

How much energy can optimal control of domestic water heating save?

M.J. Booysen ^{a,b,*}, J.A.A. Engelbrecht ^a, M. Ritchie ^a, M. Apperley ^b, A.H. Cloete ^a^aDepartment of E&E Engineering, Stellenbosch University, South Africa^bFaculty of Computing & Mathematical Sciences, University of Waikato, New Zealand

Abstract

Scheduled control of domestic electric water heaters, designed to cut energy use while minimising the impact on users' comfort and convenience, has been fairly common for some time in a number of countries. The aim is usually load-shifting (by heating water at off-peak times) and/or maximising time-of-use pricing benefits for users. The scheduling tends not to be linked to actual hot water usage and depends largely on stored thermal energy. Heat losses therefore tend to be greater than if the heater ran without a break. The effect of such a control strategy is thus to worsen the energy loss and in most cases increase greenhouse gas emissions. Many developing countries have flat-pricing (no time-of-use incentives) and rely heavily on energy from fossil fuels, making these considerations even more pressing. We explore three strategies for optimal control of domestic water heating that do not use thermostat control: matching the delivery temperature in the hot water, matching the energy delivered in the hot water, and a variation of the second strategy which provides for *Legionella* sterilisation. For each of these strategies we examine the energy used in heating, the energy delivered at the tank outlet, and issues of convenience to the user. The study differs from most previous work in that it uses real daily hot-water usage profiles, ensures like-for-like comparison in delivered energy at the point of use, and includes a daily *Legionella* avoidance strategy. We tackled this as an optimal control problem using dynamic programming. Our results demonstrate a median energy saving of between 8% and 18% for the three strategies. Even more savings would be realised if intended and unintended usage events are correctly classified, and the optimal control only plans for intended usage events.

Keywords: Electric water heater; Optimal control; Scheduled control; Temperature control; Domestic energy saving

Introduction

Scheduled control of domestic water heating has been touted as a viable way to reduce electrical energy usage. But the savings that are achievable have not been purposefully and methodically analysed and quantified to find an optimal control strategy (including schedule and temperature) that takes into account actual hot water usage patterns and personal comfort and convenience.

Electricity generators, especially those in developing countries, are struggling to meet ever-increasing demand. The problem is expected to worsen with the depletion of natural resources, the trend to move away from coal-based generation, initiatives to combat climate change (necessitating a reduction in the burning of fossil fuels), and the increasing use of electric vehicles (Serra, 2012; Monigatti et al., 2014). Developing countries rely heavily on fossil fuels for electricity generation, which directly links electrical energy usage to greenhouse gas emissions (BloombergNEF, 2018). According to the Global Energy Statistical yearbook, the global domestic electrical energy usage in 2017 was just over 22 000 TWh – a 67 % increase in usage from the year 2000 (Enerdata, 2018). In Africa, it was 663 TWh – a 75 % increase from the year 2000.

A large contributor to domestic electrical energy usage in developing countries is water heating (Hohne et al., 2019). In South Africa, for example, the residential sector uses an estimated 17 % of the total electricity generated, with domestic water heaters being responsible for a combined 38.5 GWh/day (Forlee, 1998;

*Corresponding author

Email address: mjbooyesen@sun.ac.za (M.J. Booysen)

Beute and Delpont, 2006; Eskom, 2013). Given the region’s fossil-fuel-dependent electricity generation, this translates as the release of approximately 37 700 kg/day of CO₂ (Nersa, 2018).

Water heaters are well suited for demand response to manage peak load because they can store thermal energy for prolonged periods without significant heat loss, and their water, and consequently energy, consumption patterns are cyclical (Ericson, 2009). They can thus act as a buffer to store energy in the form of heat for delayed use when peak-shifting power scheduling schemes are applied (Du and Lu, 2011; Diduch et al., 2012; Shaad et al., 2012). Scheduling schemes need to account for the device’s thermal behaviour, water draw patterns and customer comfort and convenience (Gholizadeh and Aravinthan, 2016; Roux et al., 2018). Thermal models for water heaters, and algorithms for their control, are described extensively in the literature for smart grid applications (Goh and Apt, 2004; Nehrir et al., 2007; Du and Lu, 2011; Lu et al., 2011; Diao et al., 2012; Diduch et al., 2012; Booyesen et al., 2013; Boudreaux et al., 2014; Nel et al., 2016a; Kepplinger et al., 2015; Gholizadeh and Aravinthan, 2016; Zuniga et al., 2017; Ahmed et al., 2018; Hohne et al., 2018; Jack et al., 2018; Kapsalis et al., 2018; Lunacek et al., 2018; Kepplinger et al., 2019; Gerber et al., 2019). However, only rarely are the proposed models explicitly considered as a means to reduce the overall energy used for water heating, and the resulting greenhouse gas emissions. Mostly they are considered as a means of peak load management for the benefit of the generator, or to benefit the user through time-of-use cost optimisation.

The relevant literature and major remaining challenges for energy saving through scheduled control of storage-based water heaters are listed in Table 1 and summarised in the following section. In developing countries, where the user typically pays a time-independent flat fee per kWh, and not a tariff based on time-of-use or congestion, poorer users resort to schedule control simply to reduce their monthly bill, and possibly their environmental footprint (Nel et al., 2016b; Hohne et al., 2019), and not to shift peaks or to avoid congestion charges. In this situation it is the user who bears the burden of the increased energy usage that may result from any demand management schemes (Roux et al., 2018). And given the coal-intensive energy generation typical of these countries, any change in energy used usually implies a proportional change in greenhouse gasses emitted.

Nel et al. (2018) evaluated typical strategies that people use in developing countries to save on domestic water heating, such as insulating the heater with a blanket, using less hot water, lowering the set-point temperature, and applying an on-off schedule. They simulated various hot water usage profiles (e.g. one shower per day versus two baths per day) using a one-node (lumped-mass) model for the water heater. They found that scheduled control had the biggest impact, with an expected energy saving of 9 to 18%.

In this paper, we focus on overall *energy* savings achieved through optimal control of temperature and scheduling, rather than cost savings achieved through scheduling to avoid congestion charges, or to enable load shifting.

The paper establishes the extent to which electrical energy used for water heating can be reduced for the case of *known* water draw patterns. We propose a novel optimal water heater control strategy that minimises the electrical energy used while satisfying the user’s hot water demand profile and limiting the growth of *Legionella* bacteria, and takes into account time-varying external disturbances, such as ambient temperature and cold water inlet temperature, and input constraints, such as scheduled supply-side interruption of the electricity supply. We formulate the water heater control problem as an optimal control problem and then solve it using a dynamic programming algorithm to find the optimal switching schedule for the heating element. To make the optimal control problem tractable to be solved with dynamic programming, the electric water heater (EWH) thermal dynamics is modelled with a one-node lumped-parameter model. We envisage the use of a smart water heater remote controller to manage the temperature and heating schedule using a minute-based control cycle, such as the one developed by Brown (2015) and used by Roux et al. (2018). We evaluate the performance of the proposed algorithm against the unscheduled thermostat-controlled case using one-minute simulation step time. The simulations executed to evaluate the performance also use the one-node lumped-parameter model to model the thermal dynamics, which admittedly means that the stratification in the water heater is not taken into account. We evaluate two implementations of the optimised schedule to ensure fair comparisons: one that aims to deliver water at the same temperature for each draw event, and one that aims to deliver the same energy for each draw event. We compare the strategies by using field-measured hot water usage patterns that were sampled every minute from 30 water heaters over a period of 20 days using a simulated thermal model. The water heaters were part of a research project in which users from South Africa’s Western Cape, Gauteng, and Mpumalanga volunteered to trial smart water heater controllers

Table 1: Cross-section of remaining challenges in the literature on energy saving through scheduled water heating.

	Field-measured hot water ¹	Optimal control	Temperature matched output ²	Energy matched output ³	<i>Legionella</i> -constrained control	Savings classified ⁴	Reported energy savings (%)
Fanney and Dougherty (1996)	X	X	X	✓	X	X	4-6
Goh and Apt (2004)	X	X	X	X	X	X	5-6
Booyesen et al. (2013)	X	X	X	X	X	X	14-17
Kepplinger et al. (2014)	X	✓	X	✓	X	X	5-13
Kepplinger et al. (2015)	X	✓	X	✓	X	X	10-12
Kepplinger et al. (2016)	X	✓	X	✓	✓	X	12
Booyesen and Cloete (2016)	✓	X	X	X	X	X	29-34
Gholizadeh and Aravinthan (2016)	X	X	X	X	✓	X	6
Cloete (2016)	✓	X	✓	X	X	X	6
Nel et al. (2018)	X	X	X	X	X	X	9-18
Matos et al. (2019)	✓	X	X	X	X	X	39

¹ Field-measured hot water usage used in simulation to determine savings.

² Temperature-matched hot water used in situation with purported savings.

³ Energy-matched hot water used in situation with purported savings

⁴ Electrical energy savings split into reduced thermal losses and reduced alternative losses.

(Roux et al., 2018; Booyesen et al., 2019).

We add an additional constraint to evaluate the energy impact of ensuring daily *Legionella* sterilisation in the energy-matched strategy, which operates at a lower temperature. Besides evaluating standing losses, we introduce the concept, and evaluate the impact, of the *usage losses* that result when hot water is drawn unintentionally.

Challenges in the literature

Table 1 summarises the challenges that remain in the literature and establishes the strategies to be explored in the paper.

Fanney and Dougherty (1996) evaluated the thermal efficiencies of electrical water heaters under six combinations of usage patterns and heating schedules. The metric they used was thermal efficiency, which is problematic as a stand-alone metric of savings: a water heater that is not switched on has a thermal efficiency of 100 %, and it will have different efficiencies for high volume and low volume use. The simulation results do not state the real energy savings explicitly, but they do predict savings of between 4 % and 6 %, depending on the usage patterns and heating schedules.

Goh and Apt (2004) reported savings of 5 to 8% with schedule control, and Gholizadeh and Aravinthan (2016) reported savings of 5.9 to 6.4% with a mixture of scheduled and temperature control. However, the results of both of these studies were from simulations that did not take into account the outlet temperature and useful energy used, and only with predicted (non-real) consumption patterns.

Kepplinger et al. (2014) proposed an optimisation method that uses dynamic programming to optimise scheduling for cost and/or energy usage. The method involved an hourly optimisation and implemented an hourly control of water heaters. They used synthesised usage patterns from Jordan et al. (2001), and reported energy savings ranging from 4.5 to 13.3%. Their subsequent work (Kepplinger et al., 2015) proposed an auto-scheduling mechanism with reported energy savings ranging from 10.5 to 12.4%, based on the simulation of a stratified thermal model, similar to the work presented here in our paper. In the case of cold events, however, the state constraint approach could not be satisfied. Moreover, the approach ensured equal delivery of matched energy, but did not consider the requirement of matched temperatures, so that the required temperature at the start of each usage event matched the temperature achieved by thermostat control at the start of the same event. They later extended these results (Kepplinger et al., 2016), and also ran a field trial, for which 12.3% savings were achieved.

By implementing scheduled water heating control which was hot-water-usage driven, Booyesen et al. (2013) estimated a reduction in energy used for water heating of between 14 % and 17 %. Although they did a small experiment to confirm the results, they did not take into account the effects of the possibly reduced temperature (i.e. reduced energy) of the water drawn from the tank. Moreover, their estimates were based on

a simple lumped-mass analytical physics model (i.e. not step-wise simulated) with an assumed usage pattern for only one water heater.

Subsequently, Nel et al. (2016a) presented a more accurate model with the explicit purpose of modelling energy used for water heating under schedule control. Their model makes provision for horizontally mounted water heaters (the most common orientation found in South Africa) under thermostat-based and scheduled-based heating control. Booysen and Cloete (2016) presented preliminary work in which the same model was used to compare savings, achieved by scheduling, in a controlled field trial of four water heaters. The results were augmented by a controlled laboratory experiment with one heater. They reported savings of 29 % due to scheduling. However, the study failed to determine the savings in the scenario where the output energy (and temperature) were matched to the baseline condition. This probably invalidated the conclusion, as the utility extracted from the water heater was not comparable before and after the intervention that resulted in the savings. More importantly, the metric used to measure savings, namely the change in energy per litre, did not represent the savings fairly, since it compared the electrical energy (kWh) used per litre of hot water delivered under the two conditions, but failed to take into account the temperature of (and therefore the utility extracted from) that volume of hot water. Cloete (2016) subsequently conducted a controlled study in which the experiment was repeated in a laboratory with longer, fixed heating times, and obtained just 16 % savings. He then used an iterative adjustment of the set-point to match the output temperature retrospectively for a fairer comparison, which reduced the predicted savings to only 6 %.

In recent studies, Matos et al. (2017, 2019) evaluated the impact of flow reducers and reduced temperature on energy used and the resulting CO₂ emissions. The focus of the study was more on flow reducers than on scheduling, and limited specifically to water used for baths. Their results indicated a reduction of 39 % in energy when the temperature was decreased from 75 °C to 60 °C. They described in detail the observed differences, but did not analyse in detail the energy changes from the perspective of the heater supplying the water. Further, they did not differentiate between savings achieved from standing losses and reduced unintentional hot water usage. Finally, they did not evaluate the impact of scheduled control.

In addition to these remaining challenges only one of the studies listed in Table 1, Kepplinger et al. (2016), implemented weekly temperature control standards to limit *Legionella* growth, and none of the studies evaluated losses other than thermal losses. As far as we could determine, none of the studies used high-frequency (minutes, rather than hours) sampled water usage data in their simulations.

Heating control strategies

In our study we took thermostat control as the baseline heater control strategy and evaluated three alternative strategies, the third being a variant of the second, to determine which gives the best electrical energy savings.

0. (Baseline) Thermostat control (TC): This is the mode in which water heaters are designed to be used and how most people use them. The thermostat strives to maintain the water at a target temperature, normally set between 65 °C and 75 °C, with a small hysteresis band around the set temperature. This heating strategy is wasteful, as it maintains a high temperature between draw events that may be far apart. Of the three strategies, this one loses the most energy to the environment. For this heating strategy and set temperature, water is drawn from the water heater at a temperature that is higher than people require for most uses. To achieve a convenient temperature, the user normally has to regulate the temperature by mixing with cold water to achieve a nominal temperature of approximately 40 °C (Armstrong et al., 2014; Jacobs et al., 2018; Kepplinger et al., 2019).

1. Scheduled control with temperature matching (TM): In an attempt to save on bills and reduce loss of energy, financially sensitive and environmentally aware users resort to turning off their heaters for extensive periods between the times when they need hot water (Nel et al., 2016b). The timing applied in this strategy is individually motivated and could vary significantly between users, but the optimal scenario (for a water heater with energy storage) is to switch the heater on just sufficiently in advance of the time the hot water will be needed (Booyesen et al., 2013; Nel et al., 2016b). When applying optimisation techniques, as we have done in this study, optimisation constraints can be set to ensure that the same volume of water is drawn at the same temperatures as under thermostat control, but the thermal losses to the environment are minimised. This approach assumes that the user requires water at these high temperatures, intending

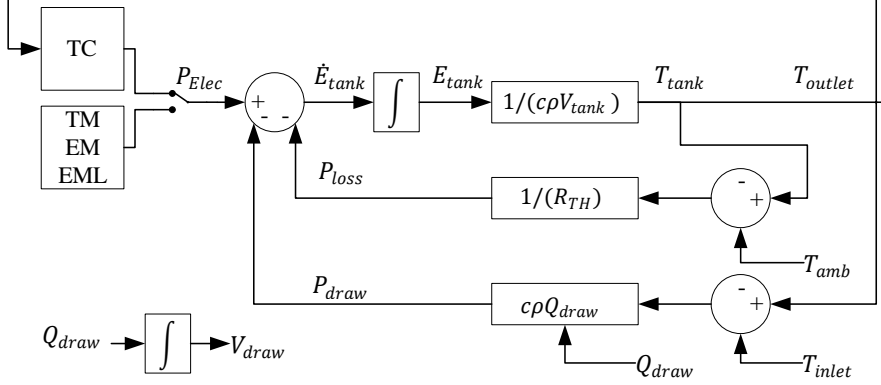


Figure 1: Closed-loop energy flow in an EWH with feedback via the thermostat.

to mix the water with cold water to achieve a desired temperature. For this control strategy, the heater will deliver the same amount of useful output energy as in the baseline thermostat control scenario.

2. Scheduled control with energy matching (EM): An alternative approach to the optimised schedule control with temperature matching is to assume that the user does not require the water at the high temperatures, but rather is satisfied with a lower, more directly usable temperature. In this case, there is less need, or no need, to add cold water, as the water is already at or just above the desired temperature for use. A lower target temperature during water draw-offs, of say 38 °C, could be used. For our study we increased the volume drawn from the heater, to ensure that the same amount of energy was delivered in the water drawn as under thermostat control (Armstrong et al., 2014; Jacobs et al., 2018; Roux et al., 2018).

3. Scheduled control with energy matching plus *Legionella* prevention (EML): Although optimising the heating schedule has energy benefits, there are health risks when the water is maintained and delivered at low temperatures. *Legionella pneumophila* thrives at temperatures between 32 °C and 42 °C and has been found in water heaters (Armstrong et al., 2014; Stone et al., 2019). According to Stout et al. (1986), to sterilise the bacteria, the heater must spend 11 min at 60 °C or 3 min at 70 °C, at least once a day.

We therefore define a *Legionella*-driven control strategy that ensures that the heater temperature reaches 60 °C for at least 11 min at least once before the largest water usage event of the day.

Electric water heater (EWH) dynamics

The EWH acts as a closed-loop system with a thermostat providing the feedback. Figure 1 shows an overview of the energy flow in this system and the function of the feedback loop. The temperature of the water inside the tank T_{tank} is fed in as input to the thermostat. The thermostat is directly connected to the heating element, which provides electrical power P_{elec} to increase the thermal energy inside the tank, E_{tank} . This energy depends on the rated power of the heating element, P_{rated} . When water is drawn from the tank at a higher temperature than the inlet temperature, the net effect is a reduction in thermal energy at rate of P_{draw} . When heat is lost to the environment because of the temperature difference between the tank and the environment, the net effect is a reduction in thermal energy at a rate of, P_{loss} .

We model the EWH thermal dynamics with the following one-node lumped-parameter model.

$$\dot{E}_{\text{tank}}(t) = P_{\text{elec}}(t) - P_{\text{draw}}(t) - P_{\text{loss}}(t) \quad (1)$$

where E_{tank} is the thermal energy of the water inside the EWH, P_{elec} is the power input delivered by the heating element, $P_{\text{draw}}(t)$ is the power output due to hot water draw (hot water leaving the EWH and being

replaced by cold water), and P_{loss} is the power output due to thermal losses to the environment. The heating element is either off or on, which means that the power input P_{elec} delivered by the element is either zero or its rated power P_{rated}

$$P_{\text{elec}}(t) \in \{0, P_{\text{rated}}\} \quad (2)$$

The electrical energy supplied by P_{rated} is given by

$$E_{\text{elec}} = \int_{t_0}^{t_f} P_{\text{elec}}(t) dt \quad (3)$$

The power output due to hot water draw P_{draw} is given by

$$P_{\text{draw}}(t) = \rho Q_{\text{draw}}(t) [\hat{h}_{\text{outlet}}(t) - \hat{h}_{\text{inlet}}(t)] \quad (4)$$

where ρ is the density of the water, $Q_{\text{draw}}(t)$ is the hot water outlet volumetric flow rate, and \hat{h}_{inlet} and \hat{h}_{outlet} are the specific enthalpy entering and leaving the water heater, respectively. Under conditions of constant pressure and constant specific heat capacity, this can be approximated by

$$P_{\text{draw}}(t) \approx c_P \rho Q_{\text{draw}}(t) [T_{\text{outlet}}(t) - T_{\text{inlet}}(t)] \quad (5)$$

where c_P is the constant pressure-specific heat capacity of the water, T_{outlet} is the hot water outlet temperature, and T_{inlet} is the cold water inlet temperature.

The power output due to thermal losses $P_{\text{loss}}(t)$ is given by

$$P_{\text{loss}}(t) = \frac{1}{R_{TH}} [T_{\text{tank}}(t) - T_{\text{amb}}(t)] \quad (6)$$

where R_{TH} is the thermal resistance of the EWH wall, T_{tank} is the water temperature inside the EWH, and T_{amb} is the ambient temperature. Note that for the one-node lumped-parameter model, the tank temperature and the hot water outlet temperature are assumed to be equal.

The relationship between the EWH energy E_{tank} and the EWH water temperature T_{tank} relative to a reference temperature where we define energy to be zero, is given by

$$T_{\text{tank}}(t) = \frac{E_{\text{tank}}(t)}{c_V \rho V_{\text{tank}}} \approx \frac{E_{\text{tank}}(t)}{c_P \rho V_{\text{tank}}} \quad (7)$$

where V_{tank} is the volume of the EWH, c_V is the constant volume specific heat capacity, which is approximately equal to c_P for water, and henceforth denoted as c .

The hot water outlet flow rate $Q_{\text{draw}}(t)$ is the superposition of the flow rate due to intentional usage $Q_{\text{usage}}(t)$ and the flow rate due to unintentional usage $Q_{\text{unintentional}}(t)$, defined below.

$$Q_{\text{draw}}(t) = Q_{\text{usage}}(t) + Q_{\text{unintentional}}(t) \quad (8)$$

The power output due to hot water draw $P_{\text{draw}}(t)$ is therefore also the superposition of the power output due

to intentional usage $P_{\text{usage}}(t)$ and the power output due to unintentional usage $P_{\text{unintentional}}(t)$

$$P_{\text{draw}}(t) = P_{\text{usage}}(t) + P_{\text{unintentional}}(t) \quad (9)$$

with

$$P_{\text{usage}}(t) = c\rho Q_{\text{usage}}(t)[T_{\text{outlet}}(t) - T_{\text{inlet}}(t)] \quad (10)$$

$$P_{\text{unintentional}}(t) = c\rho Q_{\text{unintentional}}(t)[T_{\text{outlet}}(t) - T_{\text{inlet}}(t)] \quad (11)$$

Similarly, the energy output due to hot water draw $E_{\text{draw}}(t)$ is the superposition of the energy output due to intentional usage $E_{\text{usage}}(t)$ and the energy output due to unintentional usage $E_{\text{unintentional}}(t)$

$$E_{\text{draw}}(t) = E_{\text{usage}}(t) + E_{\text{unintentional}}(t) \quad (12)$$

The distinction between usage and unintentional usage is important, because we give the usage profile to the optimal control algorithm as an objective to satisfy, while we treat the unintentional usage as a disturbance and an energy loss.

The optimal control problem formulation

In this section we briefly explain optimal control theory, using Kirk (2012) as our primary source, and then formulate the EWH control as an optimal control problem.

Optimal control theory

Optimal control theory is concerned with finding a control law for a given system so that a specified optimality criterion is achieved. This criterion is typically specified as a cost function to be minimised. The problem is therefore to find a control law that produces the optimal control input signals and resulting optimal state trajectories that together minimise a given cost function, subject to a set of dynamic constraints, terminal state constraints, and state variable and control input inequality constraints. This can be expressed mathematically as follows

$$\mathbf{u}^*(t) = \arg \min_{\mathbf{u}(t)} J(\mathbf{x}(t), \mathbf{u}(t), t) \quad (13)$$

$$= \arg \min_{\mathbf{u}(t)} \left[h(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} g(\mathbf{x}(t), \mathbf{u}(t), t) dt \right] \quad (14)$$

subject to the dynamic constraint

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) \quad (15)$$

the terminal state constraint

$$\mathbf{x}(t_f) \in \mathcal{X}_f \quad (16)$$

and the state space and input space constraints

$$\mathbf{x}(t) \in \mathcal{X} \quad (17)$$

$$\mathbf{u}(t) \in \mathcal{U} \quad (18)$$

In these equations, J is the cost function, $\mathbf{x}(t)$ is a candidate state trajectory, $\mathbf{u}(t)$ is a candidate control input signal, t is time, $\mathbf{f}(\cdot)$ represents the set of nonlinear, time-variant differential equations of the system, $g(\cdot)$ is the state transition cost function, $h(\cdot)$ is the terminal state cost function, \mathcal{X}_f is the set of admissible final states, and \mathcal{X} and \mathcal{U} are the sets of admissible states and admissible inputs. $\mathbf{x}^*(t)$ and $\mathbf{u}^*(t)$ are the optimal state trajectory and optimal control input signal that minimise the cost function subject to the constraints.

Formulation of EWH thermal control as an optimal control problem

Optimal control problem: Given a hot water usage profile in terms of flow rate $Q_{\text{usage}}(t)$ and desired minimum hot water temperature $T_{\text{usage}}(t)$ as a function of time t , the cold water inlet temperature $T_{\text{in}}(t)$ and the ambient temperature $T_{\text{amb}}(t)$ as time-varying disturbance signals, and scheduled supply-side interruption of the electricity supply to the EWH $P_{\text{max}}(t)$, we wish to determine the optimal EWH control signal $P_{\text{elec}}^*(t)$ that will satisfy the hot water usage profile, while minimising the total energy used.

System dynamics: The system dynamics are defined as the nonlinear differential equations describing the EWH thermal dynamics as described above, and specifically by equations (1), (2), (5), and (6).

System state: The state variable for the system is the thermal energy E of the water inside the EWH

$$x(t) = E(t) \quad (19)$$

Control input: The control input for the system is the power input P_{elec} delivered by the heating element

$$u(t) = P_{\text{elec}}(t) \quad (20)$$

State constraints: The physical limitations on the thermal energy inside the EWH are specified by defining the following set of admissible states

$$E(t) \in [E_{\text{min}}, E_{\text{max}}] \quad (21)$$

where the minimum energy E_{min} and the maximum energy E_{max} correspond to the minimum temperature T_{min} and the maximum temperature T_{max} specified for the EWH

$$E_{\text{min}} = c\rho V_{\text{tank}} T_{\text{min}} \quad (22)$$

$$E_{\text{max}} = c\rho V_{\text{tank}} T_{\text{max}} \quad (23)$$

Input constraints: The control input constraints are specified by defining the following set of admissible inputs

$$P_{\text{elec}}(t) \in \{0, P_{\text{rated}}\} \quad (24)$$

The heating element is either off or on, which means that the power input P_{elec} delivered by the element is either zero or its rated power P_{rated} .

Terminal state constraints: No special terminal state constraints are imposed, and the terminal state constraints are the same as the state constraints.

$$E(t_f) \in [E_{\text{min}}, E_{\text{max}}] \quad (25)$$

Cost function: The objective to minimise the total energy used is represented by the following cost function

$$J = \int_{t_i}^{t_f} P_{\text{elec}}(t) dt \quad (26)$$

where t_i is the initial time, and t_f is the final time. Note that this a specific implementation of the general cost function

$$J = \left[h(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} g(\mathbf{x}(t), \mathbf{u}(t), t) dt \right] \quad (27)$$

with the state transition cost defined as

$$g(\mathbf{x}(t), \mathbf{u}(t), t) = P_{\text{elec}}(t) \quad (28)$$

and the termination cost defined as

$$h(\mathbf{x}_f(t)) = 0 \quad (29)$$

Temperature profile constraints (usage and *Legionella* prevention): The objective of satisfying the hot water usage profile is represented by the following time-varying state inequality constraint on the thermal energy of the water inside the EWH

$$E(t) \geq c\rho V_{\text{tank}} T_{\text{profile}}(t) \quad (30)$$

The need to increase the temperature once per day to prevent the growth of *Legionella* bacteria can also be included in this time-varying state inequality constraint. The profile temperature $T_{\text{profile}}(t)$ is set to the desired usage temperature T_{usage} for times that correspond to intentional draw, to the minimum *Legionella* prevention temperature $T_{\text{Legionella}}$ for a set duration once per day, and to the minimum EWH temperature T_{min} for all other times.

$$T_{\text{profile}}(t) = \left\{ \begin{array}{ll} T_{\text{usage}} & \text{when intentional draw causes } Q_{\text{draw}}(t) > 0 \\ T_{\text{Legionella}} & \text{once per day to prevent } \textit{Legionella} \text{ growth} \\ T_{\text{min}} & \text{otherwise} \end{array} \right\} \quad (31)$$

Electricity supply constraints The constraint represented by the scheduled supply-side interruption of the electricity supply to the EWH is represented by the following time-varying input constraint

$$P_{\text{elec}}(t) \leq P_{\text{max}}(t) \quad (32)$$

The maximum power input is set to P_{rated} when the supply-side electricity is available, and to zero when the supply-side electricity is interrupted, with

$$P_{\text{max}}(t) = \left\{ \begin{array}{ll} 0 & \text{when the electricity is interrupted} \\ P_{\text{rated}} & \text{when the electricity is available} \end{array} \right\} \quad (33)$$

Constructing the temperature profile constraint: We construct the temperature profile constraint $T_{\text{profile}}(t)$ differently for temperature matching, energy matching, and energy matching with *Legionella* prevention. In constructing the constraint we take into account “unreasonable” hot water usage profiles where it is impossible to deliver hot water at the minimum desired temperature, even if the heating element is switched on permanently. For example, drawing all of the hot water from the tank and then expecting more hot water shortly thereafter is considered “unreasonable” hot water usage behaviour.

Optimal temperature matching (TM): We construct the temperature profile constraint for temperature matching so that the required EWH temperature at the start of each usage event matches the EWH temperature achieved by thermostat control at the start of the same event.

Optimal energy matching (EM): We construct the temperature profile constraint for energy matching so that the EWH temperature remains above 40 °C for the duration of each usage event. However, we increase the outlet flow rate so that the energy in the volume of hot water delivered at 40 °C matches the energy in the volume of hot water delivered at the thermostat control temperature.

Optimal energy matching with Legionella prevention (EML): We construct the temperature profile constraint for energy matching with *Legionella* prevention in exactly the same way as for energy matching (EM), except that once per day we increase the EWH temperature to 60 °C for 11 min to prevent the growth of *Legionella*. We schedule this to be just before the largest usage event for the day.

“Unreasonable” hot water usage profiles: To accommodate these profiles, we run a forward simulation for the entire usage pattern, assuming that the heating element is always switched on, to determine the best temperatures the EWH can deliver for a given usage profile. We then modify the temperature profile constraint $T_{\text{profile}}(t)$ to be the minimum of the temperature profile constructed to satisfy the hot water usage and the *Legionella* prevention and the achievable temperatures that were determined from the forward simulation.

The dynamic programming solution

Dynamic programming is a search algorithm that models an optimal control problem as a multi-stage decision process and uses the principle of optimality to find the optimal state trajectories and control inputs to minimise the cost function from all initial states. It uses the principle of optimality to drastically reduce the number of calculations required to determine the optimal control law. For systems with a large number of dimensions, the memory requirements of the dynamic programming algorithm becomes prohibitive. This is called “the curse of dimensionality”. Fortunately, the EWH system can be modelled as a first-order system, which makes dynamic programming a very suitable approach to solving the optimal control problem.

Discretisation

To apply dynamic programming to an optimal control problem, the problem has to be discretised in time to represent different decision stages, and discretised in state to represent a finite number of decisions to be made at each decision stage. The dynamic programming algorithm starts at the terminal stage and works backward in time through intermediate stages until it finds the optimal admissible path from the initial state to a terminal state.

Discrete-time dynamic model: The continuous-time differential equations describing the system dynamics are discretised to produce discrete-time difference equations that describe the state transition from one discrete time instant to the next.

$$E(k+1) = E(k) + [P_{\text{elec}}(k) - P_{\text{usage}}(k) - P_{\text{losses}}(k)]\Delta t \quad (34)$$

where Δt is the sampling period of the discrete time step.

Quantised state vector array: The continuous set of admissible states \mathcal{X} is discretised to create a finite set of state values for the dynamic programming algorithm. We therefore create an array \mathbf{X}_q of quantised state vector values $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n$ that span the set of admissible states \mathcal{X} .

$$\mathbf{X}_q = \{E_1, E_2, \dots, E_i, \dots, E_n\} \quad (35)$$

Control inputs The continuous set of admissible inputs \mathcal{U} must be discretised to create a finite set of input values. We therefore create an input array \mathbf{U}_q of quantised input vector values $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_n$ that span the set of admissible inputs \mathcal{U} .

$$\mathbf{U}_q = \{0, P_{\text{rated}}\} \quad (36)$$

Incremental state transition cost function: The continuous-time cost function is discretised by expressing the total path cost J_{ij} from the current state $\mathbf{x}_i(k)$ via the next state $\mathbf{x}_j(k+1)$ as the sum of the incremental cost ΔJ_{ij} of transitioning from the current state to the next state and the total path cost of the next state J_j to a final state.

$$J_{ij}(\mathbf{x}_i, \mathbf{x}_j) = \Delta J_{ij}(\mathbf{x}_i, \mathbf{x}_j) + J_j(\mathbf{x}_j) \quad (37)$$

The incremental state transition cost ΔJ_{ij} is obtained by discretising the state transition cost function $g(\mathbf{x}(t), \mathbf{u}(t), t)$ as follows

$$\Delta J_{ij}(\mathbf{x}_i(k), \mathbf{x}_j(k+1)) \approx g(\mathbf{x}_i(k), \mathbf{u}_{ij}(k), k)\Delta t \quad (38)$$

where \mathbf{u}_{ij} is the admissible control input that transitions the system from state \mathbf{x}_i at time k to state \mathbf{x}_j at time $k+1$.

The incremental transition cost ΔJ is therefore obtained as

$$\Delta J_{ij}(\mathbf{x}_i(k), \mathbf{x}_j(k+1)) = P_{\text{elec}}(k)\Delta t \quad (39)$$

Algorithm

Initialisation

1. Create tables to store the quantised states \mathbf{X}_q , and the optimal path cost \mathbf{J}^* from a given state to a final state, the optimal next state index \mathbf{j}^* , and the optimal control input \mathbf{U}^* , as follows:

$$\mathbf{X}_q = \{E_{ik}\} \quad (40)$$

$$\mathbf{J}^* = \{J^*_{ik}\} \quad (41)$$

$$\mathbf{j}^* = \{j^*_{ik}\} \quad (42)$$

$$\mathbf{U}^* = \{P^*_{elec\ ik}\} \quad (43)$$

where i is the quantisation index of the current state j is the quantisation index of the next state, and k is the index of the current time instant.

2. Populate the final column of the optimal cost table with zeroes to assign termination costs of zero to all final states.

$$J^*_{iN} \leftarrow 0 \quad \forall i \quad (44)$$

3. Populate the rest of the optimal cost table with infinity values so that if a finite cost is calculated for a state, it will be lower than infinity and will replace the initial cost as the new lowest path cost.

$$J^*_{ik} \leftarrow \infty \quad \forall i \quad \forall k \in [1, N-1] \quad (45)$$

4. Populate the final column of the optimal next state table so that the optimal next state (index j^*) for all final states is the state itself (index i).

$$j^*_{iN} \leftarrow i \quad \forall i \quad (46)$$

Execution

1. Set the index j of the next state $\mathbf{x}_j(k)$ to the first index of the state vector array

$$j \leftarrow 1 \quad (47)$$

2. For the given index j of the next state $\mathbf{x}_j(k+1)$, calculate the previous state $\mathbf{x}(k)$ for all admissible values of the control input $\mathbf{u}(k)$.

For each

$$P_{elec}(k) \in \{0, P_{rated}\} \quad (48)$$

$$\text{and } P_{elec}(k) \leq P_{max}(k) \quad (49)$$

calculate

$$E(k) = E(k+1) - [P_{elec}(k) - P_{usage}(k) - P_{losses}(k)]\Delta t \quad (50)$$

with

$$P_{usage}(k) = c\rho Q_{usage}(k)[T(k) - T_{in}(k)] \quad (51)$$

$$P_{losses}(k) = \frac{1}{R_{TH}}[T(k) - T_{ambient}(k)] \quad (52)$$

where

$$T(k) = \frac{E(k)}{c\rho V_{tank}} \quad (53)$$

and the inlet water temperature $T_{in}(k)$ and the ambient temperature $T_{ambient}(k)$ are known external disturbances at time instant k .

3. Check whether the previous state $\mathbf{x}(k)$ is an admissible state.

$$E(k) \geq c\rho V_{\text{tank}} T_{\text{profile}}(k) \quad (54)$$

If the previous state is an admissible state, continue to step 4. Else, continue iterating through the possible values of the control input.

4. Find the index i of the quantised previous state $\mathbf{x}_i(k)$ in the state vector array that is closest to the calculated previous state $\mathbf{x}(k)$.
5. Calculate the cost from the current state $\mathbf{x}_i(k)$ to a final state through this next state $\mathbf{x}_j(k+1)$

$$J_i \leftarrow \Delta J_{ij} + J_{*j} \quad (55)$$

where

$$\Delta J_{ij} = P_{\text{elec}}(k)\Delta t \quad (56)$$

6. If the new cost is lower than the lowest cost so far, it becomes the new lowest cost. Also, the control input $\mathbf{u}_i(k)$ and the next state index j become the new best control input $\mathbf{u}_i^*(k)$ and the new best next state j_i^* for the current state $\mathbf{x}_i(k)$. In other words, if $J_i < J_i^*$ then

$$\begin{aligned} J_i^* &\leftarrow J_i \\ \mathbf{u}_i^* &\leftarrow \mathbf{u}_i(k) \\ j_i^* &\leftarrow j \end{aligned}$$

7. If the next state index j was not the last index, increment the next state index j and return to step 2.

$$j \leftarrow j + 1 \quad (57)$$

Else, if the next state index j was the last index, and the time instant was not the first time instant, then step one time instant k backward and return to step 1.

$$k \leftarrow k - 1 \quad (58)$$

Else, terminate the execution.

Lookup table navigation

The dynamic programming algorithm produces a lookup table of the optimal state trajectories and optimal control sequences for EWH control from all time instants and admissible initial states. Given a time instant k and an initial state $\mathbf{x}_i(k)$, the optimal state trajectory $\{\mathbf{x}^*(k) : k = k, k+1, \dots, N\}$ and the optimal control sequence $\{\mathbf{u}^*(k) : k = k, k+1, \dots, N\}$ can be obtained by starting at the column index k and the row index i of the initial state in the lookup table and iteratively navigating through the lookup table by following the optimal next state indexes j^* .

Simulation setup

Simulation parameters

The hot water usage data, the software implementation in Jupyter Notebook, and the simulation output are available at <http://bit.ly/EWHSavingsESDDataset>. Table 2 lists the dataset properties, parameters and constants used in the optimisation and simulation.

Metrics

Water draw is aggregated into what we term usage *events*. An event starts when a tap is opened and stops when it is closed or, put differently, an event starts when a non-zero volume is sampled one sample after a zero volume sample, and stops when the flow returns to zero after the non-zero sample. The event definition quantifies the water usage into bundles, providing a more convenient way of referring to sections of water

Table 2: Parameters used for simulations and optimisation.

Symbol	Description	Value	Unit
EWH model parameters			
R_{TH}	Thermal resistance of EWH	0.4807	$\frac{K \cdot day}{kWh}$
c	Specific heat capacity of water	4184	$\frac{J}{kg \cdot K}$
ρ	Water density	1000	$\frac{kg}{m^3}$
T_{set}	Target temperature	68.5	$^{\circ}C$
T_{hyst}	Hysteresis (deadband)	± 1.5	$^{\circ}C$
T_{amb}	Ambient temperature	20	$^{\circ}C$
T_{inlet}	Inlet temperature of EWH	20	$^{\circ}C$
V_{tank}	Tank volume of EWH	150	L
P_{rated}	Power rating of element	3	kW
Optimisation parameters			
$T_{tank(max)}$	Maximum temperature of EWH	70	$^{\circ}C$
$T_{tank(min)}$	Minimum temperature of EWH	20	$^{\circ}C$
$T_{tank(usage)}$	Minimum target usage temperature	40	$^{\circ}C$
T_{start}	Initial boundary condition of EWH	68.5	$^{\circ}C$
Water draw dataset			
D	Duration	20	days
Δt	Sampling period	1	min
Resolution		0.5	L
Number of water heaters		30	
Average number of events per EWH per day		7.5	

Table 3: Metrics used for performance assessment.

Metric	Description	Unit
E_{elec}	Daily average electrical energy used per day for water heating as distribution of heaters.	kWh
E_{draw}	Daily average thermal energy in hot water drawn from tank (per day). Indicates effective energy used, but includes energy lost due to unintentional use.	kWh
E_{loss}	Daily average energy lost to environment through tank (per day).	kWh
$E_{unintentional}$	Thermal energy lost due to unintentional use (per day).	kWh
$\Delta E(kwh)$	Reduction in electrical energy per day compared to thermocouple control.	kWh
$\Delta E(\%)$	Reduction in electrical energy per day compared to thermocouple control.	%
\bar{T}_{usage}	Average event temperature (excludes unintentional use)	$^{\circ}C$
Cold events	Number of events with any sample $T < 40^{\circ}C$.	

usage patterns, and providing a metric for counting the number of times a user experiences an undesired temperature.

Given the plumbing between the water heater outlet and the point of use (i.e. tap), events with a volume of less than 2 L are unlikely to result in the hot water reaching the point of use. This can be shown to be true with the normal pipe diameter of 22 mm and a conservative pipe length of 5 m, which results in a volume of 1.9 L hot water drawn into the piping before the point of use. These events that draw hot water that does not reach the point-of-use are therefore considered unintended events, for which our algorithm does not instruct the element to heat, and which are excluded when counting the number of cold events. These events are likely to occur when the user uses a mixer tap with a position between hot and cold, or unwittingly opens the hot tap for a quick cold draw.

Table 3 lists the metrics we use to assess the performance of the various strategies. The electrical energy, effective energy in the drawn water, energy lost to the environment, and energy lost due to unintentional usage comprise the energy metrics. We use the change in electrical energy as a percentage reduction and as an absolute change to determine the achieved savings. We also report the average event temperatures, with the number of cold events, to establish user comfort and convenience.

The distributions of the electrical energy used per day, thermal energy drawn per day, average outlet temperature during usage events, and thermal energy losses per day over all EWHs are shown in Figure 3

(a) to (d). The average electrical energy used per day by an individual EWH is calculated using

$$\bar{P}_{\text{elec}|h} = \frac{\sum_{k=1}^{N_h} P_{\text{elec}|h}(k)\Delta t}{D} \text{ kWh/day}$$

where $\bar{P}_{\text{elec}|h}$ is the average electrical energy used by heater h per day, $P_{\text{elec}|h}(k)$ is the instantaneous electrical power used by heater h at sampling instant k , Δt is the sampling period, N_h is the total number of samples, and D is the total number of days in the data set. The average thermal energy drawn per day $\bar{P}_{\text{draw}|h}$ and the average thermal energy losses per day $\bar{P}_{\text{loss}|h}$ for an individual EWH are calculated using

$$\bar{P}_{\text{draw}|h} = \frac{\sum_{k=1}^{N_h} P_{\text{draw}|h}(k)\Delta t}{D} \text{ kWh/day}$$

$$\bar{P}_{\text{loss}|h} = \frac{\sum_{k=1}^{N_h} P_{\text{loss}|h}(k)\Delta t}{D} \text{ kWh/day}$$

The distributions of the electrical energy savings per day for each strategy (TM, EM, and EML), expressed both as a reduction in kWh per day and as a percentage reduction, are shown in Figure 3 (e) and (f). We calculate the distribution of the energy savings for a particular strategy by first calculating the individual savings for each EWH, and then plotting the distribution of the individual savings over all EWHs. For example, we calculated the energy savings per day for the EM strategy compared to the baseline TC strategy for each individual EWH using the following formulas, and then plotted the distribution of savings for the EM strategy over all EWHs:

$$\Delta\bar{P}_{\text{elec}|h,\text{EM}}(\text{kWh/d}) = \bar{P}_{\text{elec}|h,\text{TC}} - \bar{P}_{\text{elec}|h,\text{EM}} \text{ kWh/day}$$

$$\Delta\bar{P}_{\text{elec}|h,\text{EM}}(\%) = \frac{\bar{P}_{\text{elec}|h,\text{TC}} - \bar{P}_{\text{elec}|h,\text{EM}}}{\bar{P}_{\text{elec}|h,\text{TC}}} \times 100\%$$


Results and discussion

Simulation results for a single EWH


Figure 2 shows example simulation results for a single EWH, comparing thermostat control (TC), optimal temperature matching (TM), optimal energy matching (EM), and optimal energy matching with Legionella prevention (EML). For each control technique, the EWH temperature, outlet flow rate, and heating element control signal are plotted as a function of time for a period of two days (48 hours). The results are shown for the same hot water draw profile, with all simulations starting from the same initial EWH temperature of 70 °C.

Thermostat control (TC): The thermostat control keeps the temperature at about a set point of 68.5 °C (allowing for 1.5 °C hysteresis). When the temperature drops below 67 °C, the heating element switches on; when the temperature rises above 70 °C, it switches off. During usage events, when the outlet flow rate is non-zero, the temperature drops significantly. We observed a single cold event at time $t = 9$ hours, when the temperature dropped below 40 °C. This cold event was caused by an unusually large usage event, and could not be prevented even with the temperature at maximum at the start of the event, and with the heating element switched on for the full duration of the event. The electrical power consumption for thermostat control was 23.85 kWh over the two days.


Optimal temperature matching (TM): The temperature matching control ensures that the temperature matches the corresponding temperature for thermostat control, but only during usage events. Between usage events, the TM control allows the temperature to drop, and the EWH only starts heating again just before the next usage event. The temperature between the usage events is therefore lower for TM than for TC. We observed the same cold event at time $t = 9$ hours, due to the unusually large usage event. Since TC was not able to prevent the cold event, starting at maximum temperature and with the heating element switched fully on, we did not expect that TM would be able to prevent the cold event either. The electrical power consumption for optimal temperature matching was 16.7 kWh over the two days.




(a) Thermostat control



(b) Temperature matching



(c) Energy matching



(d) Energy matching with *Legionella*

Figure 2: Simulation results for thermostat control, optimal temperature matching, optimal energy matching, and optimal energy matching with *Legionella* prevention. The plots show the EWH temperature, the outlet flow rate, and the element state. The EWH temperature for thermostat control is repeated on all plots for comparison. The cold event threshold (40 °C) is indicate with a dashed blue line. The *Legionella* threshold (60 °C) is indicated with a dashed red line.

Table 4: Energy, temperature, volume, and cold event results.

		TC	TM	EM	EML
\bar{V}_{hot}	(L/day)	142	142	229	214
E_{elec}	(kWh/day)	8.2	7.9	6.8	7.2
E_{draw}	(kWh/day)	5.7	5.7	5.9	6.0
$E_{unintentional}$	(kWh/day)	0.06	0.05	0.02	0.02
E_{loss}	(kWh/day)	2.4	2.0	1.0	1.3
\bar{T}_{usage}	(°C)	65	65	42	44
ΔE_{elec} (kWh)	(kWh/day)	–	0.6	1.5	1.0
ΔE_{elec} (%)	%	–	3.6, 7.9 , 11.2	12.8, 17.8 , 26.3	9.1, 13.1 , 17.0
Average cold events per day*		0.06	0.06	0.06	0.06

Note: The distributions are reported as 25th percentile, **median**, 75th percentile

*All these cold events were generated by six heavy users out of the 30, which means the median and 75th percentile EWHs had zero cold events. The number of cold events did not increase from that of TC.

Optimal energy matching (EM): The energy matching control ensures that the temperature remains above 40 °C during usage events. We increased the outlet flow rate to represent the fact that all the water that reached the user was drawn from the EWH, and that the hot water from the EWH was not mixed with cold water before reaching the user. Between usage events, the EM control allows the temperature to drop, and only starts heating again just before the next usage event. The temperature both during and between the usage events is therefore lower for EM than for TC and TM. The same cold event is observed at time $t = 9$ hours, due to the overlarge usage event. The electrical power consumption for optimal energy matching was 16.55 kWh over the two days.

Optimal energy matching with Legionella prevention (EML): The results for the energy matching control with *Legionella* prevention look almost exactly the same as the results for energy matching control without this prevention, except that we increased the EWH temperature to the *Legionella* prevention threshold of 60 °C at time $t = 41$ hours and then held it there for 11 min. This increase in temperature was scheduled to be just before the largest usage event for the day, excluding cold events. We observed the same cold event at time $t = 9$ hours, due to the unusually large usage event. The temperature both during and between usage events was therefore almost the same for EM and EML, except that the latter had a higher temperature during the scheduled *Legionella* prevention heating. We therefore expected that EML would have a slightly higher power consumption. The electrical power consumption for optimal energy matching with *Legionella* prevention was 16.05 kWh over the two days.

The large difference of 7.15 kWh between the power consumption for TC (23.85 kWh) and the power consumption for TM (16.7 kWh) may seem surprising, given that the standing loss for an EWH is typically only 2 kWh per day. However, it should be noted that TC immediately replenishes the usage energy of the last usage event of the day, while TM, EM and EML wait until just before the next usage event, which occurs only the next day.

Distribution of results over all EWHs

The effects of the different control strategies on energy and temperature are summarised in Table 4 and shown in Figure 3. Figure 3 (a) shows that all the other strategies used less electrical energy than the baseline TC strategy. Since hot water was not needed every day, the minimums were zero for the TM and EM strategies, but not zero for the EML method because of the *Legionella* control.

Temperature-matched optimisation

The median electrical energy used for TM was 7.9 kWh/day, 0.3 kWh/day (3.7%) less than the 8.2 kWh/day median for TC. These reductions are presented as distributions in Figures 3 (e) and (f), as daily energy (kWh/day) reduction and percentage reduction, respectively.

Despite the reduction in electrical energy used to heat the water, Figures 3 (b) and (c) show that the thermal energy drawn and the outlet temperature during events, respectively, were the same for TM as for TC. The median event temperatures for TC and TM were both 65 °C. The number of cold events, a metric of negative user satisfaction, did not increase for the TM control, and actually decreased by 3 from 36 out of




Figure 3: Energy and temperature results of the different control strategies represented as distributions for all water heaters. (a) depicts electrical energy used per EWH per day, (b) depicts thermal energy drawn per EWH per day, (c) depicts outlet temperatures during usage events, (d) depicts thermal losses per EWH per day, and (e) and (f) depict the savings achieved in electrical energy per EWH as a reduction in kWh per day and percentage of total used, respectively.

a total of 4486 events. From the user’s perspective, the TC and TM strategies will therefore have no adverse effect on perceived temperatures, despite the energy savings.

Looking at the distribution of savings for the water heaters, we found the percentage reduction for TM, given as 25th percentile, **median**, 75th percentile, was 0.4, **0.6**, 0.8 kWh/day, or 3.6, **7.9**, 11.2 %. This result is in line with the savings reported by Fanney and Dougherty (1996); Goh and Apt (2004), and Cloete (2016), indicating that their approaches employed heating schedules that provided sufficient heating time to effectively achieve results that resemble temperature matching.

Figure 3 (d) shows that the median thermal losses decreased from 2.4 kWh/day to 2.0 kWh/day – a 0.4 kWh/day reduction.

Energy-matched optimisation

The average outlet temperature for EM is significantly lower than for TC and TM, as intended, with a median temperature of 42 °C. The implication is that the user will have to mix in less or no cold water to reach the desired temperature of 40 °C or below. Despite this lower temperature, the thermal energy drawn was higher than for TC – 5.7 kWh/day vs. 5.9 kWh/day – owing to a larger volume draw of 142 L vs. 229 L. The number of cold events for the EM strategy stayed at 0.06, with two fewer cold events than for TC out of 4486 events.

Figures 3 (a) and (b) show that the electrical energy used in the EM strategy was substantially less than either in TC and TM, despite its slightly higher thermal energy delivery. The median electrical energy used for EM was 6.8 kWh/day, 1.4 kWh/day (17%) less than the 8.2 kWh/day median for TC. The reduction is shown in Figures 3 (e) and (f), with respective energy and percentage reductions from the TC baseline of 1.4, **1.5**, 1.7 kWh/day and 12.8, **17.8**, 26.3 %. These results are in line with the simulated and field-measured and/or energy-matched results reported by Booysen et al. (2013); Kepplinger et al. (2015, 2016) and Nel et al. (2018). The median standing loss for this strategy was a mere 1.0 kWh/day, a reduction of 1.4 kWh/day (58%) from TC’s 2.4 kWh/day.

Legionella control

The results for the *Legionella* strategy lie between the TM and EM strategies’ results, as expected. The median electrical energy usage for *Legionella* control was 7.2 kWh/day, 12% less than the median for TC. The results in Table 4 shows that the respective percentage and absolute reductions from TC are 9.1, **13.1**, 17.0 % and 0.8, **1.0**, 1.2 kWh/day. The median standing losses for this strategy was 1.3 kWh/day, a reduction of 1.1 kWh/day from that of TC. The median event temperature for this strategy was 44 °C, while the median number of cold events was the same as for EM.

Unintentional usage loss

The unintentional usage loss, albeit very small, was 0.06 kWh/day for TC and 0.05 kWh/day for TM, indicating that TM did not entirely counter the unintended usage losses, and that most of the savings were due to a reduction in standing losses as a result of the temperature-matched optimised scheduling.

The unintentional energy usage for EM was only 0.02 kWh/day, compared to the 0.06 kWh/day for TC. This result is important, because it demonstrates that EM control limits the detrimental effect of unintentional usage on the electrical energy used by essentially applying a lower temperature control. This result also makes intuitive sense – the amount of energy lost due to unintentional usage loss will be less when the water is at a lower temperature, which is also true for simply operating at a lower set point temperature. It further shows that more savings can be achieved by optimal control than by merely reducing standing losses. Although these numbers are quite small, it stands to reason that the unintended usage is underestimated, which may explain the significantly higher savings reported by some recent field trials (Cloete, 2016; Booysen and Cloete, 2016; Matos et al., 2019). Consider, for example, the warm water that remains in the piping network between the heater and the point of use after an event – the energy in this water also constitutes usage losses, since that water will cool down if not drawn shortly after the initial event. Assuming a conservative volume of 2 L in the piping between the heater and the point of use, each event will in fact comprise legitimate usage and 2 L unintentional usage loss. If the residual warm water in the piping is at a lower temperature than the TC set point, the net losses will be significantly less: the average volume used per user was 142 L/day with an average of 7.5 events/day, which means as much as 10 % of usage would have resulted in usage losses. As an example of unintended use, we have anecdotal evidence of users of a smart EWH controller who became aware

that their casual workers were unnecessarily washing cars or cleaning floors with hot water – behaviour that was quickly remedied (Roux and Booyesen, 2017; Booyesen et al., 2019). It may be argued that this control is essentially the same as setting the thermostat to a lower temperature, but the EM control does better than just supplying a lower temperature on occasion; it ensures that hot water is delivered when needed. However, these savings will be highly dependent on the specifics of each user’s behaviour, each installation’s plumbing setup, including the spouts, and warrants further investigation.

Conclusion

Heating water is a substantial contributor to total domestic electrical energy usage in developing countries. Cost is a concern for users in these countries. The tend to pay a flat fee per kWh, which makes them sensitive to how much energy they use rather than when they use it. A broader concern is that electricity in these countries is generated mainly by burning fossil fuels, which leads to emission of greenhouse gasses. To investigate the potential electrical energy savings that can be achieved by applying optimal (schedule and temperature) control to energy-storing electric water heaters, we examined 30 water heaters over a period of 20 days. We used thermostat control (always on), which is the default mode of water heaters, as our baseline, and compared this strategy with three optimal control strategies. Our three strategies, using dynamic programming to ensure optimal heating, ensure comparable delivery at the outlet of the water heater. The paper describes our problem formulation and our algorithm for applying dynamic programming to the water heater. The first strategy ensures a temperature-matched output with equal volume, the second provides an energy-matched output with decreased but still-hot temperature and increased volume, and the third adapts the second method to ensure that *Legionella* is sterilised despite the lower temperatures. We found that temperature-matching gave a median saving of 7.9%, without adversely affecting the temperature at which water is delivered; the energy-matching method gave a median saving of 17.8%, with a reduction in energy lost to the environment as well as a reduction in energy lost due to unintentional use; and the energy-matching method with daily *Legionella* sterilisation gave a median saving of 13.1%. For all three strategies, the number of cold events did not increase from the baseline strategy. By using real world, not synthesised, hot water usage profiles, we determined the absolute best energy savings, with the fewest cold events, that can be achieved with scheduling and temperature control while causing the minimum of inconvenience to the consumer. However, it should be noted that the results are based on simulations performed with a one-node lumped-parameter model for the EWH thermal dynamics, which means that thermal stratification in the water heater was not taken into account. Further work is therefore needed to evaluate the effect of stratification on the energy savings and number of cold events.

Acknowledgement

The authors thank MTN South Africa for funding through grant S003061, the Water Research Commission (South Africa) for funding the water heater project through grant K1-7163, and Eskom for funding through the tertiary educational support programme.

Reference

- Ahmed, M.T., Faria, P., Vale, Z., 2018. Financial benefit analysis of an electric water heater with direct load control in demand response, in: IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), Aalborg, Denmark, p. TBD.
- Armstrong, P.M., Uapipatanakul, M., Thompson, I., Ager, D., McCulloch, M., 2014. Thermal and sanitary performance of domestic hot water cylinders: Conflicting requirements. *Applied Energy* 131, 171 – 179. doi:10.1016/j.apenergy.2014.06.021.
- Beute, N., Delpont, G., 2006. An historic overview of controlling domestic water heating, in: Proceedings of the 14th Domestic Use of Energy Conference, Citeseer. pp. 41–46.
- BloombergNEF, 2018. Emerging Markets Outlook 2018. <http://global-climatescope.org/assets/data/reports/climatescope-2018-report-en.pdf>. Accessed: February 2019.

- Booyesen, M.J., Cloete, A.H., 2016. Sustainability through intelligent scheduling of electric water heaters in a smart grid, in: 2016 IEEE 2nd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress, pp. 848–855. doi:10.1109/DASC-PICom-DataCom-CyberSciTec.2016.145.
- Booyesen, M.J., Engelbrecht, J.A.A., Molinaro, A., 2013. Proof of concept: Large-scale monitor and control of household water heating in near real-time, in: Proc. of International Conference on Applied Energy (ICAE), Pretoria, South Africa. pp. 1–8. URL: <http://hdl.handle.net/10019.1/95703>.
- Booyesen, M.J., Visser, M., Burger, R., 2019. Temporal case study of household behavioural response to cape town’s day zero using smart meter data. *Water Research* 149, 414 – 420. doi:10.1016/j.watres.2018.11.035.
- Boudreaux, P., Jackson, R., Munk, J., Gehl, A., Dinse, D., Lune, C., 2014. Effect of setup thermostat schedule on heat pump water heater energy consumption, coefficient of performance and peak load, in: 2014 ACEEE Summer Study on Energy Efficiency in Buildings, pp. 1–1 – 1–13.
- Brown, J., 2015. Design and implementation of an intelligent water heater control module for feedback and demand-side management. Master’s thesis. University of Stellenbosch. www.bridgiot.co.za.
- Cloete, A., 2016. A domestic electric water heater application for Smart Grid. Master’s thesis. University of Stellenbosch.
- Diao, R., Lu, S., Elizondo, M., Mayhorn, E., Zhang, Y., Samaan, N., 2012. Electric water heater modeling and control strategies for demand response, in: 2012 IEEE Power and Energy Society General Meeting, IEEE. pp. 1–8.
- Diduch, C., Shaad, M., Errouissi, R., Kaye, M., Meng, J., Chang, L., 2012. Aggregated domestic electric water heater control - building on smart grid infrastructure. 2012 IEEE 7th International Power Electronics and Motion Control Conference - ECCE Asia conference proceedings : June 2-5, 2012, Harbin, China .
- Du, P., Lu, N., 2011. Appliance commitment for household load scheduling. *Smart Grid, IEEE Transactions on* 2, 411–419.
- Enerdata, 2018, 2018. Global Energy Statistical Yearbook 2018. <https://yearbook.enerdata.net/#electricity-domestic-consumption-data-by-region.html>. Accessed: Feb 2019.
- Ericson, T., 2009. Direct load control of residential water heaters. *Energy Policy* 27, 3502–3512. doi:DOI: 10.1016/j.enpol.2009.03.063.
- Eskom, 2013, . Eskom Geyser fact sheet 2013. http://www.eskom.co.za/sites/IDM/Documents/1317_geyser_fact_sheet_no_rmr.pdf. Accessed: May 2018.
- Fanney, A., Dougherty, B., 1996. The thermal performance of residential electric water heaters subjected to various off-peak schedules. *ASME. J. Sol. Energy Eng.* 118, 73–80. doi:<https://doi.org/10.1115/1.2848010>.
- Forlee, C., 1998. Water heating in South Africa facts and figures from the 1997” notch testing” program, in: 5th International Domestic Use of Electrical Energy Conference, Cape Town, pp. 265–270.
- Gerber, S., Rix, A.J., Booyesen, M.J., 2019. Combining grid-tied PV and intelligent water heater control to reduce the energy costs at schools in South Africa. *Energy for Sustainable Development* 50, 117 – 125. doi:10.1016/j.esd.2019.03.004.
- Gholizadeh, A., Aravinthan, V., 2016. Benefit assessment of water-heater management on residential demand response: An event driven approach, in: 2016 North American Power Symposium (NAPS), pp. 1–6. doi:10.1109/NAPS.2016.7747831.
- Goh, C.H.K., Apt, J., 2004. Consumer strategies for controlling electric water heaters under dynamic pricing, in: Working paper CEIC-04-02, pp. 1–8.

- Hohne, P., Kusakana, K., Numbi, B., 2018. Scheduling and economic analysis of hybrid solar water heating system based on timer and optimal control. *Journal of Energy Storage* 20, 16 – 29. URL: <http://www.sciencedirect.com/science/article/pii/S2352152X18303153>, doi:<https://doi.org/10.1016/j.est.2018.08.019>.
- Hohne, P., Kusakana, K., Numbi, B., 2019. A review of water heating technologies: An application to the South African context. *Energy Reports* 5, 1 – 19. URL: <http://www.sciencedirect.com/science/article/pii/S2352484718301495>, doi:<https://doi.org/10.1016/j.egy.2018.10.013>.
- Jack, M., Suomalainen, K., Dew, J., Eysers, D., 2018. A minimal simulation of the electricity demand of a domestic hot water cylinder for smart control. *Applied Energy* 211, 104 – 112. doi:[10.1016/j.apenergy.2017.11.044](https://doi.org/10.1016/j.apenergy.2017.11.044).
- Jacobs, H., Botha, B., Blokker, M., 2018. Household Hot Water Temperature An Analysis at End-Use Level, in: *International WDSA / CCWI 2018 Joint Conference*, Kingston, Ontario, Canada July 23-25, pp. 1–14.
- Jordan, U., Vajen, K., Physik, F., Solar, F., 2001. Realistic domestic hot-water profiles in different time scales. Technical Report. Marburg University.
- Kapsalis, V., Safouri, G., Hadellis, L., 2018. Cost/comfort-oriented optimization algorithm for operation scheduling of electric water heaters under dynamic pricing. *Journal of Cleaner Production* 198, 1053 – 1065. doi:[10.1016/j.jclepro.2018.07.024](https://doi.org/10.1016/j.jclepro.2018.07.024).
- Kepplinger, P., Huber, G., Petrasch, J., 2014. Demand side management via autonomous control-optimization and unidirectional communication with application to resistive hot water heaters, in: *enova 2014 Nachhaltige Gebaeude Versorgung - Nutzung - Integration*, Pinkafeld, Austria, pp. 1–8. doi:[10.23919/DUE.2017.7931855](https://doi.org/10.23919/DUE.2017.7931855).
- Kepplinger, P., Huber, G., Petrasch, J., 2015. Autonomous optimal control for demand side management with resistive domestic hot water heaters using linear optimization. *Energy and Buildings* 100, 50 – 55. doi:<https://doi.org/10.1016/j.enbuild.2014.12.016>.
- Kepplinger, P., Huber, G., Petrasch, J., 2016. Field testing of demand side management via autonomous optimal control of a domestic hot water heater. *Energy and Buildings* 127, 730 – 735. doi:[10.1016/j.enbuild.2016.06.021](https://doi.org/10.1016/j.enbuild.2016.06.021).
- Kepplinger, P., Huber, G., Preißinger, M., Petrasch, J., 2019. State estimation of resistive domestic hot water heaters in arbitrary operation modes for demand side management. *Thermal Science and Engineering Progress* 9, 94 – 109. doi:[10.1016/j.tsep.2018.11.003](https://doi.org/10.1016/j.tsep.2018.11.003).
- Kirk, D.E., 2012. *Optimal Control Theory An Introduction*. Dover Books on Electrical Engineering, Dover Publications, Newburyport.
- Lu, S., Samaan, N., Diao, R., Elizondo, M., Jin, C., Mayhorn, E., Zhang, Y., Kirkham, H., 2011. Centralized and decentralized control for demand response, in: *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES, IEEE*. pp. 1–8.
- Lunacek, M., Ruth, M., Jones, W., Ding, F., 2018. Understanding the impact of electric water heater control on the grid, in: *2018 IEEE Power Energy Society General Meeting (PESGM)*, pp. 1–4. doi:[10.1109/PESGM.2018.8586582](https://doi.org/10.1109/PESGM.2018.8586582).
- Matos, C., Bentes, I., Pereira, S., Faria, D., Briga-Sa, A., 2019. Energy consumption, co2 emissions and costs related to baths water consumption depending on the temperature and the use of flow reducing valves. *Science of The Total Environment* 646, 280 – 289. doi:[10.1016/j.scitotenv.2018.07.290](https://doi.org/10.1016/j.scitotenv.2018.07.290).
- Matos, C., Briga-S, A., Bentes, I., Faria, D., Pereira, S., 2017. In situ evaluation of water and energy consumptions at the end use level: The influence of flow reducers and temperature in baths. *Science of The Total Environment* 586, 536 – 541. doi:[10.1016/j.scitotenv.2017.02.008](https://doi.org/10.1016/j.scitotenv.2017.02.008).

- Monigatti, P., Apperley, M., Rogers, B., 2014. Smart energy interfaces for electric vehicles, in: Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces, ACM, New York, NY, USA. pp. 413–416. doi:10.1145/2598153.2602229.
- Nehrir, M.H., Jia, R., Pierre, D.A., Hammerstrom, D.J., 2007. Power management of aggregate electric water heater loads by voltage control, in: 2007 IEEE Power Engineering Society General Meeting, pp. 1–8.
- Nel, P.J.C., Booysen, M.J., van der Merwe, A.B., 2016a. A computationally inexpensive energy model for horizontal electric water heaters with scheduling. *IEEE Transactions on Smartgrid* doi:10.1109/TSG.2016.2544882.
- Nel, P.J.C., Booysen, M.J., van der Merwe, A.B., 2016b. Energy perceptions in South Africa: An analysis of behaviour and understanding of electric water heaters. *Energy for Sustainable Development* 32, 62–70. doi:10.1016/j.esd.2016.03.006.
- Nel, P.J.C., Booysen, M.J., van der Merwe, A.B., 2018. Saving on household electric water heating: What works best and by how much?, in: Proceedings of IEEE ISGT 2017, Auckland, pp. 1–6. doi:10.1109/ISGT-Asia.2017.8378439.
- Nersa, 2018. Nersa. <http://www.nersa.org.za>. Accessed: January 2019.
- Roux, M., Apperley, M., Booysen, M., 2018. Comfort, peak load and energy: Centralised control of water heaters for demand-driven prioritisation. *Energy for Sustainable Development* 44, 78 – 86. doi:10.1016/j.esd.2018.03.006.
- Roux, M., Booysen, M.J., 2017. Use of smart grid technology to compare regions and days of the week in household water heating, in: 2017 International Conference on the Domestic Use of Energy (DUE), pp. 276–283. doi:10.23919/DUE.2017.7931855.
- Serra, J., 2012. *Electric Vehicles Technology, Policy and Commercial Development*. London: Routledge. doi:10.4324/9780203125755.
- Shaad, M., Momeni, A., Diduch, C.P., Kaye, M., Chang, L., 2012. Parameter identification of thermal models for domestic electric water heaters in a direct load control program, in: Electrical & Computer Engineering (CCECE), 2012 25th IEEE Canadian Conference on, IEEE. pp. 1–5.
- Stone, W., Louw, T.M., Gakingo, G.K., Nieuwoudt, M.J., Booysen, M.J., 2019. A potential source of undiagnosed legionellosis: Legionella growth in domestic water heating systems in south africa. *Energy for Sustainable Development* 48, 130–138. doi:10.1016/j.esd.2018.12.001.
- Stout, J.E., Best, M.G., Yu, V.L., 1986. Susceptibility of members of the family legionellaceae to thermal stress: implications for heat eradication methods in water distribution systems. *Applied and environmental microbiology* 52, 396–399. doi:10.1016/j.semcd.2010.11.003.
- Zuniga, M., Agbossou, K., Cardenas, A., Boulon, L., 2017. Parameter estimation of electric water heater models using extended kalman filter, in: IECON 2017 - 43rd Annual Conference of the IEEE Industrial Electronics Society, pp. 386–391. doi:10.1109/IECON.2017.8216069.