Model Predictive Control of Heat Pump Water Heaters for Energy Efficiency

Xin Jin[1], Jeff Maguire and Dane Christensen, National Renewable Energy Laboratory

ABSTRACT

Water heating accounts for 17.7% of total residential energy use in the United States and is the second largest end use after space heating and cooling. There is great potential to improve the energy efficiency of residential water heaters and reduce their energy use in new and existing homes. Previous efforts primarily focused on improving the insulation and combustion efficiency of water heaters, whereas little effort has been made from the control perspective. Heat pump water heaters (HPWHs) provide an energy efficient solution for water heating. Instead of generating heat directly, HPWHs transfer heat from the environment into the water in the tank. The heat pump is two to three times more energy efficient than resistance elements, although HPWHs typically include elements for backup and high demand situations. We propose a model predictive control (MPC) framework that aims to achieve maximum energy savings while maintaining thermal comfort. This framework uses algorithms that automatically learn users' hot water consumption patterns and adaptively increase the use of the heat pump to avoid resistance element use. It has been tested through simulations with hot water draw profiles collected from field tests. Simulation results indicate that this technique will save up to 20% for a low user or more than $20 per year.

Introduction

The building sector is the largest consumer of energy and accounts for 40% of primary energy use in the United States, more than the industrial and transportation sectors (U.S. Energy Information Administration 2012). Residential buildings' energy use exceeds that of commercial buildings, consuming 22.5% of the total primary energy. Water heating comprises 17.7% of energy use in residential buildings, and is the second largest end use after space heating and cooling (U.S. Energy Infromation Administration 2009). There is great potential to improve the energy efficiency of residential water heaters, which is a key step to achieving significant energy reduction in new and existing homes.

Many improvements can be made to increase the energy efficiency of residential water heaters. Previous efforts primarily focused on improving the jacket insulation, adding insulation to the inlet and outlet piping, and increasing the combustion efficiency of fuel-fired water heaters (Hirst and Hoskins 1978). Little effort has been made to use more advanced control strategies to increase efficiency. Most water heaters use a hysteresis strategy that uses setpoints and deadbands to control the heat sources reactively. These simple control methods are widely adopted in modern residential water heaters because of their low implementation cost. However, most use a fixed thermostat setpoint and are unsuitable for implementing advanced control algorithms to generate additional energy savings.

Heat pump water heaters (HPWHs) offer an energy efficient solution for water heating. Unlike conventional storage water heaters that generate heat directly, HPWHs transfer heat from

[1] Current Affiliation: A. O. Smith Corporate Technology Center

the ambient air into the water in the tank. HPWHs usually have a rated efficiency (energy factor) between 2 and 2.75 and are the highest efficiency electric water heaters available other than solar. HPWHs in the United States typically feature both a heat pump and at least one electric resistance element for heating, which is used for backup and high demand situations. Their multiple heat sources make them perfect candidates for implementing advanced control algorithms to reduce energy use. Algorithms can be used to minimize the use of the electric resistance element and increase the use of the heat pump, thereby improving the overall energy efficiency of the water heater. However, there are technical challenges in implementing these energy saving measures. One major challenge is predicting the users' hot water consumption. If high use periods could be anticipated, the heat pump could preheat the tank, reducing element use. We propose a novel framework to bridge the current gaps in water heater control, event prediction, and advanced control theory. Based on model predictive control (MPC), this framework aims to optimally modulate the thermostat setpoint and increase heat pump use without impacting thermal comfort. Under this framework, algorithms have been developed to identify patterns from past hot water use and estimate future hot water draw events.

Recent research on water heater controls can be broadly divided into two categories: low level (on and off) control of water heaters for demand side management and high level control of individual water heaters for thermal comfort or energy efficiency. The simplified low order (single node) tank models used in demand side management (Paull, Li, and Chang 2010; Kondoh, Lu, and Hammerstrom 2011) may not be suitable for studies where energy efficiency and thermal comfort of individual water heaters are the primary focus. For low level control, Healy et al. concluded that a linear time invariant model relating input rate and output rate may not be appropriate for fully describing water heater performance (Healy, Ullah and Roller 2011). Henze and Yuill (2009) successfully demonstrated MPC on a tankless water heater in real time; however, this technique does not apply to storage water heaters because it minimizes outlet temperature error, which is not an issue for storage water heaters.

Energy efficiency of residential water heaters can be further improved by using advanced control techniques. Standby losses can comprise a significant amount of the total energy consumption. Lowering the thermostat setpoint temperature when hot water demand is low can reduce these losses. One research question for our paper is how to modulate the thermostat setpoint to reduce energy consumption while maintaining thermal comfort. This leads to the other research question of this paper: how to effectively coordinate the heat sources of the HPWH to improve energy efficiency. To answer these research questions, it is critical to predict hot water usage patterns from past events such that optimal control strategies can be created.

Modeling of Heat Pump Water Heaters

Modeling of storage water heaters is extensively studied by the research community, and various models have been developed to meet different research objectives (Kondoh, Lu, and Hmmerstrom 2011; Paull, Li, and Chang 2010; Nehrir, Jia, and Pierre 2007). These models are not suitable for our study because they either do not consider thermal stratification or are too computationally intensive. Dynamic one-dimensional (1-D) models for storage tanks have been used in several simulation tools (Klein 2010; Crawley 2001). These tools apply mass and energy conservation to isothermal spatial zones called *nodes*. Radial and axial variations in temperature in the tank are neglected. These models allow for thermal stratification and provide a good balance between run time and model complexity (Burch 2011).

Modeling of Thermally Stratified Tanks

A dynamic 1-D model is developed for the storage water heaters in this project. A generic tank model is shown schematically in the left part of Figure 1. It is divided into a series of isothermal vertically stacked nodes. The number of nodes is chosen as a tradeoff between model accuracy and computational workload (Maguire 2012).

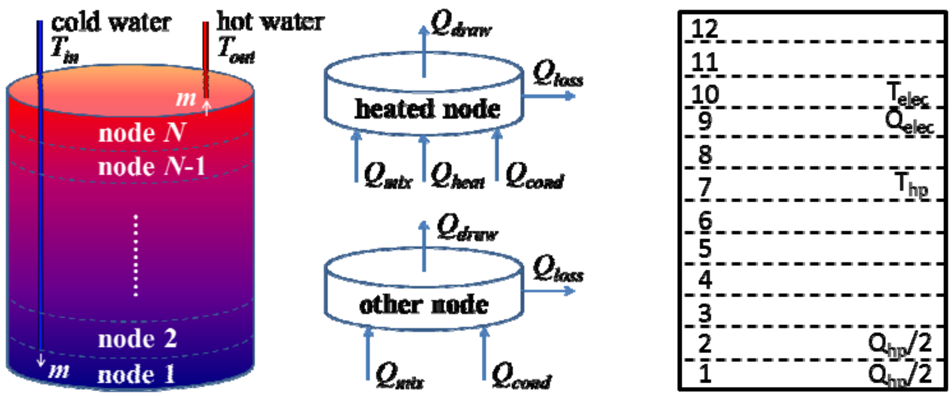

Figure 1. Schematics of thermally stratified tank models. Left: generic model; Right: HPWH model.

In this model, higher nodes always have a temperature greater or equal than lower nodes to capture stratification caused by buoyancy. Mixing occurs between nodes to ensure that this stratification is always maintained.

The following 1-D model is based on an instantaneous energy balance for each node:

$$C_i \frac{dT_i}{dt} = \dot{Q}_{heat,i} - \dot{Q}_{loss,i} - \dot{Q}_{draw,i} + \dot{Q}_{cond,i} + \dot{Q}_{mix,i} \quad (1)$$

where C_i is defined as the product of the mass of water in node *i* and heat capacity of water, T_i is the temperature of node *i*, and $\dot{Q}$ denotes the rate of change in thermal energy. Equation (1) models heating from either source $\dot{Q}_{heat}$, standby loss $\dot{Q}_{loss}$, draw $\dot{Q}_{draw}$, conduction $\dot{Q}_{cond}$, and mixing in the tank $\dot{Q}_{mix}$. Detailed explanation of individual terms can be found in (Jin, Maguire and Christensen 2014).

Table 1. Parameters of the HPWH model

Items	Values
Nominal volume	66 gallons (0.238 m^3)
Heat pump deadband	18°F (10°C)
Electric element deadband	15°F (8.3°C)
Power of electric element	4.5 kW
Rated compressor power	0.80 kW
Standby power	2.39 W

Equation (1) is a generic model for storage water heaters. Locations of heat sources and heat input rates vary for different water heaters. By specifying the location and the input rate of the heat sources, Equation (1) can be applied to any type of storage water heater. In this work, an HPWH model is developed for a specific HPWH. As shown in the right part of Figure 1, the entire tank is uniformly divided into 12 nodes. The immersed condenser coil is located at the bottom of the tank in nodes 1 and 2. The temperature sensor for the heat pump is located in node 7. The element is located in node 9, providing 4.5 kW of heat capacity in case the heat pump cannot keep up with the demand. The temperature sensor for the element is located in node 10. Table 1 summarizes the parameters of the HPWH model.

State Space Model

State space models are widely used to study the properties of dynamic systems. To simulate the non iterative part of Equation (1) for this study, a continuous time state space model can be developed by neglecting the conduction and mixing terms:

$$\dot{x}(t) = A_c(t)x(t) + B_c(t)u(t) \tag{2}$$
$$y(t) = C_c x(t) \tag{3}$$

where state $x = [T_1, T_2, \cdots, T_{12}]^T$, control $u = [v_{hp}, v_{elec}, T_{env}, T_{in}]^T$, T_i is the temperature of node *i*, T_{env} is the ambient temperature, T_{in} is the inlet water temperature, and v_{hp} and v_{elec} are the binary control signals of the heat pump and the element, respectively. The definitions of A_c, B_c and C_c matrices can be found in (Jin, Maguire and Christensen 2014). Hysteresis controllers are used to control the HPWH based on temperature differentials and user specified setpoints. The conduction and mixing terms are accounted for through iteration.

The state space model is a linear time variant system because matrices A_c and B_c vary with time if flow rate changes. To incorporate flow rate and temperature information that is recorded in discrete time, a discrete time state space model is more suitable for our study. Therefore, Equations (2) and (3) can be discretized, assuming zero-order hold for the input *u*, to

$$x[k+1] = A_d[k]x[k] + B_d[k]u[k] \tag{4}$$
$$y[k] = C_d x[k] \tag{5}$$

where A_d and B_d can be computed by utilizing the following property with the sample time τ

$$e^{\begin{bmatrix} A_c & B_c \\ 0 & 0 \end{bmatrix}\tau} = \begin{bmatrix} A_d & B_d \\ 0 & I \end{bmatrix} \tag{6}$$

Because the system is time variant, A_d and B_d need to be updated at every time step. C_d is the same as C_c. The choice of τ is a trade-off between model accuracy and implementation speed. $\tau = 1$ minute is adopted as the sample time in this paper.

The state space model of the HPWH has been implemented in MATLAB and validated with laboratory test data whereby a full day draw profile was imposed on a test article under closely monitored operation. Both consumed electric energy and delivered energy of the HPWH model are compared with the data collected from laboratory test. As shown in Table 2, differences for both terms are within 4%, indicating the model is a reasonable approximation of the HPWH.

Table 2. HPWH model validation

	Consumed energy	Delivered energy
MATLAB model	7.39 kWh	16.84 kWh
Laboratory data	7.66 kWh	16.25 kWh
Difference	-3.81%	3.65%

Model Predictive Control of Heat Pump Water Heaters

MPC refers to a range of control methods that use a model to obtain the control signal by minimizing an objective function (Camacho and Bordons 2007). Typical MPC structure consists of a dynamic model of the process or system, a history of past control actions, and an objective function over the receding prediction horizon to calculate optimal control strategies. MPC is an iterative optimization process over a finite prediction horizon. It seeks the control strategies that minimize the objective function, subject to certain constraints. Also known as *receding horizon control*, MPC implements only the first step of the obtained control strategy and repeats the entire process as the prediction horizon is shifted forward.

Different types of models have been used as the prediction model in the MPC structure. Linear input/output models are among the most popular because of their simplicity. The model can be either derived analytically from a physics based model (Henze and Yuill 2009), or generated numerically using system identification techniques (Ma et al. 2012). Storage water heaters are complex systems with inherent nonlinearities caused by thermal stratification. Time invariant models are typically used in MPC. However, tank dynamics change fast and are strongly affected by draws. As a result, a time invariant model cannot capture the dynamics of the tank.

Most MPC algorithms use a quadratic cost function to minimize the weighted sum of control actions and deviation from reference signals. Inspired by this cost function structure, we created a custom cost function that considers both energy savings and thermal comfort. The cost function contains a weighted sum of an energy consumption term and a temperature sag term, which are analogous to the control action and deviation from reference signals in the original structure, respectively.

Framework of Model Predictive Control for Heat Pump Water Heaters

As shown in Figure 2, the MPC takes measurements from the HPWH simulated via the state space model, calculates the optimal control strategies based on measurements, and sends the setpoint profiles back to the HPWH as the control signal. In our MPC framework, we choose to control the setpoints rather than heat sources to reduce instrumentation cost and prevent overheating. The right half of Figure 2 shows the detailed structure of the controller, which consists of three modules: draw volume prediction, a simplified HPWH model, and setpoint optimization.

MPC of water heaters is a challenging problem because disturbances such as hot water draws are the dominating factor affecting operation. To effectively control the water heater, proactively responding to the upcoming draw events is more desirable than reactively turning on the heat sources when the tank temperature is too low. For this reason, it is important to estimate the hot water draw volume based on past hot water consumption and incorporate the estimation into the MPC framework. Flow rate is the input of the draw volume prediction model, and can be measured by a flow meter or estimated based on changes in tank temperature.

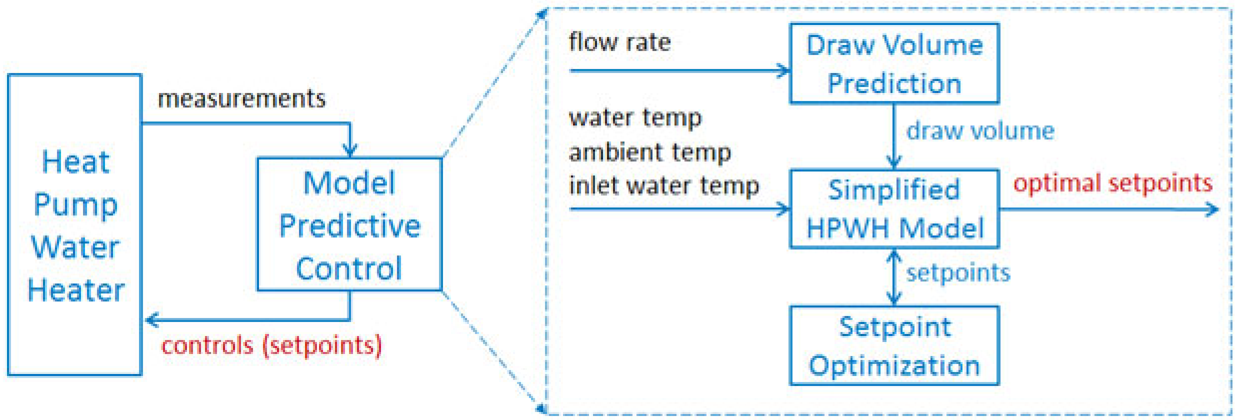

Figure 2. Framework of model predictive control for HPWH.

A model is needed to predict the tank water temperature given the estimated draw volume and setpoint temperature. Although the 12 node model described above captures the tank dynamics well, it is meant to be a virtual test bed of the actual HPWH and is too computationally intensive to serve as the prediction model. As a tradeoff between computational load and model accuracy, we use a simplified 2 node model as the prediction model, which is described below.

The goal of the MPC controller is to generate an optimal setpoint profile for the heat pump and element. A simple but effective approach is exhaustive search. The costs associated with all possible combinations of setpoints are calculated, and the profile that generates the minimum cost is selected as the optimal. The computation load of this method depends on the size of the search space, the length of the prediction horizon, and the complexity of the prediction model. The objective function used in setpoint optimization is defined as follows:

$$R^{opt} = \underset{i}{\operatorname{argmin}} \sum_{j=1}^{H} \{\alpha c_{rate}(j)\left[P_{hp}(R(i)) + P_{elec}(R(i))\right] + \beta Sag(j)\} \tag{7}$$

where R^{opt} is the optimal setpoint profile, H is the prediction horizon (in terms of number of steps), α is the weighting function for consumed energy, c_{rate} is the electricity rate (a flat rate is used here, but time-of-use [TOU] rates or critical peak pricing [CPP] could be accommodated), P_{hp} is the power consumed by the heat pump, P_{elec} is the power consumed by the element, R is the current setpoints being evaluated, β is the weighting function for thermal comfort, and Sag is an indicator of thermal comfort.

Equation (7) aims to find the optimal setpoint profile R^{opt} that minimizes the weighted sum of energy consumption and thermal comfort over the prediction horizon H. When using a flat rate, MPC minimizes the consumed energy and equivalently minimizes the energy cost. If a TOU pricing structure is used, the proposed framework could be used with demand response analysis, thereby saving energy cost for end users and shifting load to off peak periods.

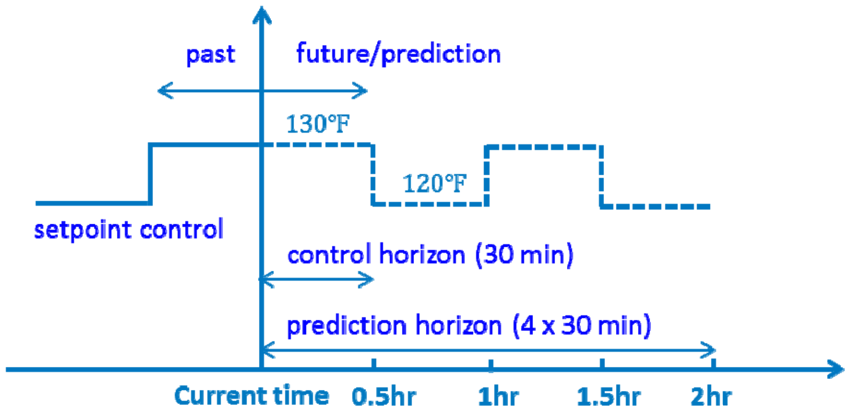

Figure 3. Illustration of the receding horizon concept in model predictive control.

Figure 3 shows the receding horizon concept in the proposed MPC framework. The controller time step is 30 minutes and the prediction horizon is four time steps, totaling two hours into the future. At the current time, the setpoint profile that minimizes the cost over the next two hours is found, and the first step of the setpoint profile is implemented for the next 30 minutes. The entire process is repeated at the end of the first step, and a new profile is generated.

Draw Volume Prediction

Accurate prediction of future draw events is critical for the success of the proposed MPC framework. We present a new prediction algorithm that is designed to predict the draw volume for each time step. The entire day is uniformly divided into 48 bins, each with a width of 30 minutes. The average daily draw volume of the past 10 days in each bin is calculated and serves as the predicted draw volume for each time step. The 10 day window is sliding and updated at the end of each day, so the prediction algorithm can capture new patterns in hot water consumption. Weekday and weekend data are considered separately, as hot water consumption patterns are often significantly different for weekdays vs. weekends.

Simplified Heat Pump Water Heater Model

A prediction model is needed to estimate the future status of the physical system and find the control strategy that minimizes the cost function. The search process could be computationally expensive, especially when the prediction model is complex or the search space is large. System identification techniques have been used to identify numerical models of input/output relationship in HPWHs. These identified models, however, are not valid if operating conditions change due to the nonlinearity and transport delay in the tank.

As shown above, the 12 node HPWH model is a good approximation of the actual system; however, it is not a good prediction model for MPC for two reasons: (1) computational costs are high because many nodes and iterations are required; and (2) temperature measurements are needed for each node to update the model status, which makes the instrumentation cost prohibitive and reduces the reliability of the entire system. Therefore, a simplified model is needed.

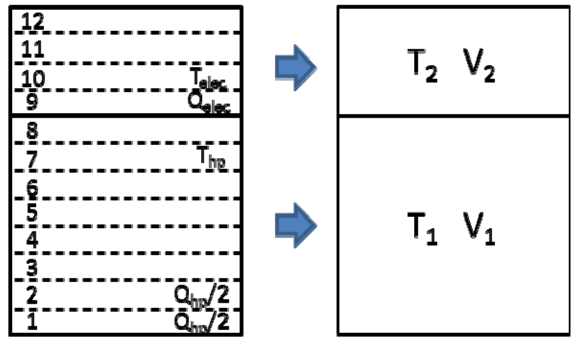

Figure 4. Layout of the simplified two node HPWH for prediction.

Figure 4 presents the layout of a 2 node HPWH prediction model, which is formulated by combining the nodes in the original 12 node HPWH model. As shown in Figure 4, we combine nodes 1-8 to form a new node of volume V_1, and combine nodes 9-12 to form a new node of volume V_2. We divide the nodes in this way mainly because the temperature sensors for the heat pump and the element are located in node 7 and node 10, respectively. By combining nodes in this way, the temperature measurements taken from the sensors could be directly used to update the states of the prediction model, so no additional temperature sensors are needed.

The 2 node model has been compared with the 12 node model using a realistic 30 day long draw profile generated by the Domestic Hot Water Event Schedule Generator (DHWESG) (Hendron, Burch and Barker 2010). The validation results are shown in Table 3, where the total energy use, heat pump energy use, electric element energy use, and delivered energy are compared between the 12 node model and the 2 node model. Overall, the results of both models are close enough to consider the 2 node model a reasonable approximation for MPC: the energy use difference is 2.40% and the delivered energy difference is -5.21%. The 2 node model overestimates the heat pump energy use by more than 8%, because T_1 in the new model is lower than T_{hp} in the original model, so the heat pump turns on more often. Similarly, T_2 is in general higher than T_{elec}, so the element of the 2 node model consumes less energy.

Table 3. Model validation using DHWESG profiles (30 days)

Model type	12 node	2 node	Difference
Total energy use (kWh)	134.80	138.04	2.40%
Heat pump (kWh)	95.12	104.11	8.63%
Electric element (kWh)	38.25	36.38	-4.89%
Delivered energy (kWh)	327.61	310.53	-5.21%

Setpoint Optimization

An exhaustive search algorithm is used to find the optimal setpoints for the HPWH. Two candidate setpoints are defined for both the heat pump and the element. In the current study, we constrained the heat pump to switch between high and low values (in this study, combinations of 120°F/130°F and 120°F/140°F were used), and the element may switch between 120°F/115°F. These setpoints were chosen to minimize electric element use and maximize heat pump use. More setpoints can be included in the procedure at the cost of increased computation, which might require a different approach than an exhaustive search. There are in total $4^{2\times 2} = 256$ setpoint profiles to be explored and the associated cost to be computed. Additional constraints can be applied to reduce the number of candidate setpoint profiles and expedite the optimization process. As an initial attempt, these constraints are not considered in the current paper.

Simulation Results

This section presents the simulation results of the proposed framework using field test data collected from two studies: a low use (35 gal/day) home in Boulder, CO (Barley, Hendron and Magnusson 2010) and a high use (92 gal/day) home in Sacramento, CA (Maguire et al. 2011). Long term monitoring equipment was installed to collect water use data from these homes for more than one year. Two performance metrics—energy savings and temperature sag—are used to evaluate the effectiveness of the proposed MPC algorithm. Temperature sag is defined as the additional amount of energy needed to heat the water to a level that would be acceptable for users when the outlet temperature drops below a comfortable level (assumed here to be 110°F). Tracking this metric is important to ensure that the control algorithms do not provide energy savings at the expense of providing inadequate hot water.

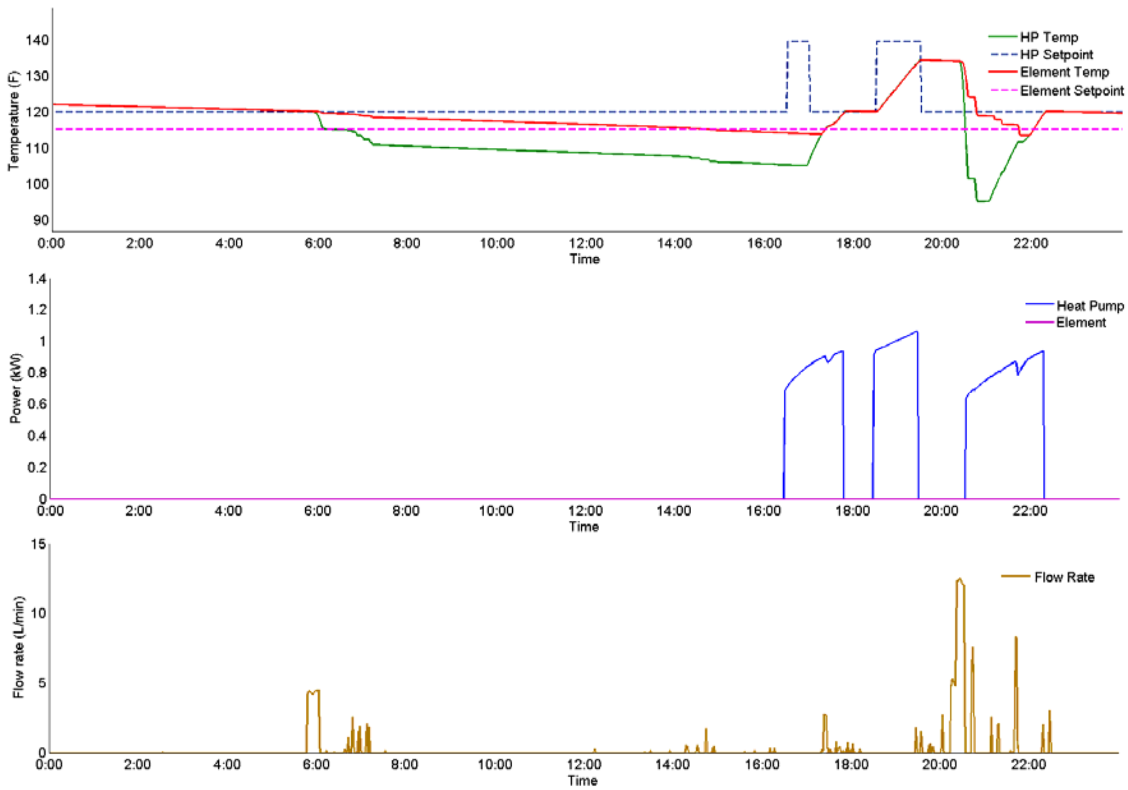

Figure 5. Temperature, setpoint, power consumption, and draw for one day with MPC.

Figure 5 shows the performance of the MPC method for one day in the high use home. On this day, past draw patterns predict large draws in the evening. Based on this prediction, the heat pump setpoint is raised to 140°F for two periods. Although there is only a small draw after the first period when the heat pump setpoint is raised, there is a large draw after the second period. Preheating the tank ahead of this draw allows the heat pump to meet this load instead of using the electric resistance element. This illustrates that although day to day variations in hot water use prevent the model from perfectly predicting large draw events, it can identify when hot water events are likely to occur and respond appropriately.

For comparative evaluation of the advanced controls, two baseline simulations are also conducted using different setpoints for each home. The first baseline uses 120°F as the setpoint for both the heat pump and the element. The second baseline uses the highest temperature the heat pump used (130°F for the low use home and 140°F for the high use home) and 120°F as the setpoint for the element. A higher temperature was used in the high use home to ensure that the HPWH was able to provide adequate hot water.

Table 4 shows the results for the MPC cases and baseline cases for both low and high end uses. Total energy consumption, as well as energy consumed by the heat pump and the element, is presented, along with temperature sag. Table 4 shows that MPC consumes less total energy than either baseline in both cases. The MPC case can provide better thermal comfort than baseline 1; however, baseline 2 provides the best thermal comfort of the cases considered here at the expense of much higher energy consumption. Thermal comfort is a highly subjective measure, based on occupant preferences. However, simulation results for an average use home (using a draw profile from the DHWESG) with a typical electric storage water heater would in a moderate climate would have a larger temperature sag than any of the cases explored here.

Table 5 summarizes the annual savings achieved by MPC in terms of both energy and cost. Cost savings were calculated using the 2012 average rate of $0.1188/kWh. Compared to baseline 1, MPC yields up to 40 kWh annual energy savings. Life cycle cost savings should be able to recover the instrumentation cost needed to implement the MPC technique in an HPWH relative to this baseline and still achieve greater thermal comfort. Compared to baseline 2, MPC achieves about 170-190 kWh of annual energy savings and more than $20 annual cost savings, which would provide net life cycle cost savings for a homeowner.

Table 4. Results of annual simulation for the two draw profiles

Case	Method	Setpoint (°F)		Energy consumption (kWh)			Temperature sag (kJ)
		Heat pump	Electric	Total	Heat pump	Electric	
Low use	Baseline 1	120	120	736.7	599.9	117.6	121.8
	Baseline 2	130	120	867.8	821.9	27.2	2.9
	MPC	130/120	120/115	697.3	650.9	27.4	46.0
High use	Baseline 1	120	120	1754	1192	541.1	1798
	Baseline 2	140	120	1906	1740	145.0	332.2
	MPC	140/120	120/115	1715	1478	215.9	1548

Table 5. Comparison of energy and cost savings between MPC and baseline methods

Case	Metric	MPC vs. baseline 1	MPC vs. baseline 2
Low use	Annual energy savings	39.37 kWh (5.3%)	170.48 kWh (19.6%)
	Annual cost savings	$4.68	$20.25
High use	Annual energy savings	38.5 kWh (2.2%)	190.7 kWh (10.0%)
	Annual cost savings	$4.57	$22.65

Conclusions

We present a novel MPC framework for HPWHs in this paper. Unlike traditional MPC, the proposed method has a modified structure that is customized for HPWH control. The proposed framework uses an exhaustive search method to find the optimal setpoint profile that maximizes energy savings and thermal comfort based on the draw volume estimated from past hot water draw events. Significant energy and cost savings can be achieved with only a modest decrease in thermal comfort. The proposed framework is very flexible, and can incorporate TOU or CPP pricing structures in its optimization cost function to create additional cost savings. It could also be slightly modified to incorporate demand response (load add/shed) capabilities.

Future research directions include improving computational efficiency, reducing the instrumentation cost and testing on actual hardware. A more efficient search method is needed to reduce the computation complexity and make the algorithm suitable for implementation in embedded systems. We are also actively looking for alternative techniques to measure temperature and flow rate in a cost effective manner. Finally, the framework should be implemented in prototypes for laboratory and field demonstrations.

References

Barley, D., B. Hendron, and L. Magnusson. 2010. "Field Test of a DHW Distribution System: Temperature and Flow Analysis (Presentation)." *2010 ACEEE Hot Water Forum.* Ontarion, California: ACEEE, May 13, 2010.

Burch, J. 2011. "Deriving simulation parameters from rating data for stratified storage tank water heaters." *Proceedings of the Solar 2011 Conference.* 2011. 111-118.

Camacho, E.F., and C. Bordons. 2007. *Model Predictive Control.* Springer.

Crawley, D. 2001. "EnergyPlus: Creating a New-Generation Building Energy Simulation Program." *Energy and Buildings* 33, no. 4: 443-457.

Henze, G.P., and D.P. Yuill. "Development of a Model Predictive Controller for Tankless Water Heaters." *HVAC&R Research* 15, no. 1 (2009): 3-23.

Healy, W.M., T. Ullah, and J.F. Roller. "Input-Output Approach to Predicting the Energy Efficiency of Residential Water Heaters – Testing of Gas Tankless and Electric Storage Water Heaters." *ASHRAE Transactions.* 2011. 521-532.

Hendron, R., J. Burch, and G. Barker. 2010. "Tool for Generating Realistic Residential Hot Water Event Schedules." *SimBuild 2010.* New York, NY, pp. 1-8.

Hirst, E., and R. A. Hoskins. 1978. "Residential Water Heaters: Energy and Cost Analysis." *Energy and Buildings* 1, no. 4: 393-400.

Jin, X., J. Maguire, and D. Christensen. 2014. *Advanced Control of Residential Water Heaters for Energy Efficiency.* Technical Report, Golden, CO: National Renewable Energy Laboratory. (Publication pending)

Klein, S. *TRNSYS Users Manual.* Madison, WI: University of Wisconsin Solar Energy Laboratory, 2010.

Kondoh, J., N. Lu, and D. Hammerstrom. 2011. "An Evaluation of the Water Heater Load Potential for Providing Regulation Service." *IEEE Tran. on Power Systems* 26, no. 3: 1309-1316.

Ma, J., J. Qin, T. Salsbury, and P. Xu. 2012. "Demand Reduction In Building Energy Systems Based On Economic Model Predictive Control." *Chemical Engineering Science* 67: 92-100.

Maguire, J., L. Earle, C. Booten, and C.E. Hancock. 2011, *Condensing Hybrid Water Heater Monitoring Field Evaluation.* Golden, CO: National Renewable Energy Laboratory.

Maguire, J. 2012. *A Parametric Analysis of Residential Water Heaters.* Boulder, CO: MS Thesis, University of Colorado Boulder.

Nehrir, M., R. Jia, and D. Pierre. 2007. "Power Management of Aggregate Electric Water Heater Loads by Voltage Control." *Power Engineering Society General Meeting.* IEEE, 1-6.

Paull, L., H. Li, and L. Chang. 2010. "A Novel Domestic Electric Water Heater Model for a Multi-Objective Demand Side Management Program." *Electric Power Systems Research* 80, no. 12: 1446-1451.

U.S. Energy Information Administration. 2012. *Annual Energy Outlook 2012.* Washington, D.C.: Office of Integrated and International Energy Analysis, U.S. Department of Energy, Report No. DOE/EIA-0383(2012).

U.S. Energy Infromation Administration. 2009. *Residential Energy Consumption Survey (RECS).* Washington, DC: U.S. Department of Energy.