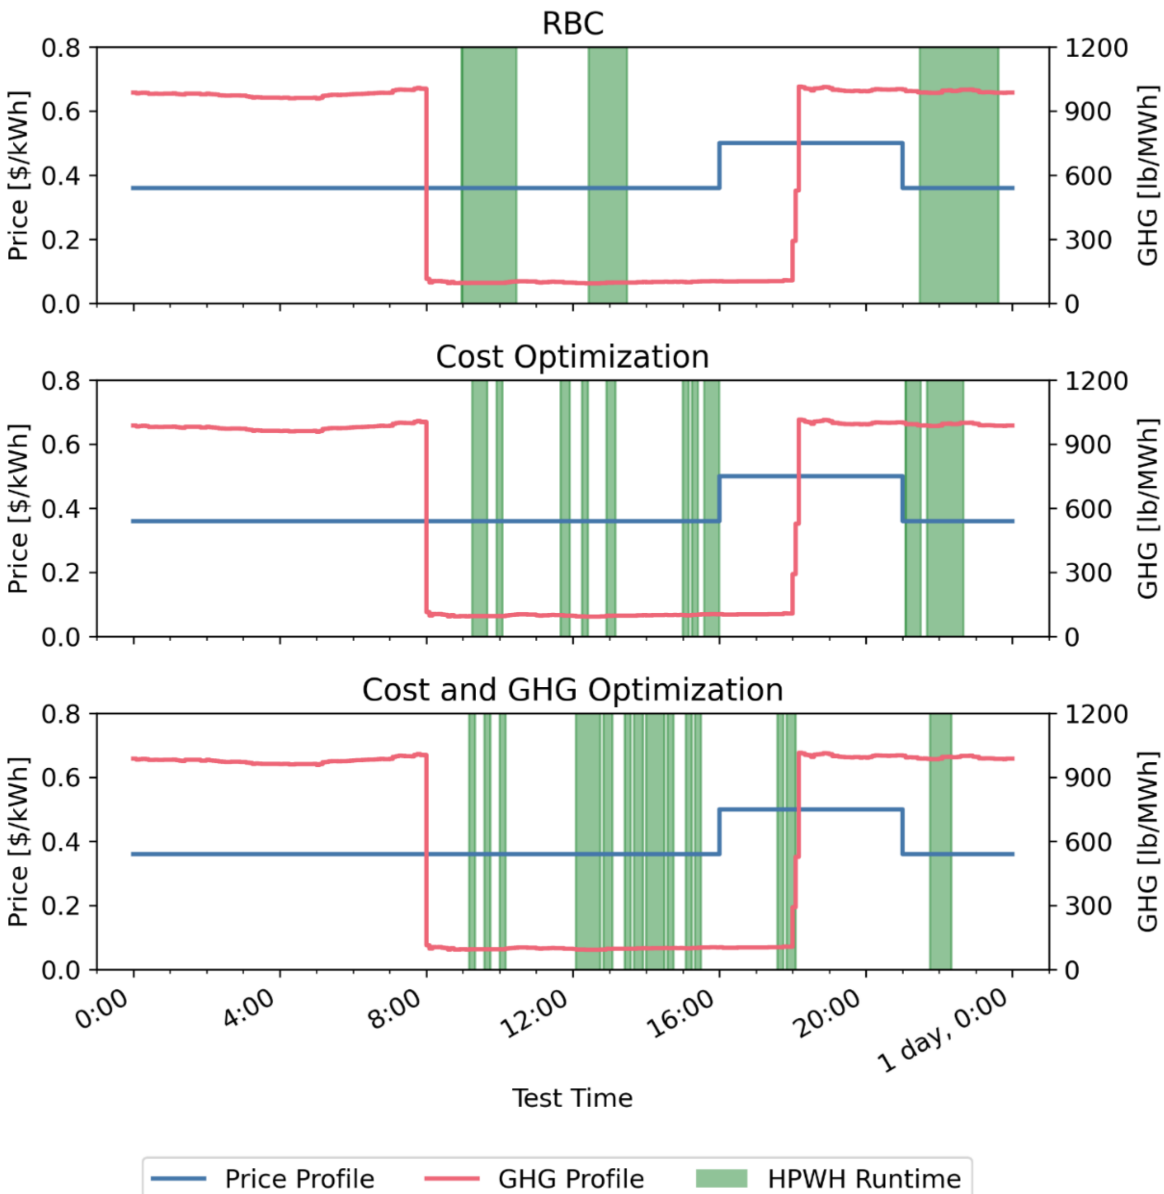

Figure 6: Results for Profile #1 using the arithmetic mean for average tank temperature. Each plot includes the electricity tariff and GHG profile for the test day. Top: baseline RBC result. Middle: MPC result with TOU-only cost function. Bottom: MPC result with marginal GHG emissions included.

Table 3: Tabulated results for Profile #1 using the arithmetic mean.

	RBC Baseline	MPC cost only	Percent change	MPC cost and GHG	Percent change
Cost [$]	0.557	0.397	-29	0.481	-14
CO_2 [lb]	0.771	0.514	-33	0.301	-61
Peak Price Runtime [min]	0	0	-	25	-
Peak GHG Runtime [min]	129	85	-34	35	-73

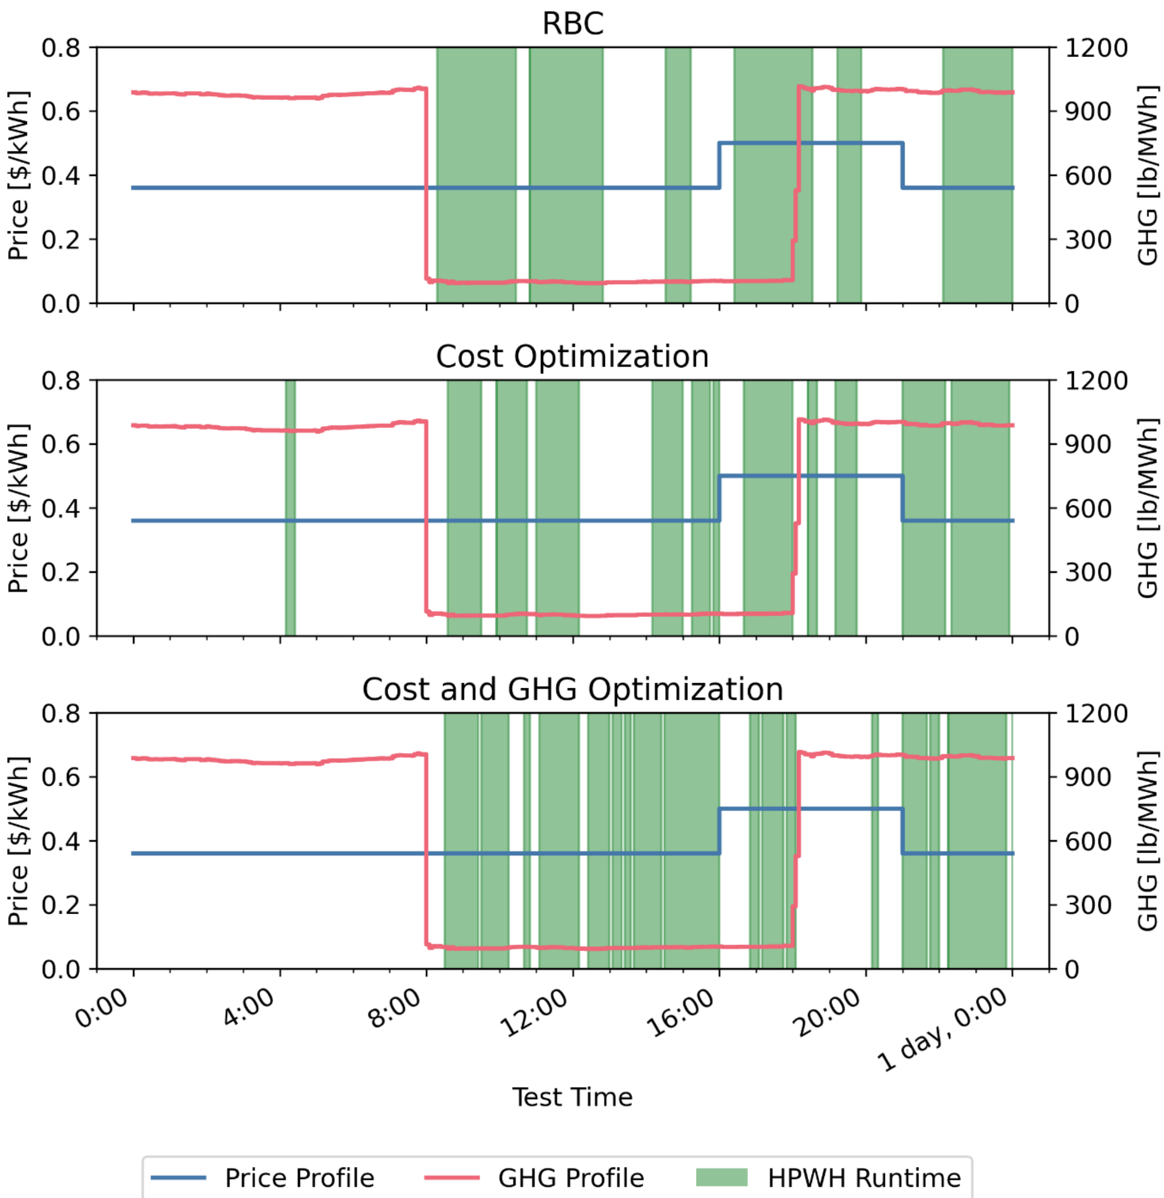

Figure 7: Results for Profile #2 using the arithmetic mean for average tank temperature. Each plot includes the electricity tariff and GHG profile for the test day. Top: baseline RBC result. Middle: MPC result with TOU-only cost function. Bottom: MPC result with marginal GHG emissions included.

Table 4: Tabulated results for Profile #2 using the arithmetic mean.

	RBC Baseline	MPC cost only	Percent change	MPC cost and GHG	Percent change
Cost [$]	1.266	1.223	-3	1.270	+0.3
CO2 [lb]	1.175	1.175	0	1.110	-6
Peak Price Runtime [min]	167	130	-22	75	-55
Peak GHG Runtime [min]	175	230	+31	159	-9

Geometric Mean

This section presents the results for Profile #2 when MPC used the 80/20 geometric mean to calculate the average tank temperature.

Profile #2

Figure 8 shows the results for Profile #2 when the MPC used an 80/20 geometric mean to calculate the average tank temperature. Overall, a reduction in HPWH runtime was observed. This result was anticipated, as the previous arithmetic mean results for Profile #2 showed that after 10-12 gallons were drawn from an idle tank, the MPC would significantly underestimate the amount of hot water remaining and begin preheating sooner than necessary to maintain comfort. Switching to the geometric mean allows the single average tank temperature to better represent the actual amount of hot water remaining in the tank.

The average tank temperature method does not impact RBC operation, so the baseline result of 6 compressor cycles and 167 minutes of runtime did not change for this set of tests. For cost only, the HPWH went through 14 compressor cycles and reduce the peak runtime to 115 minutes. For cost and GHG, the compressor did 10 cycles during the test and reduced peak runtime to 35 minutes.

Overall, the geometric mean produced an improvement in the MPC performance. The cost and GHG case outperformed RBC in all metrics. For comparison, the RBC had a total operating cost of $1.266, generated 1.175 lbs of CO_2, and had a peak cost runtime of 167 minutes. The cost only case reduced the operating cost to $1.023 (-19 percent), slightly increased the amount of CO_2 emissions to 1.222 lbs (+4 percent) and reduced peak runtime to 115 minutes (-31 percent). For the cost and GHG case, the operating cost reduced to $1.075 (-15 percent), the CO_2 emissions dropped to 0.999 lbs (-15 percent), and the peak runtime was reduced to 35 minutes (-79 percent). The tabulated results are shown in Table 5

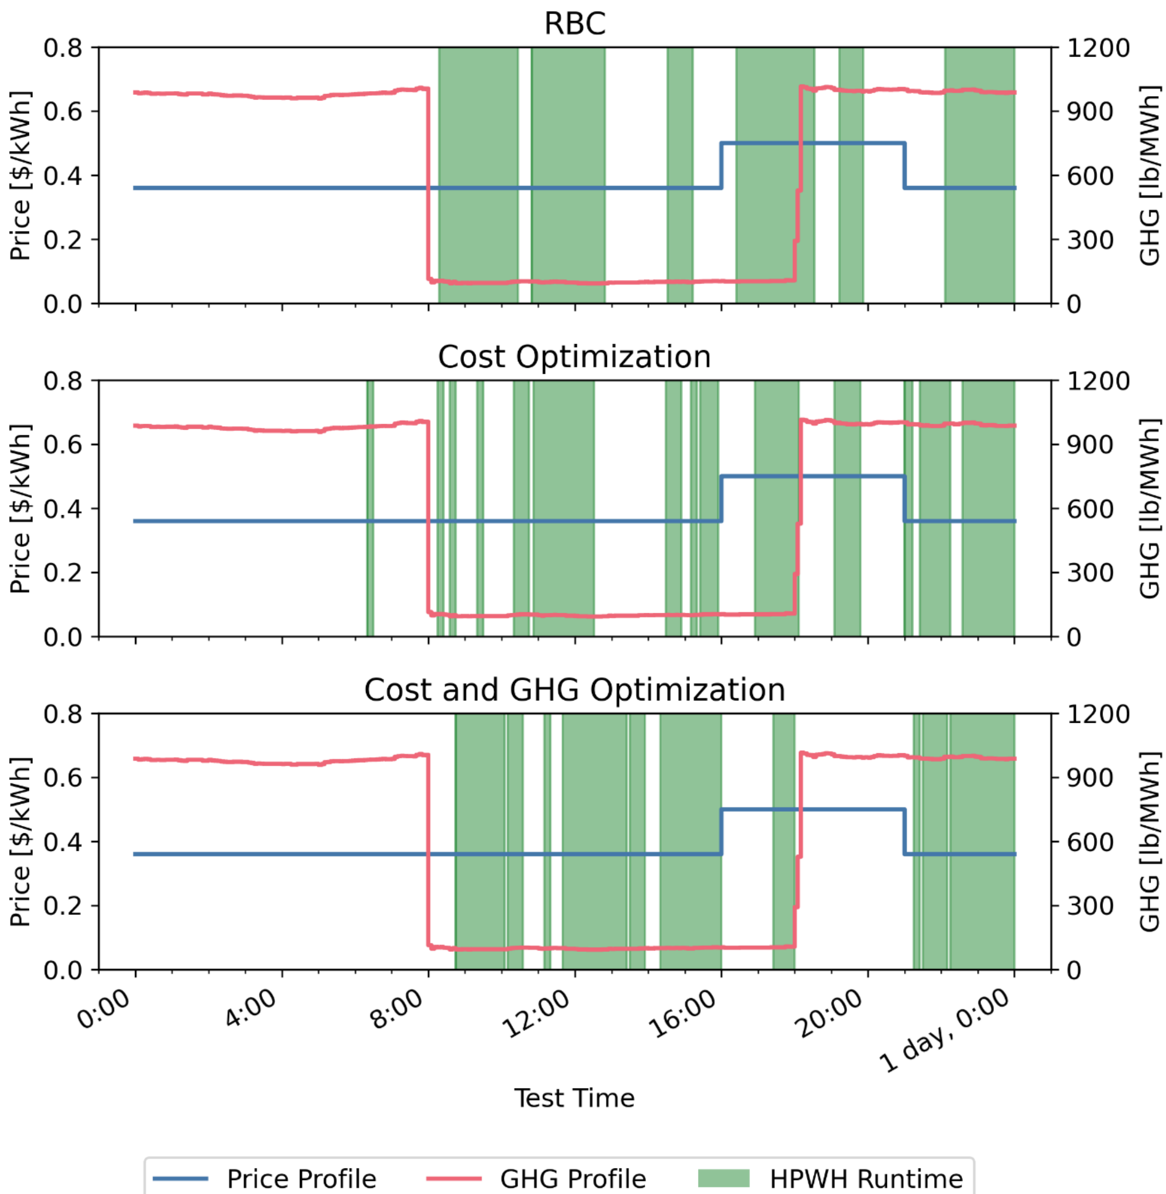

Figure 8: Shows the results for Profile #2, using the 80/20 geometric mean for average tank temperature. Each plot includes the electricity tariff and GHG profile for the test day. Top: baseline RBC result. Middle: MPC result with TOU-only cost function. Bottom: MPC result with marginal GHG emissions included.

Table 5: Tabulated results for Profile #2 using the 80/20 geometric mean.

	RBC Baseline	MPC cost only	Percent change	MPC cost and GHG	Percent change
Cost [$]	1.266	1.023	-19	1.075	-15
CO2 [lb]	1.175	1.222	+4	0.999	-15
Peak Price Runtime [min]	167	115	-31	35	-79
Peak GHG Runtime [min]	175	201	+15	155	-11

Conclusion

HPWHs are an efficient way to heat domestic hot water, but their overall performance can be further improved with advanced control strategies. This paper presents the results of a laboratory study that compares the performance of two control strategies on a HPWH: setpoint-tracking RBC and supervisory MPC. MPC is a more sophisticated and proactive control strategy that uses a mathematical model of the HPWH, the electricity tariff, and forecasts for exogenous inputs (such as weather and marginal GHG emissions) to optimize the operation of the HPWH to minimize energy costs and GHG emissions without compromising user comfort. In this study, the MPC was tested using two variations of the blended cost function. The first considered electricity cost only and the second balanced electricity cost with the standardized marginal GHG rate.

This study compared the performance of the MPC against the OEM RBC by conducting laboratory tests on a 65-gallon HPWH. Two different hot water draw profiles, derived from field data, were used. The MPC controlled the HPWH by sending new setpoints via an API. Additionally, the 65-gallon HPWH used in laboratory testing has two water temperature sensors inside the storage tank. To keep the MPC calculations simple, the two layers are combined into a single average temperature, keeping the model linear and reducing computational complexity. The laboratory testing results show the impact, of using the arithmetic and geometric means to represent the average tank, on MPC performance.

The result of this study shows that MPC can reduce HPWH operating cost and GHG emissions. In the cost only case, cost was reduced by 4 to 29 percent, peak runtime was reduced by 22 to 31 percent, and CO_2 emissions ranged from a 33 percent reduction to a 4 percent increase. For the cost and GHG case, cost ranged from a 15 percent reduction to a 0.3 percent increase, CO_2 emissions were reduced by 6 to 61 percent, and peak runtime was reduced by up to 79 percent. These results highlight the significant potential for the MPC strategy compared to OEM RBC in terms of energy cost, GHG emissions and user comfort.

Acknowledgements

Funding from the California Energy Commission under agreement #EPC-19-015 and EPC-20-025 is gratefully acknowledged.

References

Bailey, S. 2020. Adopted 2019 integrated energy policy report. https://www.whitehouse.gov/wp-content/uploads/2021/10/US-Long-Term-Strategy.pdf

Brooks, A. 2021. Thermal load shifting with heat pump water heaters: technology brief. Association for Energy Affordability.

DOE (Department of Energy). 2018. Heat pump water heaters. http://www.energy.gov/energysaver/water-heating/heat-pump-water-heaters.

dela Rosa, L., C. Mande, and M.J. Ellis. 2023. Supervisory multi-objective economic model predictive control for heat pump water heaters for cost and carbon optimization. ASHRAE Transactions, 129(1):517-525.

EIA (Environmental Impact Assessment). 2023. 2020 residential energy consumption survey. https://www.eia.gov/consumption/residential/data/2020/c&e/pdf/ce3.1.pdf.

Ellis, M., H. Durand, and P.D. Christofides. 2014. A tutorial review of economic model predictive control methods. Journal of Process Control 24:1156-1178.

Faruqui, A., R. Hledik, and S. Sergici. 2019. A survey of residential time-of-use (TOU) rates. https://www.brattle.com/wp-content/uploads/2021/05/17904_a_survey_of_residential_time-of-use_tou_rates.pdf

Leung, J. 2018. Decarbonizing U.S. buildings. Center for Climate and Energy Solutions. Tech. Rep.

Lissa, P., M. Schukat, M. Keane, and E. Barrett. 2021. Transfer learning applied to DRL-Based heat pump control to leverage microgrid energy efficiency. Smart Energy 3:100044.

Pean, T.Q., J. Salom, and R. Costa-Castello. 2019. Review of control strategies for improving the energy flexibility provided by heat pump systems in buildings. J. Process Control 74:35-49.

WattTime. 2022. Marginal emissions modeling: WattTime's approach to modeling and validation. https://www.watttime.org/app/uploads/2022/10/WattTime-MOER-modeling-20221004.pdf