Efficient Heuristics for Power Constrained Planning of Thermostatically Controlled Loads

Frits de Nijs, Mathijs M. de Weerdt, and Matthijs T. J. Spaan
Delft University of Technology
The Netherlands
{f.denijs, m.m.deweerdt, m.t.j.spaan}@tudelft.nl

ABSTRACT
Thermostatic loads are promising energy buffers that could be unlocked through Demand Response. Existing research has shown that an aggregation of thermostatic loads can be made to follow a power curve using reactive control. In this paper we investigate the use of planning to proactively control an aggregation of Thermostatic Loads as storage to overcome temporary dips in power availability.

We first present a formal problem definition of the planning problem under consideration, and apply Dynamic Programming to it to obtain an optimal solution. However, because we can prove the problem to be NP-complete, we need efficient heuristic solutions to solve practical instances. Therefore, we then extend the Dynamic Programming solution with approximations of the state and action space to obtain efficient heuristic solutions. We evaluate our proposed heuristics through simulation on instances with a range of power capacity constraints, and show that they are able to obtain a factor two improvement over the reactive approach under the hardest power constraints.

1. INTRODUCTION
A large potential storage capacity can be found in the various heat buffers operated by consumers: Houses, refrigerators and hot water reservoirs all need to be maintained at a certain temperature offset from the environment temperature. All of these buffers decay to the environment temperature over time, which requires constant action to counteract. But precisely when these buffers are heated or cooled can be shifted in time. By storing more energy in the heat buffer now, we can later 'extract' this heat from the buffer in the form of reduced loads. In this sense, we may think of heat buffers controlled by thermostats as a kind of batteries [4]. The storage potential in these batteries can be exploited through Demand Response.

These Thermostatically Controlled Loads (TCLs) can themselves be controlled automatically by the system operator, provided that an ICT infrastructure is in place. Existing work [2, 4] has mainly focused on using Control Theory to make an aggregation of TCLs closely follow an available power signal. There, control is applied reactively on a second timescale to obtain a solution for the immediate grid balancing, which means that both fluctuations up and down must be followed, and that the power available is expected to cover the demand on average.

Instead of using TCLs for balancing, we investigate the potential of TCLs for buffering temporary drops in power production, such as those caused by the fluctuations in renewable generation. These fluctuations typically occur on longer timescales [6], potentially violating the consumers comfort constraints if no preventive action is taken. This motivates using planning to respond to fluctuations before they occur, to minimize the discomfort experienced by the end-users.

In this paper we present a novel Power Constrained Planning (PCP) problem formulation, in which the deviation from the TCL set-point is optimized under power availability constraints. There are two reasons for considering available power as a constraint. Firstly, while excess power can be discarded, without power the TCLs cannot operate. Secondly, there may be too little power to satisfy the comfort constraints for a period of time. But even in this case we want to *minimize* the total discomfort. Unfortunately, we also prove that PCP is an NP-complete problem, which means that computing an optimal solution is intractable. Therefore, we present two heuristics for an optimal Dynamic Programming [3] algorithm and evaluate their performance against an existing Control Theory solution.

To evaluate the proposed heuristics, we simulate how the temperatures in an aggregation of 500 heterogeneous TCLs evolve under a number of scenarios with increasingly tight power constraints. For each scenario we compare the total error resulting from the control actions of the proposed algorithms with a control scheme. We find that as the power constraints become severe, the heuristics outperform the conventional control system by a factor two, while requiring less than 10 seconds of planning time.

The organization of this paper is as follows. First the mathematical model of a TCL is presented in Section 2. Then Section 3 defines the PCP problem we are studying here. An optimal Dynamic Programming algorithm is given in Section 4, and extended to two heuristic solutions in Section 4.3. Sections 5 and 6 evaluate the performance of the algorithms on respectively a single device and an aggregation. Section 7 describes the related work, while Section 8 states our conclusions and future work.

2. BACKGROUND

A thermostatic load is any device that is able to consume (electric) power for the heating or cooling of a body in relation to its natural temperature, such as refrigerators or central heating systems. The goal of the thermostat is to operate the device so that the temperature of the body remains as close as possible to the set-point at all times.

A Markov chain model of thermostats controlled by hysteresis controllers was presented by Mortensen and Haggerty [8]. In their model, the temperature of the body in the next time step $\theta_{i,t+1}$ is deduced from the current temperature $\theta_{i,t}$, the current outside temperature θ_t^{out}, a temperature input from the device θ_i^{pwr}, and a random temperature shift $\theta_{i,t}^{\text{rnd}}$ modeling exogenous actions such as opening a door.

The temperature input from the device is determined through the interaction of the effective thermal power produced P_i (kWh) with the thermal resistance R_i (°C / kW) as $\theta_i^{\text{pwr}} = R_i P_i$. It signifies the maximum temperature offset obtainable when the device is permanently switched on. The hysteresis controller switches the device state depending on the measured temperature $\theta_{i,t}$ relative to the set-point $\theta_{i,t}^{\text{set}}$. To protect the device from short-cycling, the set-point is surrounded by a dead-band θ_δ. Hysteresis control is modeled through the boolean variable $m_{i,t}$. The controller maintains the previous state $m_{i,t+1} = m_{i,t}$, unless the temperature boundaries are violated. In case of a heating device:

$$m_{i,t+1} = \begin{cases} 0, & \left(\theta_{i,t} > \theta_{i,t}^{\text{set}} + \frac{\theta_\delta}{2}\right) \\ 1, & \left(\theta_{i,t} < \theta_{i,t}^{\text{set}} - \frac{\theta_\delta}{2}\right) \\ m_{i,t}, & \text{otherwise} \end{cases} \tag{1}$$

The size of the time step Δ, together with the thermal constants R_i and C_i (the thermal capacitance, kWh / °C) determine how quickly the current temperature responds to the external factors through the fraction $a_i = \exp \frac{-\Delta}{R_i C_i}$, resulting in the model:

$$\theta_{i,t+1} = a_i \theta_{i,t} + (1 - a_i)\left(\theta_t^{\text{out}} + m_{i,t}\theta_i^{\text{pwr}}\right) + \theta_{i,t}^{\text{rnd}}. \tag{2}$$

Finally, the power extracted from the grid is determined through the efficiency μ_i of the heater, thus $P_i^{\text{eff}} = \frac{P_i}{\mu_i}$. Then, we can determine the power consumption of i at time t to be $m_{i,t} P_i^{\text{eff}}$.

3. PROBLEM DEFINITION

In this section we formally define the problem of planning the activation of thermostatic loads under power constraints. First we present the mathematical model of the planning problem. Subsequently, we prove that it is an NP-complete problem, showing that computing the optimal solution is intractable.

3.1 Power Constrained Planning (PCP)

We adapt the unconstrained model from Section 2 into a Power Constrained Planning problem. We use boldface characters to represent vectors of device parameters over all devices, i.e. $\boldsymbol{\theta}_t = \begin{bmatrix}\theta_{0,t} & \theta_{1,t} & \cdots & \theta_{n,t}\end{bmatrix}$. Then, using the Hadamard product $\boldsymbol{b} = \boldsymbol{a} \circ \boldsymbol{\theta}_t \implies b_i = a_i \times \theta_{i,t} \quad \forall i$, we can define a state transition function to computes $\boldsymbol{\theta}_{t+1}$ as $f(\boldsymbol{\theta}_t, \boldsymbol{m}_t, \theta_t^{\text{out}}) = \boldsymbol{a} \circ \boldsymbol{\theta}_t + (1 - \boldsymbol{a}) \circ \left(\theta_t^{\text{out}} + \boldsymbol{m}_t \circ \boldsymbol{\theta}^{\text{pwr}}\right)$

With this function, we can define a planning problem using a given horizon h, the thermal properties of the n thermostatic loads with initial temperatures $\boldsymbol{\theta}_0$, the predicted outdoor temperature θ_t^{out}, and the predicted power production L_t. A solution is a device activation schedule that never consumes more power than is available while minimizing the cost function $c(\boldsymbol{\theta}_t)$. The entire planning problem becomes:

$$\begin{aligned} \underset{[\,\boldsymbol{m}_0\ \boldsymbol{m}_1\ \cdots\ \boldsymbol{m}_h\,]}{\text{minimize}} \quad & \sum_{t=0}^{h} c(\boldsymbol{\theta}_t) \\ \text{subject to} \quad & \boldsymbol{\theta}_{t+1} = f(\boldsymbol{\theta}_t, \boldsymbol{m}_t, \theta_t^{\text{out}}) \\ & \boldsymbol{m}_t \cdot \boldsymbol{P}^{\text{eff}} \leq L_t \\ & m_{i,t} \in [0, 1] \qquad\qquad \forall i, t \end{aligned} \tag{3}$$

Due to the generality of the model, the controlled loads can have different objectives which can be expressed through the cost function. Besides typical functions such as the squared error $c(\boldsymbol{\theta}_t) = \sum_{i=0}^{n}\left(\theta_{i,t} - \theta_{i,t}^{\text{set}}\right)^2$ or maximum set-point deviation $c(\boldsymbol{\theta}_t) = \max_i \left(\theta_{i,t} - \theta_{i,t}^{\text{set}}\right)$, we might consider more application specific functions, such as the comfort level of a home owner which might be affected more by low temperatures, or a refrigerator which only incurs high penalties when the temperature gets too high.

3.2 PCP is NP-complete

We demonstrate that the Power Constrained Planning problem is an NP-hard problem via reduction from Partition, which is an NP-complete problem [5].

Given an instance of the Partition problem consisting of a set of n integers $s_i \in S$, we construct an instance of PCP with a horizon of one time-step. The power production in the only time-step is equal to $L_0 = \frac{1}{2}\sum_{i=0}^{n} s_i$. The power consumption of the n thermostatic loads is mapped to the weight of the equivalent integer in S as $P_i^{\text{eff}} = P_i = s_i$, while the thermal constants of each load are set to $R_i = 2, C_i = 1$ for all i. The time step Δ is set such that $a_i = \exp \frac{-\Delta}{R_i C_i} = \frac{1}{2}$, thus $\Delta = -2 \ln \frac{1}{2}$. The initial temperature $\theta_{i,0}$ and the outside temperature θ_0^{out} are both set to 0 for all i. Finally, the set-point is equal to s_i, and the cost function is simply the sum of all errors, thus $c(\boldsymbol{\theta}_t) = \sum_{i=0}^{n}(s_i - \theta_{i,t})$. Does there exist a schedule with error at most $\epsilon \leq \frac{3}{2}\sum_{i=0}^{n} s_i$?

It is easy to see that the generated instance has a solution iff S contains a partition. Consider how the temperature at the horizon $\theta_{i,1}$ depends on the state of the device $m_{i,0}$:

$$\begin{aligned} \theta_{i,1} &= a_i \theta_{i,0} + (1 - a_i)\left(\theta_0^{\text{out}} + m_{i,0} R_i P_i\right) \\ &= \begin{cases} \frac{1}{2} \times 0 + \frac{1}{2} \times (0 + 2s_i) = s_i, & m_{i,0} = 1 \\ \frac{1}{2} \times 0 + \frac{1}{2} \times (0 + 0) = 0, & m_{i,0} = 0 \end{cases} \end{aligned} \tag{4}$$

At time $t = 0$, all loads incur an error of size s_i because $\theta_{i,0} = 0$. The error in time step $t = 1$ depends on $m_{i,0}$. Those devices that were turned on do not contribute to the error, while those that remained off again contribute s_i. Because the contribution to the power demand $\sum_{i=0}^{n} m_{i,0} P_i$ is the inverse of the contribution to the error, the power limit L_0 imposes the constraint that *at most* half of the total load can be active at the same time, while the decision requires *at least* half of the total load to be active. A partition of S can thus be represented by setting $m_{i,0} = 1$ for all i in the 'left' partition, since this requires exactly the available power while imposing exactly the required error. In the other direction, if the generated PCP instance contains a solution, then that solution cannot have more than half of the total load active. But in order to be a solution, it must activate at least half the total load to match the required error. Thus the decision partitions the set into exactly half on, half off.

To prove that PCP is NP-complete, we must further show that we can verify a certificate in polynomial time. To verify that the certificate schedule $\begin{bmatrix} \boldsymbol{m}_0 & \boldsymbol{m}_1 & \cdots & \boldsymbol{m}_h \end{bmatrix}$ is a solution, we need to test that it never exceeds the power constraints L_t, and that the total error is less than the decision value ϵ. To test the power constraints, we must compute $\boldsymbol{m}_t \cdot \boldsymbol{P}^{\text{eff}} \leq L_t$ for all t, which takes $O(nh)$. Testing the decision requires evaluating $\sum_{t=0}^{h} c(\boldsymbol{\theta}_t) \leq \epsilon$, which consumes $O(c(\boldsymbol{\theta}_t)h)$. Provided that h and $O(c(\boldsymbol{\theta}_t))$ are at most polynomial in n, verifying a certificate for PCP takes polynomial time in n.

4. DYNAMIC PROGRAMMING

The defining aspect of the PCP is the Markov model of the temperature progression. Such a state transition model is easy to embed in a Dynamic Programming solution, which we will present in this section. Unfortunately, since PCP is NP-complete, finding an optimal schedule for a PCP instance is an intractable problem. Therefore, in this section we also propose heuristics that we expect will allow the Dynamic Programming solution to return good plans quickly.

4.1 Optimal Dynamic Program

Using the notation developed in the previous section, we can define the following recurrence for the power constrained planning problem:

$$\begin{aligned} p(\boldsymbol{m}_t) &= \begin{cases} 0, & \boldsymbol{m}_t \cdot \boldsymbol{P}^{\text{eff}} \leq L_t \\ \infty, & \boldsymbol{m}_t \cdot \boldsymbol{P}^{\text{eff}} > L_t \end{cases} \\ S(\boldsymbol{\theta}_t) &= \min_{\boldsymbol{m}_t} \left(c(\boldsymbol{\theta}_t) + p(\boldsymbol{m}_t) + S(f(\boldsymbol{\theta}_t, \boldsymbol{m}_t, \theta_t^{\text{out}})) \right) \\ S(\boldsymbol{\theta}_{h+1}) &= 0 \end{aligned} \tag{5}$$

By computing $S(\boldsymbol{\theta}_0)$ we obtain the schedule that results in the minimum possible error. This scheme can be implemented using the standard backward recursive dynamic programming algorithm, but such a scheme will quickly become intractable since the state space $\boldsymbol{\theta}_t$ and the action space $\boldsymbol{m}_t$ are both exponential in n. The next sections propose heuristics to reduce the size of the state and action space.

4.2 Price-Based Switching Heuristic

Since it is intractable to control each device separately, we are looking for a heuristic that abstracts away the activating and deactivating of the loads. Considering that the loads are expected to contain a thermostat, the obvious choice would be to let each device be switched by its hysteresis controller. What then remains is to decide which devices are allowed to stay on when there is not enough power to run them all.

We propose to use a greedy heuristic to control the switching of the devices. Sort the devices that want to be on by the estimated reduction in error from being on. Switch as many as feasible, starting from the highest error reduction. We assume devices can estimate their error contribution through the difference between the measured temperature and the set-point. This heuristic may be implemented as the price-based mechanism proposed by Koch, Zima, and Andersson [6], who demonstrate a controller that assigns a price to switching device state based on its estimated error. A similar heuristic was proposed by Hao, Sanandaji, Poolla, and Vincent [4], who call it a priority-stack-based control algorithm.

4.3 State Estimation of Price-Based Aggregation

Given the greedy switching heuristic, this section deals with estimating its current state. Following that, we also present two action spaces to control the aggregation. In this section we assume without loss of generality that the devices are for heating.

When the devices are similar, we may estimate their state by computing the state of the average device. In the following, we use an overline to indicate an average, thus $\bar{\theta}_t = \frac{1}{n} \sum_{i=0}^{n} \theta_{i,t}$ is the average temperature. Unconstrained and in steady state, we expect the average temperature to equal the average set-point. Otherwise, we may bound the development of the temperature as follows. The lower bound on the temperature is obtained when all devices are off. It is equal to the average decay to the outside temperature, thus:

$$\bar{\theta}_{t+1}^{\min} = \bar{a}\bar{\theta}_t + (1 - \bar{a})\theta_t^{\text{out}}. \tag{6}$$

The upper bound depends on the number of devices that can be switched on. Suppose that we know the maximum possible power consumption at the grid is $P^{\max} = \sum_{i=0}^{n} P_i^{\text{eff}}$, then we may estimate the maximum average duty cycle to be equal to $D = \frac{L_t}{P^{\max}}$. Given this duty cycle, the maximum temperature can be approximated from a period of heating of length $D\Delta$ and a period of cooling of length $(1 - D)\Delta$ as

$$\begin{aligned} \bar{\theta}_{t+1}^{\max} = \bar{a}\bar{\theta}_t + &\left(1 - \exp \frac{-D\Delta}{\bar{R}\bar{C}}\right) \left(\theta_t^{\text{out}} + \bar{\theta}^{\text{pwr}}\right) + \\ &\left(1 - \exp \frac{-(1 - D)\Delta}{\bar{R}\bar{C}}\right) \theta_t^{\text{out}}. \end{aligned} \tag{7}$$

Thus, the state transition is estimated by function $\bar{f}$:

$$\bar{\theta}_{t+1} = \bar{f}(\bar{\theta}_t, \bar{\theta}_t^{\text{set}}, \theta_t^{\text{out}}) = \begin{cases} \bar{\theta}_{t+1}^{\min}, & \bar{\theta}_t^{\text{set}} < \bar{\theta}_{t+1}^{\min} \\ \bar{\theta}_{t+1}^{\max}, & \bar{\theta}_t^{\text{set}} > \bar{\theta}_{t+1}^{\max} \\ \bar{\theta}_t^{\text{set}}, & \text{otherwise} \end{cases} \tag{8}$$

To control the aggregation we propose to use a global set-point offset θ_t^{off} as state, so that $\bar{\theta}_{t+1} = \bar{f}(\bar{\theta}_t, \bar{\theta}_t^{\text{set}} + \theta_t^{\text{off}}, \theta_t^{\text{out}})$.

This approach was taken by Callaway [2] to make an aggregation of thermostatic loads closely follow a load signal. Because maintaining a range of possible offsets as actions may be expensive, we further propose a heuristic that considers just the three extremes, heating, maintaining and cooling: $\theta_t^{\text{off}} \in \left[\bar{\theta}_{t+1}^{\min}, 0, \bar{\theta}_{t+1}^{\max}\right]$.

5. SINGLE DEVICE PLANNING

In order to demonstrate the value of planning for load balancing, this section evaluates the performance of planning a single load. Planning a single device is tractable and thus we are able to use the optimal algorithm specified in Section 4.1. This algorithm is compared with the baseline case using the hysteresis controller, and with the proposed three-action state approximate model from Section 4.3.

The experimental setup is as follows: We simulate a fictional household heated by a small heat-pump in winter, under intermittent power from a renewable source such as wind energy. In particular, thermal capacitance $C = 12$, thermal resistance $R = 14$ and heating power $P = 3.2$. The horizon is set to 6 hours, and during this time the outside temperature remains constant at 10 degrees, while the occupants request a set-point of 20 degrees inside. For these 6 hours there are predictions for three short outages, during which no power is available for running the heat-pump. In this scenario the consumers want to remain warm, and therefore the system is penalized more for temperatures lower than the set-point, compared to temperatures higher than the set-point. Small fluctuations in temperature within the 1 degree deadband of the hysteresis control are not penalized. Thus the error function becomes:

$$c(\theta_t) = \begin{cases} (\theta_t - \theta_t^{\text{set}})^2 - \frac{1}{2}^2, & \theta_t > \theta_t^{\text{set}} + \frac{1}{2} \\ (\theta_t - \theta_t^{\text{set}})^4 - \frac{1}{2}^4, & \theta_t < \theta_t^{\text{set}} - \frac{1}{2} \\ 0, & \text{otherwise} \end{cases} \tag{9}$$

Given this scenario, we expect the optimal algorithm to pre-heat the home before each outage, thereby reducing the amount by which it drops below the set-point. This should result in a reduced error compared to the hysteresis control. In order to not unfairly judge the control system, we reduced its deadband to $\left[-\frac{1}{8}, \frac{1}{8}\right]$, and made it maintain an average temperature of $20\frac{3}{8}$, thereby ensuring it stays at the top of the allowed deadband under normal operation. For all experiments, the control system updates device state with an interval of 1 second. With respect to the three state model from Section 4.3, because we are planning a single device, there is no state estimation, and thus we expect it to reduce to the optimal algorithm.

Figure 1 presents the results of the deterministic simulation. For these scenarios each plan was produced within 10 seconds. In this figure the two left-most graphs show the performance of optimal planning, while the middle and right graphs show how performance deteriorates when the planning system is limited to taking an action every two or five minutes respectively. As the bottom graphs show, the cumulative error is much smaller when using the planning algorithms, achieving around $\frac{1}{4}$ of the error incurred by control only. As expected, this reduction in error is obtained by 'over-heating' the house before the predicted outages, which can be seen in the rise in temperatures before the periods of no power indicated by the vertical bands.

Two interesting observations can be made from the results. The dynamic programming solutions use the error free zone as a buffer, by trying to stay at the top of the band at all times. When the intervals between actions become too large at five minutes (right figures), we see that the DP algorithm decides to increase the temperature above the deadband to incur the least possible error. Secondly, we observe that the three state algorithm (DP3) is able to avoid this temperature peaking by delegating to the high frequency control mechanism by using the 'Maintain' action during the periods of power availability. As a consequence, it actually incurs a smaller total error than the DP algorithm in the five minute interval case.

Using planning interval sizes that are larger than the modeled time step size has two advantages. In the first place, it reduces the number of states to evaluate, and thereby reduces the time needed to compute the plan. Secondly, and more importantly from a practical standpoint, larger planning intervals prevent short-cycling the controlled device by ensuring it stays in one state for the duration of the interval. Fewer commands also reduces the constraints on the required communication interface. We conclude that the good performance for the larger planning intervals demonstrates that planning has merit outside a theoretical scenario.

6. MULTI DEVICE PLANNING

To see if the positive results obtained in the previous result carry over to the multiple device scenario, this section evaluates the performance of the heuristic planning algorithms proposed in sections 4.3 and 4.3. The scenario considered here is similar to the scenario proposed in the previous section except for the number of devices and the nature of the outages. In this multi device setting, we consider houses of varying build quality by simulating 500 houses with capacitance $C \in U[10, 14]$ and resistance $R \in U[12, 16]$. The heat-pumps are assumed to be identical for all houses.

When multiple devices are considered, a reduction in the available power may switch off a subset of the controlled devices. Then, some of the devices may be getting too cold even when the average temperature is within the deadband. We thus expect that the difficulty of the problem depends on the magnitude and frequency of the power outages. Therefore, the size of the outages is varied on three levels: during 'Mild' outages, approximately two-thirds of the average unrestricted load remains available, while during 'Severe' outages at most one-third of the needed power is available. 'Full' outages are like the binary case considered in the single device setting, which means that during an outage no power remains. The frequency of the outages is controlled through the probability of an outage occurring on each time step (the planning algorithms are still provided with accurate predictions of the outages before planning).

Under these varying constraints we expect the planning algorithms to perform better than just the price-based control system whenever the power constraints are severe enough to cause dips in the temperature below the deadband. Further, because the performance of the group of devices must

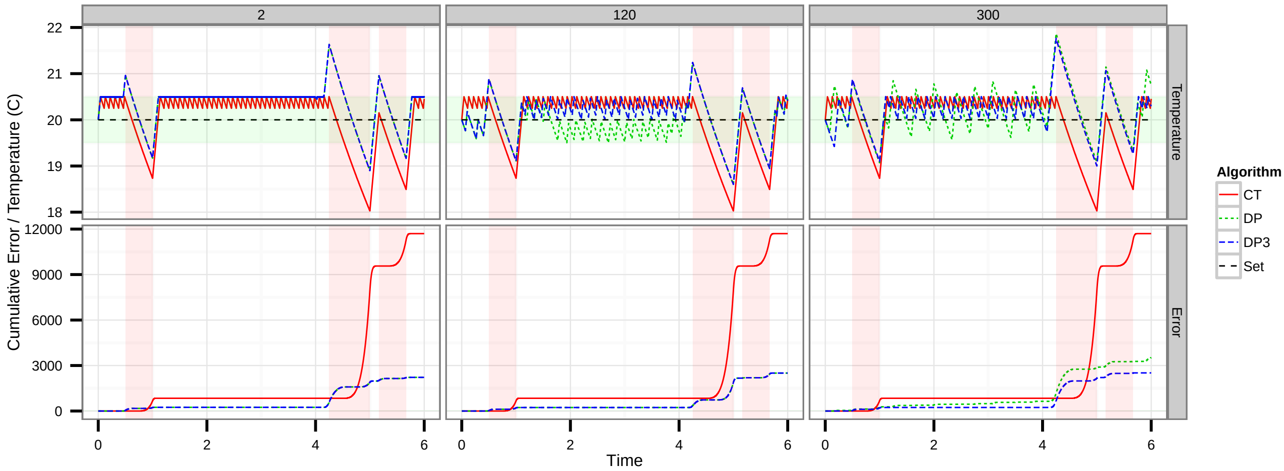

Figure 1: Comparison of (optimal) planning (DP) with hysteresis control (CT) and the three-state planning algorithm (DP3), under three planning interval sizes. Left: two second interval, middle: two minute interval, right: five minute interval. The top figures present the temperature over time, while the bottom figures present the cumulative error. The vertical (red) bands highlight the three power outages, while the horizontal (green) band shows the error-free region.

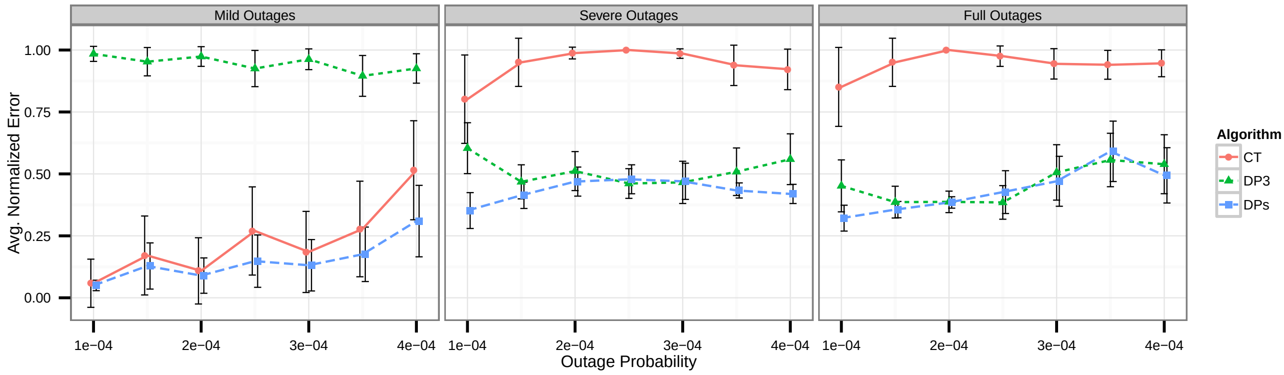

Figure 2: Performance of the heuristic Dynamic Programming models DP3 (Equation 4.3) and DPs (Equation 4.3) under varying levels of power constraints.

be estimated from the power availability, we expect the best performance when outages are 'Full' outages, which are estimated accurately by $\bar{\theta}^{\min}_{t+1}$.

Figure 2 presents the results of the multi-device experiments. Since the heuristics plan one 'average' device, running times are still below 10 seconds. In these figures, the y-axis presents the normalized error: The method with the highest error is assigned a normalized error of 1, and the others are then computed through their ratio to the worst method. The results are averaged over 20 runs per setting. We make a number of observations: As the severity of the outages increases, the benefit of planning becomes larger compared to just using control. But when the outages become very frequent, errors start becoming unavoidable which translates into an increasing error for increasing outage frequency for the DP heuristics. The three-state planning algorithm does not perform well under 'Mild' outages, which is a consequence of the continuous error incurred from staying just inside the error-free band 'on average'. The more fine-grained offset-based algorithm (DPs) instead consistently outperforms the control mechanism.

7. RELATED WORK

A similar setting is investigated by Rogers, Maleki, Ghosh, and Jennings [9], who consider optimizing the energy cost and carbon intensity of a single household by planning the activation of a heater using a smart predictive thermostat. Because our algorithm builds plans for an aggregations of households, our solution is able to take into account the interactions that occur due to periods of high simultaneity. Planning only individual households risks overloading the grid capacity during periods of cheap electricity caused by high renewable production.

Mathieu, Kamgarpour, Lygeros, and Callaway [7] propose controlling thermostatic loads as arbitrage, essentially making them follow the inverse of a real-time price signal. They consider a control scenario where loads are only allowed to be controlled inside their deadbands, which is unobtrusive to the end-user. Their focus is on reducing the energy costs of consumers, under the assumption that sufficient power is available to maintain temperature inside the deadbands as required.

Stadler, Krause, Sonnenschein, and Vogel [10] present the dynamics of a system of set-point controlled refrigeration units. They examine the requirements of the controllers embedded into the units by considering two different types of controller, one which only accepts binary signals requesting an increase or decrease in load, and a more powerful controller that allows to specify a start time and duration during which reduction of load is requested. They demonstrate that when devices are equipped with the advanced controller, a 100% load reduction can be sustained for more than 45 minutes if given a 20 minute lead time to cool the refrigerators. Another contribution of their paper is a mathematical approximation model of the state of the aggregated fridges. Their results demonstrate that planning is necessary to get the most out of the storage capacity of these devices, which matches our observations.

A planning problem definition related to our own is presented by Bosman, Bakker, Molderink, Hurink, and Smit [1]. The authors consider the planning of a fleet of MicroCHP units, so that their profit is maximized while the heat demand and minimum and maximum power production constraints are satisfied. The local search algorithm used first builds plans for each device individually using Dynamic Programming, which are then combined into a global plan and checked against the power constraints. If constraints are violated, the objective function of each house is updated and the devices replanned until a feasible solution is found. Compared to our contribution the authors do not consider time-variable power constraints. Their heuristics are also much more computationally intensive, which limits their experiments to planning interval sizes and number of devices planned of 15 minute intervals over 24 hours and 100 units respectively, while requiring more than an hour of computation time.

8. CONCLUSIONS AND FUTURE WORK

In this paper we investigated the algorithmic challenges to using Thermostatically Controlled Loads for buffering for periods of low power availability. To this end we presented a planning problem definition for the optimal control of Thermostatically Controlled Loads under power constraints. Because this is an NP-complete problem, we extended an optimal dynamic programming approach to two heuristic algorithms. From evaluating these heuristics on a range of different power constraints we can conclude that using planning techniques can strongly reduce the deviation from the set-point that occurs as a consequence of intermittent power drops. Planning becomes particularly effective when the operating constraints are more severe.

Building upon this work we want to look at a number of extensions to the solutions proposed in this paper. In the first place, we may try to make the proposed heuristics more accurate through the use of clustering. By grouping devices with 'similar' heat response together, we can reduce the error incurred by averaging their properties. Secondly, we want to look at modeling and handling uncertainty in the power availability predictions. Topics include investigating how uncertainty can best be represented, and whether these algorithms can be made suitable for handling this uncertainty. Related to this, we also want to investigate how we can build online versions of these algorithms. One question that needs to be answered for an online version is how large the planning interval should be. It is likely that the length of the planning horizon must depend somehow on the thermal properties of the device being planned.

References

[1] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. "Planning the production of a fleet of domestic combined heat and power generators". In: *European Journal of Operational Research* 216.1 (2012), pp. 140–151.

[2] D. S. Callaway. "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy". In: *Energy Conversion and Management* 50.5 (2009), pp. 1389–1400.

[3] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. "Dynamic Programming". In: *Introduction to Algorithms.* 3rd. The MIT Press, 2009, pp. 359–413.

[4] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. "A Generalized Battery Model of a Collection of Thermostatically Controlled Loads for Providing Ancillary Service". In: *Annual Allerton Conference on Communication, Control and Computing.* 2013.

[5] R. M. Karp. "Reducibility among Combinatorial Problems". In: *Complexity of Computer Computations.* 1972, pp. 85–103.

[6] S. Koch, M. Zima, and G. Andersson. "Active Coordination of Thermal Household Appliances for Load Management Purposes". In: *Symposium on Power Plants and Power Systems Control.* 2009, pp. 149–154.

[7] J. L. Mathieu, M. Kamgarpour, J. Lygeros, and D. S. Callaway. "Energy Arbitrage with Thermostatically Controlled Loads". In: *European Control Conference.* 2013, pp. 2519–2526.

[8] R. E. Mortensen and K. P. Haggerty. "A stochastic computer model for heating and cooling loads". In: *IEEE Transactions on Power Systems* 3.3 (1988), pp. 1213–1219.

[9] A. Rogers, S. Maleki, S. Ghosh, and N. R. Jennings. "Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming". In: *Second International Workshop on Agent Technology for Energy Systems.* 2011, pp. 71–78.

[10] M. Stadler, W. Krause, M. Sonnenschein, and U. Vogel. "Modelling and evaluation of control schemes for enhancing load shift of electricity demand for cooling devices". In: *Environmental Modelling & Software* 24.2 (2009), pp. 285–295.