Preprint of https://doi.org/10.1016/j.esd.2018.03.006

Comfort, peak load and energy: Centralised control of water heaters for demand-driven prioritisation

M. Roux[a], M. Apperley[b], M.J. Booysen[a,*]

[a] *MTN Mobile Intelligence Lab and Department of E&E Engineering, Stellenbosch University, South Africa*
[b] *Faculty of Computing & Mathematical Sciences, University of Waikato, New Zealand*

Abstract

Recent advances in smart grid technology enable new approaches to address the problem of load control for domestic water heating. Since water heaters store energy, they are well-suited to load management. However, existing approaches have focused on the electrical supply side, ignoring the obvious link between the user and the grid: individual hot water consumption patterns. This paper proposes a load spreading approach in which water heaters compete for access to the heating medium. The proposed smart grid solution takes grid load limits, real-time temperature measurements, water usage patterns, individual user comfort, and heater meta-data into consideration. The scheduler only turns on the heaters with the highest level of need, but limits the number of *on* heaters to ensure that the grid load stays below a set limit for a set time. The method is evaluated by simulation against various heater set temperature levels, and for various load limits, and compared with ripple control and actual consumption measured in a field trial of 34 water heaters. The proposed algorithm reduces the load from 62kW to 20, 30, 40, and 50kW (vs. 106kW for full ripple control). The resulting number of unwanted cold events is fewer than for ripple control, and only slightly more than no control, while reducing the total energy by 14% from a user-optimised natural experiment.

Keywords: Demand-side management, water heating, peak demand reduction, thermal energy storage, smart grid, user satisfaction.

1. Introduction

Electricity generators are responsible to supply sufficient generation capacity for the worst-case demand, and therefore strive to limit the power demand during peak times for smoother demand profiles [1, 2] Electricity supply is a recurrent concern, particularly in developing countries that experience capacity constraints and limited reserve margins [3, 4, 5, 6]. This challenge is compounded by peak demand periods of the residential load sector coinciding with the peak demand periods experienced by the supply grid [1, 7, 8].

In South Africa, for example, constrained power has led to rolling blackouts and the implementation of large-scale demand-side management (DSM) techniques, such as ripple control on electric water heaters (EWHs), to limit the burden on the grid [1]. In South Africa, which hosts 5.4 million electric water heaters and where instantaneous water heaters are not commonly used, electric water heating is responsible for 7% of the grid load and between 30% and 50% of the residential energy demand [1, 9].

EWH-controlling DSM can be implemented in a variety of ways with financial incentives and direct control being popular choices. Direct control in its simplest form, which is employed in developing countries such as South Africa, is established by means of ripple controllers that are connected to the electric supply of each EWH and controlled centrally, using a signal superimposed on the supply lines to disable EWHs over whole neighbourhoods, towns, or districts. [1, 5]. Direct control is useful during critical demand periods to remove loads from the grid to reduce the demand. However, direct control fails to take the consumer into account,

*Corresponding author
Email address: mjbooysen@sun.ac.za (M.J. Booysen)

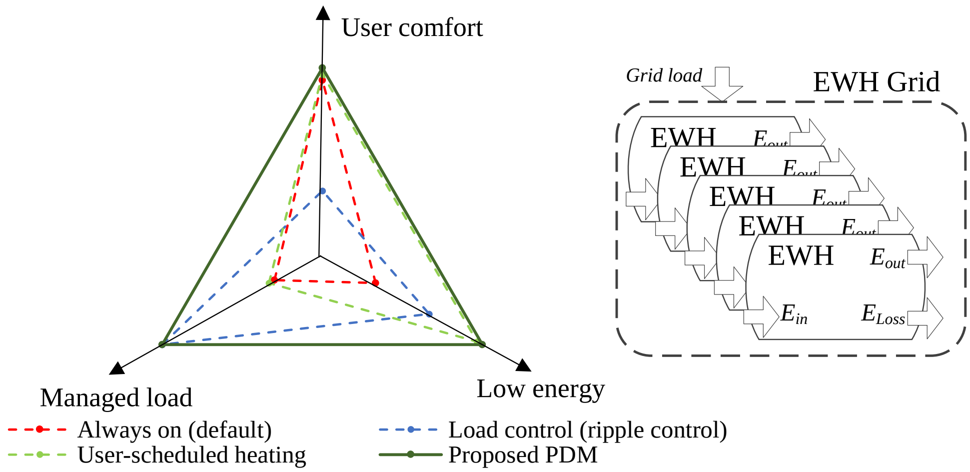

Figure 1: The balancing act with electric water heaters between the impact of high grid loads on the supplier or reseller (*Grid load*), the cost to the individual user (and the environment) of energy consumption (E_{in}), and the impact on the individual user's comfort (E_{out}).

and could lead to users experiencing cold water, especially if the user happens to also apply a coinciding heating schedule of their own to reduce their energy costs. Moreover, since direct control does not address standing losses to the environment and only shifts the load temporally, it does not provide an energy-efficient solution.

On the contrary, consumers want to reduce the energy they consume to limit expenses (and potentially their environmental footprint), while maintaining user comfort (water being hot when needed). In South Africa, water heating is responsible for 35% of the average households' energy consumption [9]. Since cost-sensitive consumers in developing countries (such as South Africa) pay for energy consumed, and are not penalised for the time of use, many consumers resort to applying scheduled heating to only heat water in preparation for hot water usage [10]. The energy used for water heating can be reduced by as much as 29% if intelligent scheduling is employed to only heat water as and when needed, but, the impact of this self-serving scheduling has an adverse impact on the grid [7, 11, 10].

.

These apparently competing objectives for the utility provider and the consumer, depicted in Figure 1, can be managed through the use of the emerging smart grid (SG) technology. One of the key focal points of SGs is improved energy efficiency through load management [12]. The SG concept enables the use of demand response (DR) techniques that have the main objective of matching the consumer consumption to the available supply capability. This objective is achieved either through direct control of specialised devices or by encouraging consumers to adjust their consumption pattern by rescheduling their power demand to off-peak periods and to use less energy overall. However, to make this informed type of direct control possible, a centralised algorithm, bidirectional communications (measure and control) between the EWH controller and the schedule controller, and knowledge of the consumption patterns are required.

1.1. Contributions and Content

This paper presents a novel algorithm that reduces peak demand on the grid through intelligent management of EWHs that have an inherent thermal storage capacity. However, the algorithm bases heating priority on individual users' demand for heated water, and by doing so, also reduces energy by not heating when warm water is not needed. The proposed algorithm uses measured water consumption patterns and set peak load limits and periods for the grid to apply dynamically-calculated priorities for heating scheduling, using real-time measured temperature, observed consumption patterns, and individual EWH metadata as input. For this paper, the water consumed by each EWH over four weeks was used as a representative set of hot water demand for each individual EWH, and that was considered the consumption pattern to which the controller must schedule. The algorithm is simulated with a validated EWH model [13, 14], and data from a field trial of 34 electric water heaters, measured per minute over 28 days.

This paper is organised as follows. Section 2 presents the baseline EWH model and investigates SG scheduling techniques for grid optimisation and maintaining user comfort with EWHs. The proposed scheduling algorithm is developed in section 3. The results of simulations utilising the proposed algorithm is presented and compared in section 5, and section 6 concludes the paper.

2. Related Work

This section contains an overview of the existing literature on direct control of water heaters. The objectives, methods, results, and shortcomings are highlighted. The properties of the related work are summarised in Table 1, at the end of this section.

It is well-established that water heaters are suited to direct control due to their capacity for storing thermal energy vs. their relatively high power ratings and high energy consumption [15]. Various studies have explored the potential of using EWHs for that purpose [15, 16, 17, 18, 19, 20, 21, 22].

The objectives of these studies all fall in the broad categories of

- Thermal modelling of EWHs, which allows assessment of control strategies,
- Utility's peak load shaving by shifting water heaters' load through direct control,
- Users' cost reduction through load shifting under Time of Use (ToU) tariffs,
- Users' cost reduction through intelligent heating schedules, which reduces energy usage,
- Managing users' comfort despite load shifting.

Gholizadeh et al. presents a simple "event-driven" demand response solution to manage EWHs in a smart grid with distributed control in [16]. The objective is to shift the load from peak times by heating during off-peak periods. The proposed algorithm fundamentally controls to high target set temperatures for EWHs out of peak grid load periods, leveraging the energy storing capacity of EWHs. The method is to control the temperature inside water heater to the highest setting when the outdoor ambient temperature is low, and water consumption is less. The paper uses random variations on the ASHRAE 90.2 water consumption profiles to simulate demand. The paper assumes that the user has knowledge of the real-time energy pricing, and makes decisions based on that. Successful deployment of the method is thus dependent on user interaction. The paper claims to achieve lower energy consumption, however, this is compared to an "always on" strategy, rather than the oft-employed user scheduling which reduces energy consumption by as much as 29% [10]. Also, given that the specifics of the thermal model are not given, it is difficult to explain why the higher internal temperatures at lower external temperatures do not contribute to greater standing losses under the algorithm's control. A concern with the paper is that the metric that is controlled to is (only) a target temperature for each EWHs, without concern for either users' temporal demand load, the nett effect on the grid, and optimised energy consumption.

Belov et al. in [22], proposes a solution similar to the one in [16], but with a different definition of user comfort and more exhaustive evaluation against different pre-defined consumption profiles. Again, since comfort is the focus of the algorithm, the impact on energy lost through standing losses, and the effect on the grid are not prioritised. For example, the heating stage before a direct load control (similar to ripple control) causes a significant load spike on the grid, as evidenced by the results in [22]. A shortcoming of this approach is that the higher standing losses due to the higher and earlier heating, requires a longer time spent pre-heating, which may not be available if an event occurs prior to direct load control. Moreover, despite that energy consumption is mentioned as a design objective, the water heaters are individually set to strive for a maximum temperature, which means that unnecessary heating will occur (eventual lost energy), and the effect on the grid load is greater than what is necessary to achieve user comfort.

In [20], Kondoh et al. proposes a method for providing regulation service through direct control of EWHs. The paper presents a thermodymanic model for a vertically-oriented, two-element water heater, that is used in the smart grid simulation for centralised regulation. The control circuitry (two thermostats) of the EWH is overridden and replaced by a controller that is centrally controlled. An algorithm is proposed to control the

Table 1: Features of state-of-the-art solutions for direct control of EWHs, compared with the PDM proposed in this paper.

	[16]	[22]	[20]	[21]	Proposed PDM
Centralised control	✗	✗	✓	✓	✓
No user interaction required	✗	✗	✓	✗	✓
Individual water draw profiles for evaluation	✗	✗	✗	✗	✓
Individual EWH properties and meta-data used	✗	✗	✗	✗	✓
Validated multi-nodal thermal EWH model	✗	✗	✗	✗	✓
Level limits on grid peak load	✗	✗	✗	✗	✓
Temporal limits on peak loads	✓	✓	✓	✓	✓
Time of Use optimisation (Cost)	✓	✗	✓	✓	✗
Energy optimisation (Environmental and cost)	✓	✗	✓	✗	✓
User comfort via event temperature	✗	✓	✗	✗	✓

power consumed by each EWH through implementation of a total power limit, with consideration of recent thermostat switching. In essence, the two parameters controlled for are target internal temperature of the EWH and grid load. However, the objective is to store as much energy in the EWHs as possible, rather than to optimise for the actual consumption pattern and to reduce individual's energy losses to the environment through standing losses. The assumption is made that users do not normally schedule their water heaters for energy efficiency. For the assessment of the algorithm, water draw is extrapolated from measured electric load, assuming "always on" scheduling by the users, which is conducive to loss of accuracy.

In [21], Cui et al. considers a way to integrate solar energy solutions, the supply of which is temporally limited, with water heaters that have the capacity to store the energy. A cooperative gaming strategy is proposed to encourage users to bid for time on the grid in a ToU based set-up. A simple thermal model is used to model the EWHs in the simulation, and a retailers' settlement model that uses the energy spot market and the Elbas market is used to determine the cost for energy. The paper shows financial benefit to the users and improved predictability for the utility. However, the solution requires active participation from users, does not take into account total energy consumed, and does not take into account user comfort (hot water when needed).

2.1. Remaining Challenges

Although various DSM solutions have been proposed for direct control of EWHs ([16, 20, 21, 22]), none of the proposed solutions present a solution with a temporal absolute limit on the peak load (for any time and any power level). This is crucial for demand response in a developing country, where supply could be spurious and conditions dynamic. Moreover, none of the existing solutions provide individualised control based on real-time temperatures, taking into account the individual user's consumption patterns, and the parameters metadata) of the individual user's EWH. Another remaining challenge is that none of the existing methods include a validated multi-nodal thermal model for modelling of the EWH. Since many of the existing solutions have been developed for the developed world, they assume an existing ToU infrastructure and consumer mindset, which is not in place for emerging countries such as South Africa. Importantly, the method proposed in this paper is evaluated against the temperatures experienced by the user and actual water consumption data, which is a crucial metric when evaluating user comfort.

3. Proposed Peak Demand Manager

This section describes the proposed method of the hot water consumption-based Peak Demand Manager (PDM), illustrated in Figure 2. Although merely managing the peak demand is relatively simple and effective from a utility or supplier perspective, it does not take into account individual needs and customer satisfaction, and could lead to unintended power spikes.

The crux of the method is explained here and then expanded upon in the following subsections. The method fundamentally relies on 1) the configurable maximum allowed peak for the SG, 2) for each individual EWH the time required to heat from current temperature~~,~~ to the set temperature, 3) for each individual EWH the time until expected hot water use, and 4) for each individual EWH some metadata (e.g. volume of the vessel and power rating of the element). Heating time: The PDM uses the measured real-time internal temperature and EWH metadata (e.g. power rating and volume) of each of the EWHs to calculate the time

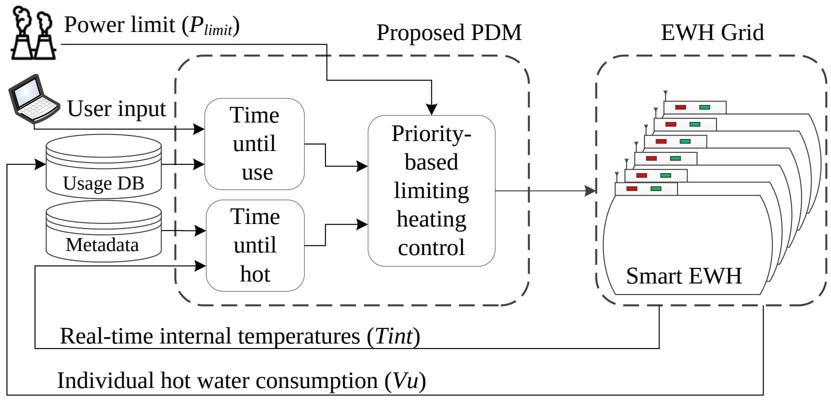

Figure 2: Proposed consumption-based Peak Demand Manager (PDM), which uses heating time and time until expected use to control the heating schedules of each individual EWH, while ensuring the total power limit imposed by the grid is adhered to.

required to heat the water to a temperature that is considered to be hot. Time until use: The PDM uses historic hot water consumption patterns to determine the time left until hot water will be required. However, the method uses a "heat as soon as feasible" approach, which makes it less reliant on the accuracy of the prediction model. It also allows for user-overriding adaptations in case of inadvertent hot water consumption apparent from the pattern.

The PDM uses the time until use and the heating time for each individual EWH to create individual heating priorities for all the EWHs in the SG. The algorithm then uses predefined maximum SG load limits and meta data (power rating of each EWH) to instruct only a limited number of of the highest priority EWHs to turn on for that control cycle.

The objective of the centralised PDM is to limit the peak demand while supporting user satisfaction and user energy cost, by individually controlling the heating of each individual EWH.

3.1. Centralised Peak Limit Management and Prioritisation

The control takes into account for each EWH: when hot water will be needed based on historical consumption patterns or as indicated by the user, and the current temperature of the EWH. This *need* for heated water at a known time (from water currently at a known lower temperature) is then traded off against the available SG peak limit.

For every control cycle of 1 minute, the measured temperatures of each EWH are calculated and the consumption pattern assessed for that control cycle. The PDM uses a time-domain cost function described in section 3.2 to assign to each EWH a heating priority. The N EWHs with highest priority are instructed to then be on for the control cycle, such that

$$P_{\mathrm{EWHs}} = \sum_{n=1}^{N} P_n \leq P_{limit} \tag{1}$$

Where P_n represents the power demand of each individual EWH (based on metadata containing the power rating), P_{EWHs} represents the total power demand of all the EWHs and P_{limit} represents the set peak limit for each time step derived from the remaining power margin determined by the SG.

3.2. Time-Domain Cost Function

An event is defined as the process which starts as soon as the hot water consumption is non-zero (event starting time) and continues until the hot water consumption drops down to zero (event end time). User comfort is defined in terms of event temperature.

To achieve this objective within reasonable time and due to the scale of the data, a heuristic approach was selected. As noted by [23], this approach enables faster, more frugal decision making with potential for higher accuracy than more complex methods.

Table 2: Commercially available EWHs volume vs. power ratings, and distribution of these configurations in data set.

Configuration	Volume (L)	Power (kW)	Count in 34
A	100	2	5
B	150	3	21
C	200	4	8

The load managing scheduler has to calculate a measure of the urgency with which each EWH has to be heated - this measure of urgency is defined by a cost function. The two key factors that drive the cost function are 1) the amount of time before the next water use, and 2) the amount of heating time that is needed to heat the EWH from its current temperature to the hot threshold temperature. The cost function is defined as

$$k_n = -\alpha t_{event} + \beta t_{heat} \tag{2}$$

Where t_{event} is the time before the next expected usage event, and t_{heat} is the heating time.

The amount of time before the next event is determined from the historical hot water consumption pattern, but the conceptual design allows for it to be overridden by each individual user. The amount of heating time required by each individual EWH is determined by 1) the measured internal temperature, 2) the heating element rating, and 3) the volume of the EWH. Using the power rating of the element to heat the water and the volume of the EWH, the time required to heat the body of water, t_{heat}, is given by

$$t_{heat} = \frac{c\rho V(T_{hot} - T_{int})}{P} \tag{3}$$

With T_{hot} (in °C) representing the threshold above which water is considered to be hot, T_{int} (in °C) is the measured internal temperature of the EWH, c (in $Wh/(kg°C)$) the specific heat capacity of water, ρ (in kg/L) the density of water, V (in L) the volume of the EWH, and P (in W) the power rating of the EWH element.

Rewriting (2) gives the generic cost function as

$$k_n = -\alpha t_{event} + \beta \frac{c\rho V(T_{hot} - T_{int})}{P} \tag{4}$$

Table 2 lists nominal EWH volumes vs. power ratings [24], which illustrates a $\frac{V}{P}$ ratio of 50 L/kW. The table also shows the distribution in the dataset from the field trial used in this paper.

A simple initial approach to the cost function is to assume that the coefficients have the same weight, implying that each heating minute has the same urgency weight as a minute preceding the next usage event.

Under this assumption, (2) reduces to

$$k_n = -t_{event} + 210(T_{hot} - T_{int}) \tag{5}$$

This cost function is evaluated at each control step of 1 minute to determine the relative priority of each EWH and establish which EWHs should be allowed to receive power from the limited supply at which time step. The thermal model of the EWHs in the simulation, described in section 2 is run every minute. Due to the numerous time steps that are involved with a high resolution time series simulation, in addition to the multitude of EWHs in the simulation, it is imperative that the cost function is computationally inexpensive. The scalar model was therefore vectorised, allowing concurrent simulation to provide comparative analytic data.

3.3. Event Temperature as Proxy for User Satisfaction

The user comfort metric of event temperature for each EWH is calculated as the weighted average of the per time stamp temperature and hot water usage, as shown in (6).

$$T_{event} = \left(\sum_{u=1}^{U} T_u V_u\right) / \left(\sum_{u=1}^{U} V_u\right) \tag{6}$$

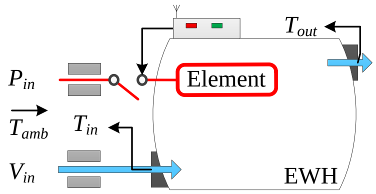

Figure 3: EWH measurement and control setup used for the 34 EWHs, adapted from [7]. The controller controls the electrical supply switch, measures the power (P_{in}), measures the water flow (V_{in}), and measures the inlet, outlet, and ambient temperatures (T_{in}, T_{out}, and T_{amb}, respectively).

Where T_{event} represents the weighted mean event temperature in °C, V_u represents the volume of water withdrawn from the EWH in litres, T_u represents EWH outlet temperature in °C for each time step u during a usage event with duration U.

The objective of the PDM is to ensure that the weighted average event temperature in (6) is above a threshold temperature, $T_{threshold}$, below which the event temperature is considererd to be cold.

4. Experimental Setup and Metrics

The experimental setup, which is adapted from the proposal in Fig. 2, uses measured hot water consumption with simulated EWHs.

The setup consists of the proposed PDM, metadata and measurements from pilot field EWHs, and the validated thermal EWH model, one simulated for each of the field EWHs.

4.1. Measured Data From Pilot in Field

The input data for the simulation is sourced from a pilot study in South Africa, in which smart controllers are retrofitted onto existing EWHs [7]. The experiment uses measurements from individual EWHs measured over a four-week period from 34 smart EWH controllers. The configuration of each EWH controller, shown in Fig. 3, exposes through a remote interface per-minute control of the power supply of the element based on a server-set target temperature (zero for off), while measuring and reporting every minute the following telemetry: volume of hot water consumed (replaced by cold water), temperature of consumed hot water, temperature of cold inlet water, ambient temperature, and electrical energy consumed.

A database is used with metadata, such as the EWH volumes, element power ratings, and thermal resistances of the EWHs.

The individual hot water consumption patterns and metadata are used to stimulate one EWH simulation model for each field EWH, with a simulation time step of one minute. The measured cold inlet temperatures and measured ambient temperatures are also used in the simulation as required by the thermal model. The measured electrical energy is also used as a baseline for comparison against simulation results.

By centralising the PDM scheduling authority, the peak demand on the grid can be managed and limited to a predefined limit – in this experiment, the limit was set to range from 20kW to 50kW in increments of 10kW for the SG of 34 EWHs.

4.2. EWH Model

Nel et al. presents two EWH models, a one- and a two node variety [13]. The one node model is focused on the energy balance equation and considers the body of water contained within the EWH as a single node. The base assumption of this model is a uniform water temperature distribution. The two node model builds on the one node model in that it has the same properties as the one node model during steady state, ie. when the water has had sufficient time to mix and enters a one node state. However, the two node model splits into two nodes, separated by a thermocline, if a sufficiently large hot water usage event occurs. In this state the nodes are stacked vertically due to the buoyancy difference of water at low and high temperatures, with the upper node being the hot, less dense, node and the lower node being the cold, more dense, node. This state is governed by two first-order differential equations which makes it more computationally expensive than the one node state.

Table 3: Experiment parameters for PDM performance analysis.

Parameter	**Value**	**Unit**
Number of EWHs	34	–
Peak limits, P_{limit}	20, 30, 40, 50	kW
EWH volume, V	100, 150, 200	L
EWH element power, P	2, 3, 4	kW
Set point temperature, T_{set}	50, 60, 70	°C
Hot temperature threshold, T_{hot}	40	°C
Specific heat capacity of water $\boldsymbol{c}$	1.1628	$Wh/(kg°C)$
Period	19 Sept–17 Oct	–
Ripple control time period	6–8pm weekdays	–

4.3. Parameters

Table 3 shows the selected parameters for the experiment.

To ensure representative environmental conditions, measured ambient temperature and cold inlet water temperatures were used for each EWH during each simulated minute. The actual tank volume and actual element rating were used. The different combinations and distribution in the used sample set are given in Table 2.

The power limits were chosen as a reasonable distribution stretching from 40% of the measured peak demand for the field units (20kW) to the measured peak demand (50kW).

4.4. Evaluation Metrics

The performance of the algorithm is determined by key metrics from both a grid and consumer perspective. The performance of the PDM, with the various power limits and target set temperatures, is compared to the load profiles measured by the field EWHs, which is labelled “measured” in the results. An implementation of ripple control (as used in South Africa) is also used to provide a comparison~~,~~. This ripple control is active (power turned off) between 18:00 and 20:00 on weekdays. The ripple control is indicated by the label Ripple in the plots and tables.

4.4.1. Peak Demand and Demand Profile

One of the primary goals of the algorithm is to reduce the peak demand of the SG of EWHs, and is one of the direct constraints applied during simulation. As such, this is one of the key metrics to be evaluated, to ensure that the peak demand does not exceed the specified peak limit. As a result of the PDM, the demand profile of the SG will differ from the measured reference value due to the demand shifting that occurs as the EWHs are utilised as thermal energy storage devices. The resultant demand profile provides insight into how the demand has been smoothed by the PDM.

4.4.2. Mean Event Temperature

From the consumer’s perspective, the resultant temperature of the hot water during an event is typically the highest priority. If the experienced temperature is too low, especially during winter months, it will diminish the incentive for the consumer to participate in the peak reduction strategy. This metric is used as one of the consumer comfort indicators.

4.4.3. Cold Event Count

Another consumer comfort indicator is the total number of cold events experienced. A cold event is defined as an event with a mean temperature below that of the selected threshold.

4.4.4. Total Energy Consumption

One of the goals of SGs is to enable more efficient control of energy consumption, which is important for the consumer and the utility. Since user comfort and grid load are optimised, an increase in total energy consumed is expected. The total consumed energy for the various peak limits are compared to the measured reference to provide a baseline for energy consumption performance.

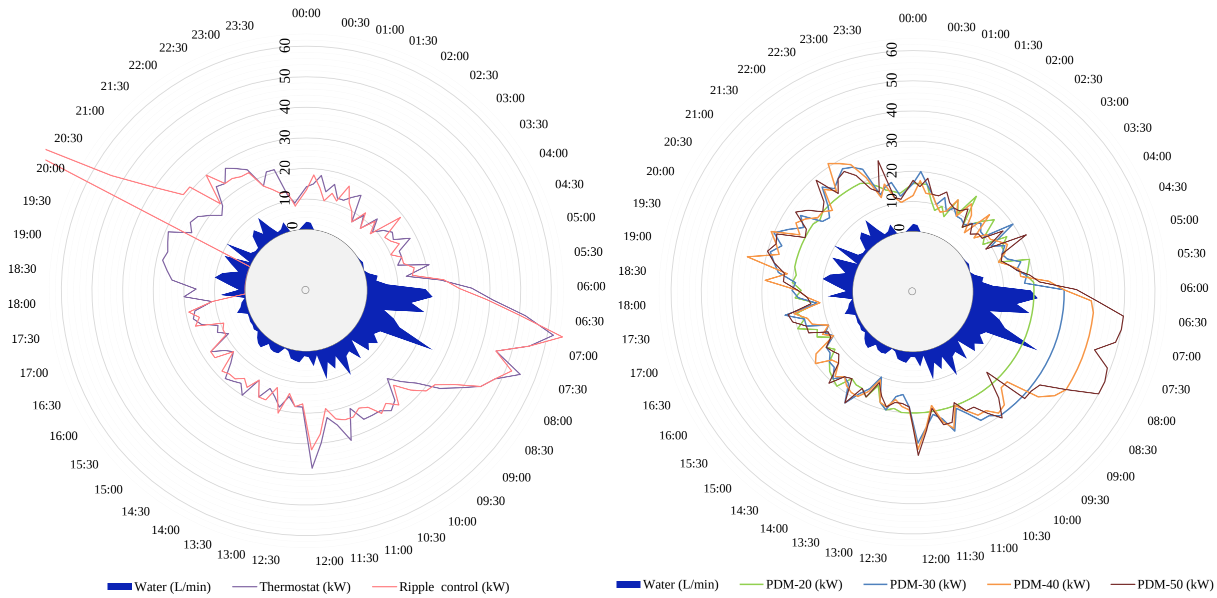

Figure 4: Grid power for a single representative weekday for the different power limits and a temperature limit of 60°C, showing hot water consumption and resulting demand profile. The water consumed by the 34 EWHs is shown in blue. The left plot indicates thermostat control (purple) and ripple control (red). The thermostat control's load almost exactly follows the ripple control's in the morning. A 62kW peak is observed in the morning and an 106kW peak when ripple control ends in the evening. The same day is shown on the right under PDM control, showing the smoothing effect at each level.

5. Results and Discussion

This section reports the results obtained using the experimental setup and selected metrics presented in Section 4.

5.1. Peak Demand and Demand Profile

The demand profile for a typical day, running under $T_{set} = 60$ °C, is illustrated in Figure 4. The peak load under thermostat and under control was 62kW for the 43 EWHs just after 06:30 in the morning, lagging the first peak in water consumption as water is reheated. A second peak of 51kW is observed an hour later, following another peak in consumption. . For the ripple control, the demand peaks at 106kW in the evening, directly following the ripple control period and evening water consumption, and 63kW during the morning hours. Note that the graph does not show the full extent of the 106kW. For the PDM-controlled SG, per design, the peak seen in this period was equal to the peak limit for all four set limits, indicating that the entire available power margin is utilised by the PDM to store energy in the EWHs without surpassing the limit. It is evident from Fig. 4 that a lower power limit increases the duration that the SG is at peak, since the reduced rate requires more time for the same energy delivery. This has the effect of smoothing out the demand profile.

Despite the load profile being smoothed, a potential remains for undue stress on the power switching circuitry due to the PDM's 1 min control cycle. However, the remaining switch count was less under PDM than under thermostat and ripple control for all cases – 12.6% fewer switches per day for the worst case of $P_{lim} = 20$kW and $T_{set} = 50$°C than for thermostat control. This is most likely because cyclic energy top-ups that are required due to standing losses, are less under PDM control.

5.2. Mean Event Temperature

Table 4 captures the mean event temperatures experienced during the month period for each of the power and EWH set temperatures. In the 50°C set limit case, month periods achieved similar mean temperatures than the thermostat and ripple controlled cases. This indicates a high measure of user comfort is maintained even at lower power and lower set point temperatures. At the higher set point of 60°C, the user comfort is maintained at within 2°C of those achieved by the thermostat control. This trend continues at the set limit of

Table 4: Mean event temperatures in°C for the given power limits and target set temperatures.

Set Point (°C)	**50**	**60**	**70**
Power Limit (kW)	Month	Month	Month
Thermostat	48	57	66
Ripple	48	57	65
20	47	55	61
30	48	56	64
40	48	57	65
50	48	57	66

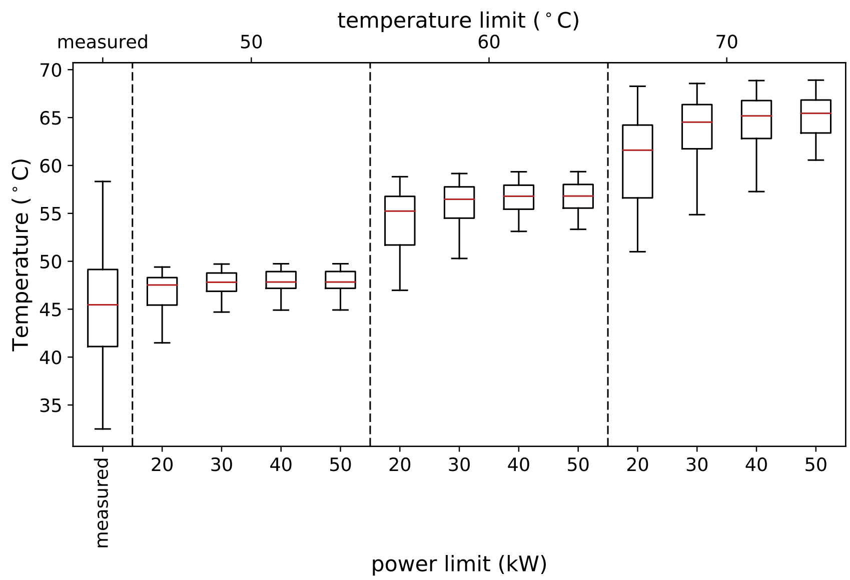

Figure 5: Distribution of event temperatures over a month: Measured, and simulated at various power and temperature limits.

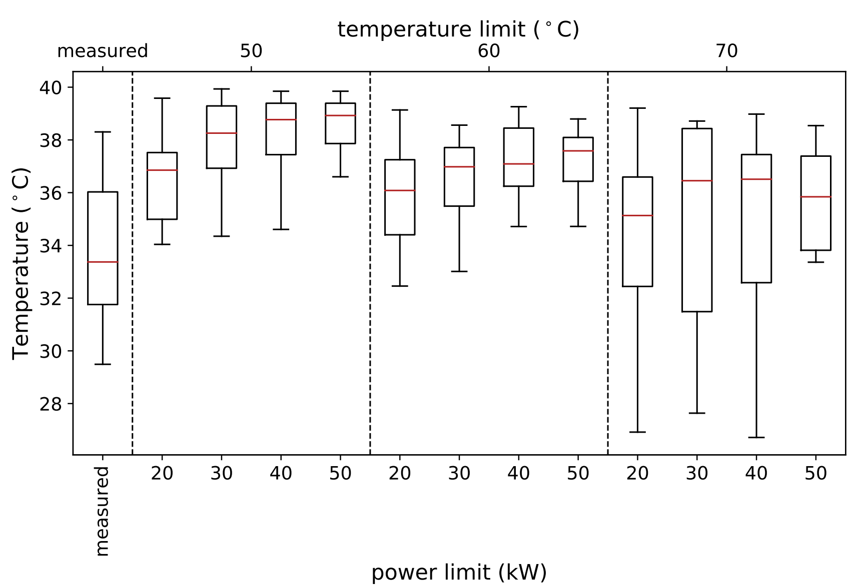

Figure 6: Distribution of cold event temperatures over a month: Simulated at various power and temperature limits.

Table 5: Total number of cold events experienced (as a percentage of 12614) for the given power and temperature limits.

Set Limit (°C)	**50**	**60**	**70**
Power Limit (kW)	Month	Month	Month
Thermostat	3.7	2.0	1.5
Ripple	5.0	2.4	1.7
20	10.4	10.3	11.0
30	4.6	3.8	4.2
40	3.8	2.4	2.1
50	3.7	2.1	1.6

70°C, with the temperatures over 21°C above the hot threshold, T_{hot}, but 5°C colder than the thermostat case – indicating that less water is unnecessarily heated for the lower power limits, while water is still delivered above the threshold. Fig. 5 indicates that per unit this trend continues, with all units maintaining a sufficiently high mean event temperature, even with low set- and peak limit constraints. With ripple control, the mean event temperatures are equivalent to a power limit of 40kW over the three temperature limits.

Although T_{hot} is used as a metric for the comfort through the temperature of usage events, it is important to note that the EWHs are still controlled to the higher set point temperatures (T_{set}) listed in Table 3. The risk of *Legionella* growth was evaluated and found to be mitigated in all the permutations. The shortest period continuously spent in any day at or above 50 °C was observed to be 190 minutes (observed under the most constrained permutation of P_{lim} = 20kW and T_{set} = 50°C), which is longer than the *Legionella* sterilisation time of 111 min at that temperature [25, 26].

Considering the temperature spread of events with event temperature below the threshold (cold events), indicated in Fig. 6, the cold events have a mean temperature of 38°C and a median temperature of 37°C for a temperature limit of 50°C. This is close to the temperature threshold of 40°C which indicates many cold events are just below this threshold. As expected, the higher the temperature and power limits are set, the fewer number of cold events are experienced. The mean and minimum cold event temperatures decrease as the temperature limit is increased, this is most likely due to fewer, but colder events experienced as a result of excessively large event sizes outside of the EWH specifications.

Although the *measured* events seem significantly lower, this may be misleading because measured outlet temperature is captured with a temperature sensor that is strapped onto the outlet pipe, which is not ideal for accurate measurement of transient temperature flows. The measured values tend to be a few degrees lower due to the non-ideal thermal coupling and thermal inertia of the sensor. They are reported only as an indication of spread.

5.3. Cold Event Count

Table 5 lists the number of cold events for the different scenarios as a percentage of the total of 12614 hot water usage events recorded during the sample period. In the most energy-constrained case with the lowest set point and power limit of 50°C and 20kW respectively, the PDM resulted in 1309 cold events.Increasing the power limit to 40kW, the cold events were 487, a reduction of 62.8% from the 20kW limit. Increasing the set point to 60°C, with a 20kW limit, only slightly reduces the number of cold events, which indicates that the available energy is the main constraint, and not a too low temperature (i.e. low potential thermal capacity). This is substantiated by the increase in cold events for a 70°C temperature and limit of 20kW, indicating a prioritisation that is affected by the high set temperature and resulting in a less equitable distribution of energy. This reduction and then increase in cold events with increasing set temperatures is also seen for the 30kW limit – the thermal energy shortfall is throttled by the available grid peak power.

However, for the 40kW grid limit, the increase in temperature does monotonically result in a decrease in cold events, indicating that sufficient supply is provided from the grid to source the thermal energy despite higher set point vs. measured temperature differentials.

Also interesting to note is that the number of cold events does not reduce significantly from 40kW to 50kW for the 50°C set temperature(3.8% to 3.7%), indicating that the limitation in that instance is not primarily the power available from the grid, but rather the energy storage capacity of the EWH at that set point temperature. This is substantiated by the cold event reduction from 40kW to 50kW for the 70°C set temperature (from 2.1% to 1.6%), where more energy storage capacity reduces the load burden while maintaining hot water supply.

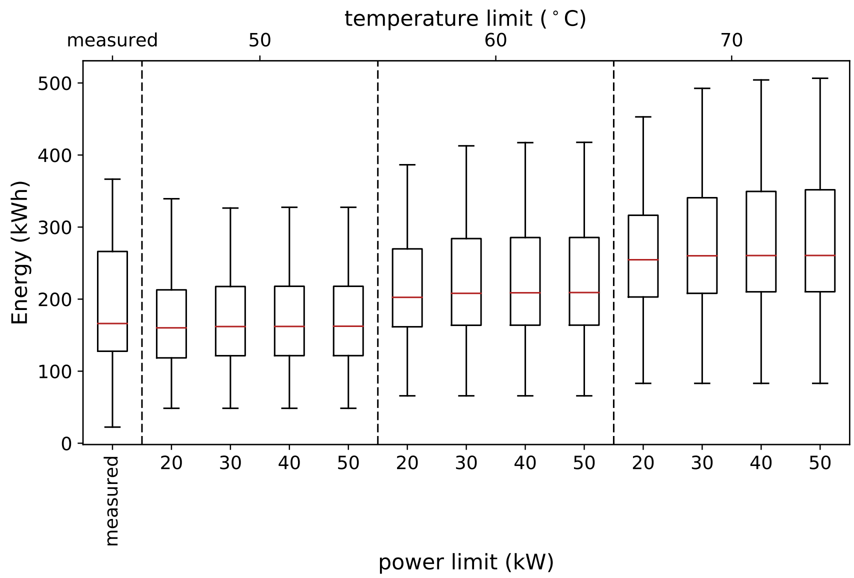

Figure 7: Consumed energy distribution for a month at various power and temperature limits.

The ripple control caused more cold events than the algorithm controlled instances above and equal to 30kW limits for 50°C. In the 50°C temperature limit case, the ripple control incurred 629 cold events for the month, while the 30kW power limited case only incurred 589. Importantly, this indicates that a lower temperature limit with an always available but level-limited power source under PDM control, performs better than a ripple controlled instance with no power limit. Furthermore, the number of cold events for ripple control with no power limit is consistently higher than always on, and also higher than or equal to power limited control for a power limit of 40kW at a temperature set point of 60°C. At 70°C, the ripple control incurs 215 cold events compared to 206 cold events for a 50kW limit. This result indicates that mean event temperature. the metric used in the previous section, is not a sufficient condition for user comfort.

Also important to note is that the majority of cold events were caused by two high-consumption EWHs – without these two units, the number of cold events for the month with a power limit of 20kW and set point of 60°C, reduces from 1303 to 684.

5.4. Total Energy Consumption

Table 6 lists the energy consumed by the SG for the selected set points and peak limits. As expected the set limit of 70°C significantly exceeded the energy consumption of the measured SG due to the higher energy storage capacity. Furthermore, with ripple control, the energy consumption was slightly less than the second highest power limit, 40kW, regardless of the temperature limit.

Of note is that the energy measured in the natural experiment falls between the 60°C and 50°C levels, indicating that consumers are effectively operating their EWHs at those temperatures.

An improvement to the approach presented in this paper is to limit the heating period to a shorter time before the consumption events, which will reduce the energy consumption further, and move the grid load to lead the water consumption profile, similar to that seen in [11] and [7].

6. Conclusion

Although electric water heaters have high power levels and consume large amounts of energy, their inherent thermal storage capacity makes them ideal for direct control. Various solutions have been proposed to leverage this property using direct control, but none of the existing approaches cater for all of the following: load on the grid with a configurable level and temporal limit, reduced energy consumed (for environmental and cost benefit) through low temperatures and heat-when-needed, user comfort through event temperature, and individualised flow demand as input. This paper presents a heuristic SG solution in which EWHs compete for

Table 6: Total energy consumed in **MWh** on the grid for the given power and temperature limits.

Set Limit (°C)	**50**	**60**	**70**
Power Limit (W)	Month	Month	Month
Thermostat	6.43	8.33	10.19
Ripple	6.37	8.25	10.09
20	6.18	7.75	9.21
30	6.40	8.20	9.86
40	6.43	8.31	10.12
50	6.43	8.33	10.18
Measured		7.4	

the heating medium, and that addresses all of the above by incorporating individual consumption patterns, knowledge of individual EWHs temperature and metadata, and configurable (time and level) grid load limits. The following metrics were used for the evaluation: load profiles, mean event temperatures, number of cold events, and energy consumed. The results show that the mechanism outperforms measured results from 34 field units, and also outperforms the oft-employed ripple control, without sacrificing significant user comfort.

References

[1] N. Beute, G. Delport, An historic overview of controlling domestic water heating, in: Proceedings of the 14th Domestic Use of Energy Conference, 2006, pp. 41–46.

[2] H. O. R. Howlader, M. Furukakoi, H. Matayoshi, T. Senjyu, Duck curve problem solving strategies with thermal unit commitment by introducing pumped storage hydroelectricity renewable energy, in: 2017 IEEE 12th International Conference on Power Electronics and Drive Systems (PEDS), 2017, pp. 502–506. doi:10.1109/PEDS.2017.8289132.

[3] H. Naude, Demand-side management by centralized ripple control system, in: 2nd Domestic Use of Electrical Energy Conference, Cape Town, April 1995, pp. 71 – 82. doi:10.1109/ISGT-Asia.2015.7387131.

[4] C.-A. Gabriel, J. Kirkwood, S. Walton, E. L. Rose, How do developing country constraints affect renewable energy entrepreneurs?, Energy for Sustainable Development 35 (2016) 52 – 66. doi:https://doi.org/10.1016/j.esd.2016.09.006.
URL `http://www.sciencedirect.com/science/article/pii/S0973082616305762`

[5] A. Wilken, G. Delport, Making money with intelligent hot water load control, in: The Domestic Use of Energy at the Eastern African Power Industry Convention 2000, Kampala,, 2000, pp. 128–146.

[6] A. Sebitosi, Energy efficiency, security of supply and the environment in south africa: Moving beyond the strategy documents, Energy 33 (11) (2008) 1591 – 1596. doi:https://doi.org/10.1016/j.energy.2008.08.003.
URL `http://www.sciencedirect.com/science/article/pii/S0360544208001916`

[7] M. Roux, N. Naude, M. Booysen, A. Barnard, Electric water heaters in smartgrids: Individual savings versus network peak load management, South African Universities Power Engineering Conference (SAUPEC).

[8] H. O. R. Howlader, M. Furukakoi, H. Matayoshi, T. Senjyu, Water heating ins south africa facts & figures from the 1997 notch testing program., in: 5th International Domestic Use of Electrical Energy Conference, Cape Town, 2017, pp. 265–270.

[9] Eskom Geyser fact sheet 2013, `http://www.eskom.co.za/sites/IDM/Documents/1317_geyser_fact_sheet_no_rmr.pdf`, accessed: March 2018 2016.

[10] M. J. Booysen, A. H. Cloete, Sustainability through intelligent scheduling of electric water heaters in a smart grid, in: IEEE DataCom 2016, 2016, pp. 848–855. doi:10.1109/DASC-PICom-DataCom-CyberSciTec.2016.145.

[11] M. Roux, M. J. Booysen, 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), 2017, pp. 276–283. doi:10.23919/DUE.2017.7931855.

[12] X. Fang, S. Misra, G. Xue, D. Yang, Smart Grid; The New and Improved Power Grid: A Survey, IEEE Communications Surveys & Tutorials 14 (4) (2012) 944–980. doi:10.1109/SURV.2011.101911.00087. URL `http://ieeexplore.ieee.org/document/6099519/`

[13] Nel, PJC and Booysen, MJ and Van Der Merwe, B, A computationally inexpensive energy model for horizontal electrical water heaters with scheduling, IEEE Transactions on Smart Grid.

[14] A. H. Cloete, A domestic electric water heater application for Smart Grid., Master's thesis, Stellenbosch: Stellenbosch University (2017).

[15] Z. Xu, R. Diao, S. Lu, J. Lian, Y. Zhang, Modeling of electric water heaters for demand response: A baseline pde model, IEEE Transactions on Smart Grid 5 (5) (2014) 2203–2210. doi:10.1109/TSG.2014.2317149.

[16] A. Gholizadeh, V. Aravinthan, Benefit assessment of water-heater management on residential demand response: An event driven approach, in: 2016 North American Power Symposium (NAPS), 2016, pp. 1–6. doi:10.1109/NAPS.2016.7747831.

[17] L. Zeng, Y. Sun, Q. Ye, B. Qi, B. Li, A centralized demand response control strategy for domestic electric water heater group based on appliance cloud platform, IEEJ Transactions on Electrical and Electronic Engineering 12 (2017) S16–S22. doi:10.1002/tee.22549. URL `http://dx.doi.org/10.1002/tee.22549`

[18] G. M. Binini, J. L. Munda, O. D. Dintchev, Comparison of hot water draw models integrated within a statistical physically-based model for electrical demand of domestic water heater, in: 2017 IEEE AFRICON, 2017, pp. 1155–1160. doi:10.1109/AFRCON.2017.8095645.

[19] H. Y. Song, S. C. Park, Y. T. Yoon, A design scheme of the smart appliances to minimize energy cost without reducing customer's utility, in: 2015 IEEE Power Energy Society General Meeting, 2015, pp. 1–5. doi:10.1109/PESGM.2015.7286369.

[20] J. Kondoh, N. Lu, D. J. Hammerstrom, An evaluation of the water heater load potential for providing regulation service, in: 2011 IEEE Power and Energy Society General Meeting, 2011, pp. 1–8. doi:10.1109/PES.2011.6039149.

[21] Q. Cui, X. Wang, X. Wang, Y. Zhang, Residential appliances direct load control in real-time using cooperative game, IEEE Transactions on Power Systems 31 (1) (2016) 226–233. doi:10.1109/TPWRS.2015.2391774.

[22] A. Belov, A. Vasenev, P. J. M. Havinga, N. Meratnia, B. J. van der Zwaag, Reducing user discomfort in direct load control of domestic water heaters, in: 2015 IEEE Innovative Smart Grid Technologies - Asia (ISGT ASIA), 2015, pp. 1–6. doi:10.1109/ISGT-Asia.2015.7387131.

[23] G. Gigerenzer, W. Gaissmaier, Heuristic Decision Making, Annual Review of Psychology 62 (1) (2011) 451–482. doi:10.1146/annurev-psych-120709-145346.

[24] Kwikot, Kwikot 600 Dual, Electric Water Heaters for both Horizontal and Vertical Installations, Slimline and Standard (2005).

[25] P. J. Dennis, D. Green, B. P. C. Jones, A note on the temperature tolerance of legionella, Journal of Applied Microbiology 56 (2) (1984) 349–350.

[26] P. M. Armstrong, M. Uapipatanakul, I. Thompson, D. Ager, M. McCulloch, Thermal and sanitary performance of domestic hot water cylinders: Conflicting requirements., Applied Energy 131 (2014) 171–179.