Title
Supervisory Multi-Objective Economic Model Predictive Control for Heat Pump Water Heaters for Cost and Carbon Optimization

Authors
dela Rosa, Loren
Mande, Caton
Ellis, Matthew J

Publication Date
2023

Peer reviewed

AT-23-C060 Published in ASHRAE Transactions, Volume 129, Part 1

Supervisory Multi-Objective Economic Model Predictive Control for Heat Pump Water Heaters for Cost and Carbon Optimization

Loren dela Rosa **Caton Mande** **Matthew J. Ellis, PhD**

ABSTRACT

Heat pump water heaters (HPWHs) offer an efficient way to heat water using electricity, which aligns with efforts toward decarbonization and better utilization of renewable energy sources. Economic model predictive control (EMPC) provides an automated way to provide load flexibility of this new electric load by accounting for exogenous inputs like time-varying electric rates or marginal greenhouse gas emissions. Furthermore, a cloud-based supervisory controller implementation of EMPC enables retrofitting existing HPWHs with cloud-connection. In this work, the formulation of a supervisory multi-objective EMPC for HPWH is presented. The formulation of a supervisory multi-objective EMPC for HPWH is presented for equipment with a single heat pump and up to two backup resistive heating elements. A temperature setpoint is computed from the EMPC decisions using a logic-based setpoint calculator so the existing HPWH rule-based control (RBC) strategy activates the desired heat sources when deemed optimal by the EMPC. The performance of the simulated testing results under the supervisory EMPC is compared against the performance under an RBC strategy and under a regulatory EMPC that directly controls the HPWH. The simulation results demonstrate that the RBC in the proposed control architecture, operates the HPWH heat sources in the optimal manner computed by the EMPC, which indicates that the setpoint calculator can translate EMPC decisions to appropriate temperature setpoints. To minimize heat pump cycling, the effect of increasing the minimum on-time for the heat pump is also considered and the results show that increasing the minimum on-time increases the operating cost.

INTRODUCTION

In the United States, water heating is the second-largest residential energy end-use and accounts for 20% of total residential consumption (EIA 2018). Heat pump water heaters (HPWHs) offer an efficient way to heat water using electricity, aligning with efforts toward decarbonization and leveraging renewable energy sources. HPWHs have inherent thermal energy storage, which can shift electricity use away from peak hours by storing hot water generated with low-cost, emissions-free renewable electricity for use later in the day (Delforge 2020). Nonetheless, determining the optimal time and amount to preheat the HPWH tank water is a non-trivial decision that should account for the current and future energy cost, water draw, grid greenhouse gas (GHG) emissions rate, and thermal losses to the ambient. Excessive preheating, for example, could lead to higher energy consumption and cost. HPWHs are typically controlled by a rule-

Loren dela Rosa is a Ph.D. student in the Department of Chemical Engineering, University of California, Davis, Davis, CA. **Caton Mande** is a Senior R&D Engineer with the Western Cooling Efficiency Center and a Ph.D. student in the Department of Mechanical and Aerospace Engineering, University of California, Davis, Davis, CA. **Matthew J. Ellis** is an Assistant Professor in the Department of Chemical Engineering, University of California, Davis, Davis, CA.

based control (RBC) strategy that maintains the tank water temperature in a region around the setpoint. While this strategy is proven and robust for maintaining a user-defined setpoint, this type of control usually does not provide any flexibility when the HPWH operates — it will run until it reaches the setpoint, regardless of the electricity cost or grid GHG emissions rate.

Economic model predictive control (EMPC) is an ideal choice for controlling HPWHs because the optimization-based predictive control technique determines control actions by predicting system behavior over a horizon to minimize an economic cost function. Thus, EMPC can automatically shift the HPWH's electric load from peak to off-peak periods. EMPC may directly control the HPWH heating sources, i.e., the heat pump and backup electric resistive heating element(s). In a previous work (Mande et al. 2022), an EMPC was proposed that directly controls the HPWH and demonstrated savings compared to an RBC strategy. However, such an implementation would require an on-premise deployment. An alternative approach is to implement EMPC as a supervisory controller. In this implementation, the EMPC does not directly control the HPWH. Instead, it determines the temperature setpoint, which it communicates to a lower-layer tracking controller. The lower-layer tracking controller, which could be an existing RBC strategy, manipulates the HPWH heating sources to achieve the setpoint. This implementation offers an advantage as it allows for a cloud-based EMPC, enabling retrofitting existing cloud-connected HPWH with EMPC without requiring the installation of additional on-premise equipment.

In this study, a supervisory multi-objective EMPC is presented. The multi-objective EMPC balances operating cost with expected grid GHG emissions associated with operating the HPWH while ensuring that the tank temperature is maintained within a comfortable range. Considering HPWHs with a heat pump and up to two backup resistive heating elements, EMPC decides the optimal operation of these heating sources. A logic-based setpoint calculator is then used to convert the operating decisions computed by the EMPC to a temperature setpoint. The temperature setpoint is communicated to the existing HPWH RBC strategy and is computed so that the RBC implements the optimal operating decisions of the EMPC. The performance under the supervisory EMPC is compared against the performance under an RBC strategy and under a regulatory EMPC that directly controls the HPWH based on simulation results. To minimize heat pump (HP) cycling, the effect of increasing the minimum on-time for the HP is also considered.

WATER HEATER TANK MODEL

In this study, the performance of the different control strategies is evaluated on a simulated HPWH that contains three heat sources to provide heat to the water. These heat sources are the heat pump (HP) condenser coil, an upper electric resistive heating element (E1), and a lower electric resistive heating element (E2). The HP condenser coil does not span the entire height of the tank and is positioned near the bottom of the tank. The HPWH tank contains two temperature sensors located above each electric resistive heating element to measure the upper and lower water temperatures, respectively. The hot water that is drawn from the top of the tank is replenished by cold make-up water that enters the bottom of the tank, maintaining a constant tank water volume.

Accurately modeling the temperature dynamics of the water in the tank is helpful for evaluating the effectiveness of load flexibility strategies for HPWHs. Inspired by previous work on modeling stratified thermal storage tanks (Kleinbach, Beckman, and Klein 1993; Nash, Badithela, and Jain 2017; Powell and Edgar 2013), the simulated HPWH storage tank is discretized into n vertically stacked volume segments referred to as nodes, with the bottom node designated as node one. The water temperature within each node is spatially uniform. The energy balance over each node can be expressed by a first-principles model given in Equation 1:

$$C_p m_i \frac{dT_i}{dt} = \dot{Q}_{i+1,i} + \dot{Q}_{i,i-1} + \dot{Q}_{conv,i} + \dot{Q}_{amb,i} + \dot{Q}_{HP,i} + \dot{Q}_{Ej,i} \tag{1}$$

where C_p is the heat capacity of the water, m_i is the mass of water in the i-th node, T_i is the water temperature of the i-th node, $\dot{Q}_{i+1,i}$ is the conductive and diffusive heat transfer rate between node i and the node above (i.e., node $i+1$), $\dot{Q}_{i,i-1}$ is the conductive and diffusive heat transfer rate between node i and the node below (i.e., node $i-1$),

$\dot{Q}_{conv,i}$ is the convective heat transfer rate from the water flowing into and out of the node i, $\dot{Q}_{amb,i}$ is the rate of heat transfer between node i and the ambient, $\dot{Q}_{HP,i}$ is the sensible heat transfer rate provided to node i by the HP condenser coil, and $\dot{Q}_{Ej,i}$ is the sensible heat transfer rate provided to node i by the j-th resistive heating element, which could be the upper or lower electric resistance heating element denoted by E1 or E2, respectively. The simulated water temperatures may experience temperature inversion (warmer water exists below cooler water) whenever a heat source is on. A correction term for temperature inversion is provided in Equation 2 (Nash, Badithela, and Jain 2017) to account for natural convection, given by:

$$\dot{Q}_{j,i} = \frac{\bar{k}A}{\Delta z}\left(T_j - T_i\right) \tag{2}$$

where A is the cross-sectional area of the tank, Δz is the node height, T_j is the temperature of node j, T_i is the temperature of node i, and $\bar{k}$ is the modified thermal conductivity according to:

$$\bar{k} = \begin{cases} k\Delta(T_i - T_j), & if\ T_j < T_i \\ k, & otherwise \end{cases}$$

where k is the thermal conductivity and Δ is a large scaling factor. The solution to the ordinary differential equation in Equation 1 with the correction term applied to handle temperature inversion in Equation 2 computes the temperature of each node as a function of time. Such solution yields an approximation to the temperature profile in the tank (Powell and Edgar 2013).

RULE-BASED CONTROL APPROACH FOR HPWHS

Rule-based control (RBC) is one of the most utilized control methods for HPWHs (Lissa *et al.* 2021). To quantify baseline performance, an RBC method, inspired by the approach used in HPWHSim (Kvaltine, Logsdon, and Larson 2016) is formulated to represent a typical RBC method employed in HPWHs. The objective of the RBC is to maintain the tank temperatures within a temperature range defined by the temperature setpoint minus a deadband, with each heat source having its own deadband. If the HP cannot meet the hot water demand, the RBC will allow the electric resistive elements to come because the heating capacity of the resistive elements is much larger than the HP. The RBC uses a set of logical conditions or rules involving temperature setpoints and deadbands to decide whether one of the three heat sources should be turned on. Only one heat source may be turned on at a moment in time.

The RBC has four deadbands, which are denoted by: $T_{DB,HP}$, $T_{DB,HP,Standby}$, $T_{DB,E1}$ and $T_{DB,E2}$. The upper tank temperature and lower tank temperature are denoted by T_{up} and T_{lo}, respectively. From the four deadbands, the values of the two temperatures may be checked against the four temperature ranges to determine how to operate the three heat sources. The basic idea of the RBC design is as follows. The RBC will turn on HP if either the upper temperature is less than $T_{sp} - T_{DB,HP,Standby}$ or if the lower temperature is less than $T_{sp} - T_{DB,HP}$. If the upper temperature is less than $T_{sp} - T_{DB,E1}$ or the lower temperature is less than $T_{sp} - T_{DB,E2}$, the RBC will turn on E1 or E2, respectively. Under typical conditions, the decision to turn off the heat source occurs once the temperature setpoint is reached. For example, if E1 is turned on, it will be kept on until the upper temperature reaches the temperature setpoint, T_{sp} , as measured by the upper tank temperature sensor. The same concept applies for turning off E2 and the HP.

SUPERVISORY EMPC FRAMEWORK FOR HPWHS

The overall control framework features a cloud-based EMPC deployment. The advantage of the cloud-based EMPC deployment over a local (on-premise) EMPC deployment is the cloud-based EMPC deployment can be used to retrofit existing cloud-connected HPWHs without requiring modifications to the HPWH. However, a disadvantage of a cloud-based deployment is the possibility of communication disruptions or delays. For this reason, many cloud-

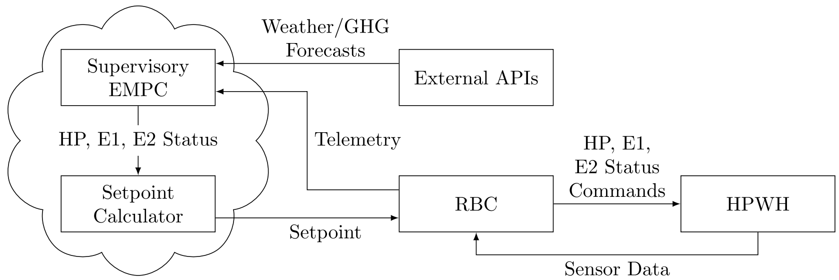

Figure 1 A block diagram of the proposed control architecture.

connected HPWHs do not enable direct control of the HPWH heat sources through the cloud. To address this issue, the EMPC is used as a supervisory controller, meaning that it sends a temperature setpoint to the local HPWH RBC. Since the EMPC computes decisions on how to operate the HPWH heat sources, an additional calculation, referred to as the setpoint calculator, is used to convert the EMPC decisions to a temperature setpoint. The setpoint calculator is also implemented in the cloud. The overall control architecture is depicted in the block diagram in Figure 1.

The overall execution and data communication is as follows. Every 5 minutes, the telemetry data from the HPWH is retrieved from the local device and forecasts of the exogenous inputs are obtained from application programming interfaces (APIs) and forecasting models. The telemetry data and forecasts are provided as input to the EMPC. The EMPC problem, described below, is solved to determine the optimal sequence of on/off decisions for the HPWH heat sources over a future horizon. The on/off decisions for the first element in the sequence are sent to the setpoint calculator to compute the temperature setpoint. The resulting setpoint is communicated from the cloud to the HPWH. The RBC uses the setpoint to make its decision on how to operate the HPWH. This process is repeated indefinitely.

EMPC Problem Formulation

Prediction model. EMPC uses a dynamic model of the system to forecast system behavior over a finite-time window, called the prediction horizon, to compute the control actions that minimize a cost function (Ellis, Durand, and Christofides 2014). In Equation 1, a detailed tank thermal dynamic model is developed for n nodes. This tank thermal dynamic model is nonlinear when $n > 1$ due to the correction term that accounts for natural convection. Thus, a one-node lumped tank thermal dynamic model is used as the predictive model in the EMPC to reduce computational complexity. By discretizing the one-node lumped tank thermal dynamic model in time and implementing a zeroth order hold on all inputs, the resulting predictive model used in the EMPC is given by:

$$\tilde{T}_{k+1} = A(\tilde{d}_k)\tilde{T}_k + B(\tilde{d}_k)\tilde{u}_k + B_d(\tilde{d}_k)\tilde{d}_k \qquad (3)$$

where $\tilde{T}_k$ is the predicted (average) tank water temperature at the kth time step along the EMPC prediction horizon, $\tilde{d}_k$ denotes the forecasted exogenous inputs at time step k, the matrices $A(\tilde{d}_k)$, $B(\tilde{d}_k)$, and $B_d(\tilde{d}_k)$ can be computed from Equation 1, and $\tilde{u}_k := \left[\tilde{u}_{HP,k}\ \tilde{u}_{resist,E1,k}\right]^T$ are the control actions at time step k and the decision variables of the EMPC. Specifically, $\tilde{u}_{HP,k}$ is a binary variable indicating if the HP is on or off at time step k and $\tilde{u}_{resist,E1,k}$ is a binary variable indicating if a resistive element is on or off at time step k. Given that the capacity of the resistive elements is usually equal, there is no difference in the use of either element in the lumped one-node model. In this work, E1 is selected as the decision variable over E2 in the EMPC as this may be preferred to ensure that the water drawn from the top of the tank meets the demand for hot water. Nonetheless, further work is needed to develop a strategy that converts the EMPC's decision of turning on a resistive element to a decision that specifically turns on either E1 or E2. At each time step, the model in Equation 3 is initialized with the average temperature of the two tank measurements taken at time step j.

Cost function. The multi-objective cost function of the EMPC is shown in Equation 4 and includes three terms. The three objective terms are the electricity cost associated with operating the HPWH, the predicted marginal GHG emissions that result from the electricity consumed by the HPWH, and violations of tank water temperature from the minimum and maximum allowable limit.

$$\sum_{k=0}^{N-1} \omega_1 \underbrace{p_{\text{elec},k}\left(\tilde{P}_{HP,k} + \tilde{P}_{resist,E1,k}\right)\Delta t}_{\text{Electricity cost}} + \omega_2 \underbrace{\tilde{p}_{ghg,k}\left(\tilde{P}_{HP,k} + \tilde{P}_{resist,E1,k}\right)\Delta t}_{\text{GHG emissions}} + \quad \omega_3 \underbrace{\tilde{T}_{viol,k}}_{\text{Comfort violations}} \tag{4}$$

where N is the number of time steps in the prediction horizon, $p_{elec,k}$ is the electric tariff at the kth time step, $\tilde{p}_{ghg,k}$ is the forecasted marginal GHG emissions at the kth time step, $T_{viol,k}$ is the violation of the minimum or maximum allowable temperature at the kth time step, $\tilde{P}_{HP,k}$ and $\tilde{P}_{resist,E1,k}$ are the predicted power consumption of the HP and E1 at the kth time step, respectively, Δt refers to the duration between two sample times, and ω_1, ω_2, and ω_3 are weights that are used to manage trade-offs between the multiple objectives that may be in conflict.

The input data required to formulate the cost function include the electric tariff and marginal grid GHG emissions. The electric tariff is assumed to follow a residential time-of-use (TOU) tariff. The marginal grid GHG emissions is provided from an external source. Although the ambient temperature impacts HP efficiency, this study uses a constant efficiency-type model to compute the predicted power consumption of the HP for simplicity. Future refinements will account for the effect of ambient temperature on HP efficiency. A constant efficiency-type model is also used to compute the power consumed by the electric resistive heating element. The weight assigned to an objective is proportional to the relative importance of the objective in the problem. For example, if the objective is purely to minimize cost (referred to as cost minimization), the weights are set to $\omega_1 = 1$ and $\omega_2 = 0$. For purely optimizing GHG emissions (referred to as GHG emissions optimization), $\omega_1 = 0$ and $\omega_2 = 1$. For all cases, ω_3 takes a non-zero value to prevent trading off thermal comfort for benefit of the two other objective terms.

Other Exogenous inputs. The exogenous inputs include hot water draw volumetric flow rate (based on field measurements), inlet water temperature (assumed to be constant at 20 °C), and ambient temperature surrounding the HPWH (from a weather forecast). The inlet water temperature is assumed to be constant with time since a model that estimates this temperature is currently unavailable. A potential future refinement would be to develop a model that predicts the inlet water temperature. A detailed mathematical formulation of the multi-objective EMPC finite-horizon optimal control problem is given in (Mande et al. 2022).

The proposed control architecture is compared to a control architecture involving an EMPC that directly controls the HPWH heat sources. The latter approach is referred to as the regulatory EMPC. The regulatory EMPC has the same optimal control problem formulation as the supervisory EMPC, but its decision is directly communicated to the HPWH. More details of this approach are given in (Mande et al. 2022).

Setpoint Calculator

The objective of the setpoint calculator is to convert the output of the EMPC to temperature setpoints that would prompt the RBC to implement the control actions computed by the EMPC. At each time step, the setpoint calculator receives the predicted input trajectory of the EMPC. The temperature setpoint that is computed over the first sampling period is sent down to the RBC. At the next time step, the process is repeated after receiving the new input trajectory of the EMPC. The basic idea of the setpoint calculator is as follows. If the EMPC decides to turn on one of the heat sources for a given time step, the setpoint calculator sends down a temperature setpoint with a large enough temperature difference to prompt the RBC to turn on the desired heat source. If the EMPC decides to turn off all heat sources for a given time step, the setpoint calculator will send down a low temperature setpoint such that none of the heat sources are expected to turn on.

The specific logic of the setpoint calculator is summarized: If the EMPC predicts the HP to be on at time step k, the setpoint is computed as $T_{sp,k} = \min(T_{max,sp}, T_{up,k} + T_{DB,HP,Standby} + \varepsilon_{HP}, T_{lo,k} + T_{DB,HP} + \varepsilon_{HP})$. If the

EMPC predicts the E1 to be on at time step k, the setpoint is computed as $T_{sp,k} = \min(T_{max,sp}, T_{up,k} + T_{DB,E1} + \varepsilon_{E1})$. If the EMPC predicts all heat sources to be off at time step k, the setpoint is $T_{sp,k} = T_{min,sp}$. $T_{up,k}$ is the upper tank temperature at time step k and $T_{lo,k}$ is the lower tank temperature at time step k. The constants ε_{HP} and ε_{E1} take on values that cause the tank temperature sensor measurements to fall below the setpoint minus the deadband temperature to prompt the desired heat source to turn on. $T_{min,sp}$ is the minimum temperature setpoint and is chosen to be a low temperature setpoint such that tank temperatures are within the temperature setpoint minus the deadband to prevent any heat sources from turning on. If a heat source does turn on after sending down a minimum temperature setpoint, this implies that either the upper or lower tank temperature is too low, in which case turning on a heat source may be desired based on comfort considerations. $T_{max,sp}$ is the maximum temperature setpoint and is selected to be near or equal to the maximum allowable temperature setpoint of the HPWH. The $T_{max,sp}$ and $T_{min,sp}$ are different from the minimum and maximum allowable temperature in the EMPC. Specifically, $T_{max,sp}$ and $T_{min,sp}$ have a wider temperature range compared to the temperature range used in the EMPC to achieve the performance described above. Moreover, it is worth emphasizing that the overall goal of the setpoint calculator is to have the RBC turn the heat source on or off at the times dictated by the EMPC, and not for the tank temperature to reach the computed setpoints. Hence, the resulting temperature setpoint profiles of this framework may seem irregular.

APPLICATION TO A SIMULATED HPWH

In this section, the results obtained from closed-loop simulations consisting of the simulated HPWH under the RBC, regulatory EMPC, and the proposed supervisory EMPC framework are presented and analyzed. The simulated HPWH has the tank configuration discussed in the earlier section. The dynamic temperature profile of the HPWH is assumed to be perfectly forecasted by the tank thermal dynamic model with one node in Equation 1. The HPWH model parameters are summarized in (Mande *et al.* 2022). Perfect forecasting of all exogenous inputs is assumed in this work to establish a baseline performance that can later be compared to the case when uncertainty in the forecasts is introduced. The exogenous input profiles are shown in Figure 2a. A residential TOU electric tariff structure is considered. The peak period occurs from 4 PM to 9 PM with a peak price of $0.50/kWh. The off peak-price is $0.36/kWh. The marginal GHG emissions data is provided by WattTime. The water volumetric flow rate data is collected from a field site in California Climate Zone 12. The weather data is obtained from OpenWeather.

For the EMPC, the minimum and maximum allowable temperature of the average water temperature is selected to be 42 ℃ (107.6 °F) and 50 ℃ (122 °F), respectively. This choice corresponds to the approximate temperature range of the HPWH under the RBC with setpoint 50 ℃ (122 °F) and the deadband $T_{DB,HP,Standby}$ = 8.33 ℃ (15 °F), which gives a range of 41.67 ℃ (107 °F) to 50 ℃ (122 °F). The minimum temperature bound of the EMPC is rounded to 42 ℃ (107.6 °F) because of the EMPC's tendency to regulate the average water temperature near the lower bound for periods to save energy (see results below), which can result in small temperature violations from the lower bound. Thus, the EMPC temperature range (for both the supervisory case and the regulatory case) is restricted to a smaller range to ensure that the average water temperature is maintained within 41.67 ℃ (107 °F) to 50 ℃ (122 °F). The minimum and maximum temperature setpoint are set to $T_{min,sp}$ = 41 ℃ (105.8 °F) and $T_{max,sp}$ = 55 ℃ (131 °F), respectively. The ε_{HP} and ε_{E1} are assigned a value of 1 ℃ (1.8 °F). The supervisory and regulatory EMPC problem are solved under cost optimization and with a weight value of ω_3= 2 applied to the comfort violations term in the cost function. The EMPC problem is solved using CPLEX. The absolute gap tolerance is set to 0.001 with a solver time limit of 20 s. The solver returns a solution that either met the gap tolerance or time limit. A controller time step of 5 minutes is used in the EMPC, with a prediction horizon of 24 hours to consider a full diurnal cycle. A minimum on dwell-time of 10 minutes is enforced to limit excessive cycling of the EMPC. The average water tank temperature is initialized at 44 ℃ (111.2 °F) for all closed-loop simulations. All simulations begin at midnight where minimal hot water demand is expected.

For baseline purposes, the RBC is applied to the HPWH, and the closed-loop results are given in Figure 2b. A large hot draw event occurs in the morning, likely caused by occupants showering. Prior to the first morning water draw, the RBC turns on the HP as the tank temperature dropped below the HP deadband. The HP turns on for several hours to return the tank water temperature to setpoint. After the setpoint is reached, the RBC turns off the HP. Over the day, several hot water draw events occur, which eventually causes the tank temperature to decrease to its minimum. Once the temperature is at its minimum, the RBC turns on the HP to return to the water temperature setpoint. However, this occurs during the peak period (Figure 2b). The estimated cost associated with this operating strategy is \$0.693. Its estimated GHG emissions is 1.78 lb.

The closed-loop results of the regulatory and supervisory EMPC are given in Figure 2c and 2d, respectively. For the supervisory EMPC, the on/off decisions implemented by the RBC matched those computed by the EMPC, indicating that the setpoint calculator can translate the EMPC's decisions to the appropriate temperature setpoints. For both the regulatory and supervisory EMPC, none of the electric resistance elements turned on and no temperature violations occurred. Thus, only the HP decisions of the RBC are displayed in Figure 2d. Over the first six hours, both the regulatory and supervisory EMPC pre-heat the tank to prepare for the first expected draw event of the day. The first draw event causes the temperature of the tank to decrease but is still above the lower bound. Between 6 AM and 9 AM, both the regulatory and supervisory EMPC keep the tank temperature near the minimum for energy savings. Prior to the peak period, both the regulatory EMPC and supervisory EMPC pre-heat the tank near the maximum allowable temperature to minimize the utilization of the HP during this time window. However, due to multiple hot water draw events that occurred in the peak period, both the regulatory EMPC and supervisory EMPC turned on the HP to prevent temperature violations. During the peak period, the regulatory EMPC total HP runtime is 35 minutes, whereas the supervisory EMPC total HP runtime is 25 minutes. The estimated operating cost of the regulatory EMPC is \$0.607, whereas the estimated operating cost of the supervisory EMPC is \$0.600. In terms of GHG emissions, the regulatory EMPC GHG emissions is 1.54 lb and the GHG emissions of the supervisory EMPC is 1.53 lb. The difference in closed-loop evolution under the two EMPCs is caused by the solver not able to completely close the optimality gap in the time needed, resulting in different solutions returned in each case.

From the simulation results in Figure 2, the operating behavior under EMPC cost-optimization (regulatory and supervisory case) is different from the operating behavior under the RBC, resulting in lower cost under EMPC but with more HP cycling. If this cycling is deemed undesirable as it causes more wear on the HP, the minimum on-time can be increased greater than 10 minutes. The effect of increasing minimum on time to the operating cost is investigated under cost-optimization. Two minimum on times for the HP are considered: 30 minutes and 1 hour. The operating cost increases as the minimum on-time increases. Specifically, for a minimum on-time of 30 minutes, the operating cost of regulatory and supervisory EMPC are \$0.625 and \$0.631, respectively. In the case of 1 hour minimum on-time, the operating cost of regulatory and supervisory EMPC is \$0.664.

CONCLUSION

A cloud-based supervisory EMPC framework for HPWHs was proposed to minimize electricity cost and GHG emissions rate. In the proposed framework, the cloud-based EMPC sends down a temperature setpoint to the local HPWH RBC as opposed to directly controlling the HPWH heat sources. The RBC, regulatory EMPC, and the supervisory EMPC were applied to a simulated HPWH to analyze the behavior of the HPWH under the three control strategies. From this analysis, the RBC (in the supervisory EMPC case) was able to implement the on and off decisions computed by the EMPC, which indicates that the setpoint calculator was able to translate the EMPC's decisions to the appropriate temperature setpoints. The effect of increasing minimum on-time was also investigated and showed that minimizing HP cycling results in increased operating cost. In all cases considered, the EMPC reduced operating costs and expected GHG emissions compared to the cost and expected GHG emissions under RBC.

ACKNOWLEDGEMENTS

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

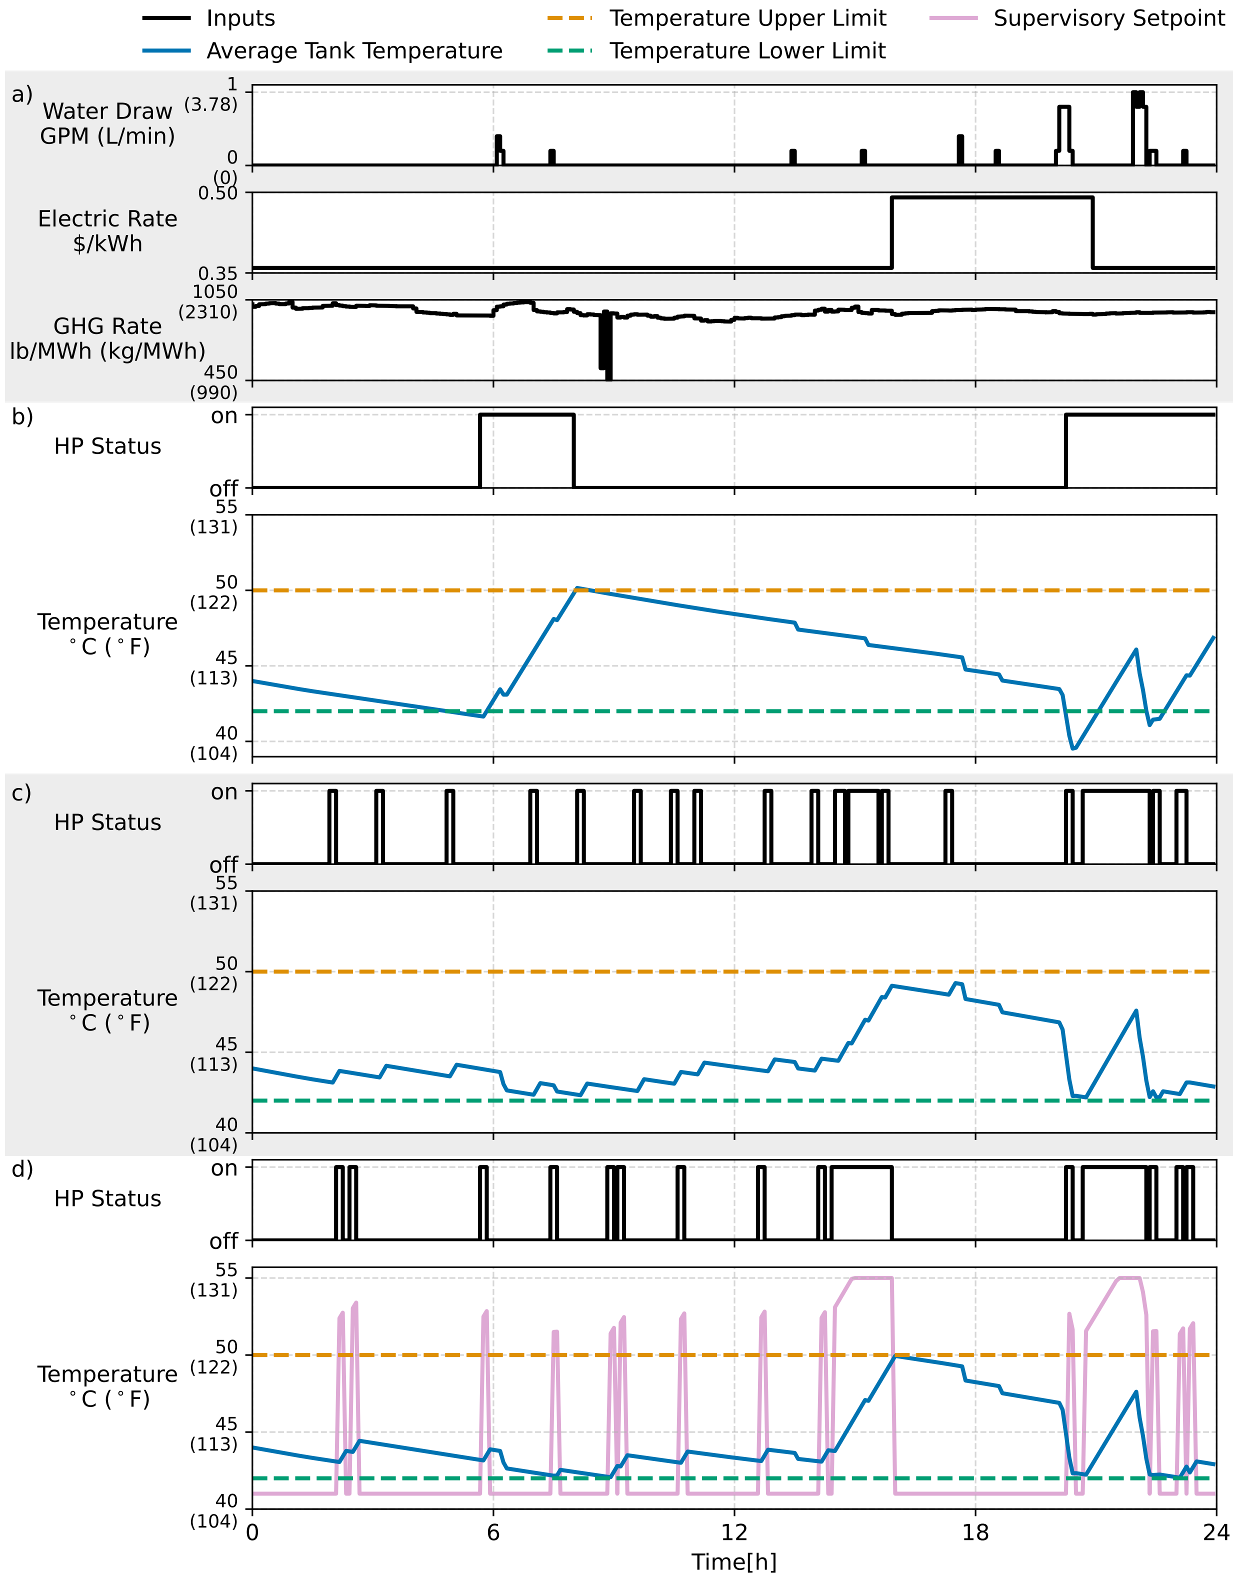

Figure 2 The (a) exogenous inputs and HP status and average water temperature under (b) RBC, (c) regulatory EMPC, and (d) supervisory EMPC under cost-optimization.

REFERENCES

Delforge, P. and J. Vukovich. 2018. Can Heat Pump Water Heaters Teach the California Duck to Fly? *Proceedings of the 20th ACEEE Summer Study on Energy Efficiency in Buildings*, *Pacific Grove, CA*.

Delforge, P. 2020. "Heat Pump Water Heaters as Clean-Energy Batteries." Natural Resources Defense Council, January 7. www.nrdc.org/experts/pierre-delforge/heat-pump-water-heaters-clean-energy-batteries.

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.

EIA (Environmental Impact Assessment). 2018. "Space heating and water heating account for nearly two thirds of U.S. home energy use." www.eia.gov/todayinenergy/detail.php?id=37433#.

Kleinbach, E. M., W.A. Beckman., and S.A. Klein. 1993. Performance study of one-dimensional models for stratified thermal storage tanks. *Solar Energy* 50:155-66.

Kvaltine, N., M. Logsdon, and B. Larson. 2016. A Guide to HPWHsim. Ecotope. https://github.com/EcotopeResearch/HPWHsim/blob/master/doc/Guide_to_HPWHsim_v1_1.pdf

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.

Mande, C., A. Aboud, L. dela Rosa, R. Colins, S. Outcault, and M.J. Ellis. 2022. Timing is everything: Optimizing load flexibility of heat pump water heaters for cost, comfort, and carbon emissions. *Proceedings of the 22nd ACEEE Summer Study on Energy Efficiency in Buildings*, *Pacific Grove, CA*.

Nash, A.L., A. Badithela, and N. Jain. 2017. Dynamic modeling of a sensible thermal energy storage tank with an immersed coil heat exchanger under three operation modes. *Applied Energy* 195:877-89.

Powell, K. M. and T. F. Edgar. 2013. An adaptive-grid model for dynamic simulation of thermocline thermal energy storage systems. *Energy Conversion and Management* 76:865-73.