

# Identifying Flexible Pool Pumps Suitable for Distributed Demand Response Schemes

Lachlan L. H. Andrew<sup>1</sup>, Draga Doncila Pop<sup>2</sup>, Reza Razzaghi<sup>3</sup>, David L. Dowe<sup>2</sup>

1 School of Computing and Information Science, University of Melbourne

2 Faculty of Information Technology, Monash University

3 Department of Electrical and Computer Systems Engineering, Monash University

E-mail: l.andrew@ieee.org

**Abstract.** Demand response will be an important tool as non-dispatchable generation is added to the grid. Swimming pool filtration pumps are a promising appliance for the grid operator to control because, unlike air conditioners, their time of operation can be shifted by a day without affecting their primary role. Although many customers set and forget the timer settings, this paper studies smart meter data and demonstrates that many change the settings frequently. A “care index” is proposed that quantifies how much a customer appears to care about the precise timing of pump operation, which indicates how willing they are likely to be to relinquish control to the grid operator.

## 1. Introduction

Electricity networks are undergoing several revolutions. On one hand, driven primarily by the need to reduce greenhouse gas emissions, generation is increasingly being provided by renewable sources such as wind and solar power, which are much less centralized than traditional generators. On the other, increased penetration of information technology and improving power electronics is allowing more distributed control.

Historically, demand was taken as given and generation was adjusted to meet it. Since wind and solar generation cannot increase in adverse weather, there is increasing need for demand side management, such as demand response. This exploits the IT revolution to allow coordinated actions by many customers to reduce the load at times of high demand or of low generation. In this paper, we will investigate the potential of controlling swimming pool filtration pumps to play a part in this demand response.

Demand response can take many forms, which will be described in Section 1.1. Currently, it is used to combat peak load, which often occurs on hot days when the air conditioning load is high. Consequently, many schemes reduce air conditioning during peak times. However, as an increasing fraction of generation is intermittent, it will be necessary to curtail load at times other than peak demand, due instead to drops in generation. Another household appliance with significant power draw is a pool filtration pump.

One of the biggest barriers to adoption of demand response is that users interact directly with many appliances, such as washing machines, and indirectly with others, such as through reduced comfort when heating ventilation and air conditioning (HVAC) settings are changed. Pool pumps are ideal in this

regard, since users do not interact with them directly, instead operating them via timers, and their core function of keeping the pool clean is not affected by even a few days of inactivity. This makes it plausible that users would not care at all when their pumps operate, provided they operate for sufficient hours on average. The main contribution of this paper is to investigate that hypothesis by studying how and how often customers change the timer settings of their pool pumps. The results suggest that surprisingly many people do seem to care enough about when their pumps operate, enough to change the timer settings frequently.

### *1.1. Frequency control via distributed demand response*

Concerns related to climate change and aging fossil fuel generators are driving the deployment of renewable energy sources (RESs). For example, in Australia around 60% of existing coal-fired generation capacity will retire by 2040 [1] and the installed capacities of medium scale solar (100 kW to 5MW) and wind have increased 5- and 3-fold respectively during the past five years [2]. The volatility and uncertainty of these RESs makes it challenging to balance the demand and supply and to maintain the reliable and secure operation of power networks. This balance is reflected in the frequency of the power system, which must be kept within a narrow band ( $\pm 0.15$  Hz in Australia).

Power system operators employ various mechanisms at different time scales to manage the imbalance introduced by load or variable generation forecast errors and by contingencies. Operating reserves are historically provided by the unused capacity of synchronous generators, with wind and solar dispatched to their full output. The progressive integration of RESs and consequently higher variability of generation requires more headroom for the generators and also faster ramp rates. In such scenarios, there is an increasing need for higher spinning reserves and fast-start generation units. At the same time, retiring synchronous generators result in higher frequency excursions due to low inertia. Thus, frequency regulation services using conventional generators might be insufficient.

A supplementary solution is providing frequency control reserves from loads. This approach has been utilized in conventional power systems by using under frequency load shedding (UFLS) [3], in which network operators use frequency relays to interrupt and curtail loads after a contingency to reduce the demand. However, this approach is typically employed during severe contingencies and can affect the reliability of supply for a large number of consumers.

The most efficient solution would be making use of the flexibility of aggregates of deferrable loads. With the advanced communication links deployed in Smart Grids, it is possible to implement Demand Response (DR) based schemes via Direct Load Control (DLC) or Indirect Load Control (ILC). Different appliances have been studied for DR including Thermostatically Controlled Loads (TCLs) such as refrigerators [4], air-conditioning and water heater systems [5], [6] or pool pumps [7], [8]. The flexibility harnessed from various appliances can provide frequency control ancillary services at different time scales [9], [10]. TCLs can contribute to primary frequency regulation using frequency-responsive load controllers [11] and stochastic control algorithms [12]. DR can provide spinning reserve using aggregated residential appliances [13]. The availability of higher spinning reserve will enable grid operators to better manage large contingency events, resulting in less load curtailment and blackouts.

### *1.2. The Pool Pump*

Pool pumps may be suitable for providing spinning reserve [2]. Compared to TCLs that have thermal inertia limitations, pool pumps have excellent load-shifting capability and can provide the reserve for many hours. They have been identified as the appliance with the most flexible operating schedule in a study attempting to quantify the potential for load shifting in household appliances [14]. However, this study was survey-based without measuring customers' likelihoods to participate in the DR schemes. Another survey-based study in Southern California identified 55% of pool maintenance providers unconcerned about a 16 week program during which a pump would be remotely turned off for 2-6 hours one day per week [15]. The same study identified potential to shift 180MW of grid load at midday.

However, this survey-based study only considered 152 houses and for large scale deployment more systematic analysis is necessary.

This paper proposes a method that can be scaled to study tens of thousands of customers. It extracts pump use from smart meter data by applying rule-based artificial intelligence, as described in Section 2. It then derives three indices to estimate how much a customer cares about the pump’s operating schedule in Section 3, and demonstrates in Section 4 that the three are in moderate agreement.

These indices can be used to target future surveys or advertising campaigns to customers. To be most useful, they should be calibrated by surveys, but that is beyond the scope of this paper.

## 2. Identifying Pumps

### 2.1. Consumption Data

This study is based on aggregate electricity consumption data from 953 residential customers in Melbourne, Australia collected in 2013 that appear to have swimming pools. The data measures the energy consumption per half hour. Figure 1 shows this for two households.

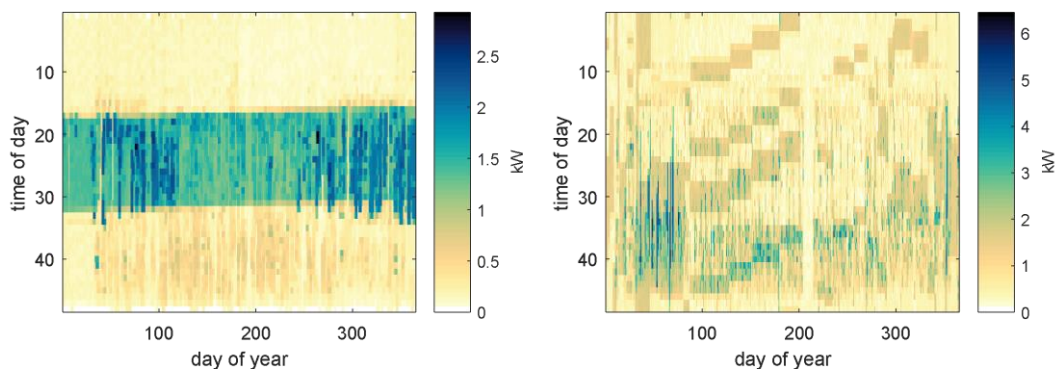


Figure 1. Half hourly consumption data for two households in 2013.

In the first subfigure, there is a band of increased power consumption from about sample 18 (9am) to sample 32 (4pm), due to a device on a timer. The timer has a drift, causing the appliance to be turned on slightly earlier each day. Based on the power consumption and the duration of use (similar to the daily recommended use of 8 hours [16]), it appears highly likely that this is a pool pump. The second also has an appliance on a timer, but the timer settings change frequently. The challenge is to detect such appliances automatically from the full collection of 60,000 households.

### 2.2. Detection of pump use

This study used a heuristic to identify pool pump usage, followed by manual correction of the estimates, based on visualizations as in Figure 1. Unfortunately, no ground truth was available. The primary heuristic used is that pumps are turned both on and off by a timer (not, say, a thermostat), and that power consumption is around 0.5 to 1 kW. If users never changed timer settings, then this would be simple, but changes are actually quite common. The heuristic works in three phases, at increasing levels of abstraction.

The first phase identifies possible turn-on and turn-off times by identifying runs of consecutive days on which the power consumption jumped up or down by a consistent large amount. This abstracts from individual samples into “runs” of days. The trustworthiness of a run is estimated, in part, by the variance over days of the jump in power consumption.

The second phase matches portions of the turn-on runs with corresponding turn-off runs to form periods of pool use. This abstracts from runs into “rectangles”. The final set of rectangles left after the cleaning described below is denoted  $R$ .

The third phase adjusts start and end dates, and turn-on and turn-off times, such as extending rectangles to include missed days, truncating them to avoid overlap, merging neighbouring rectangles and deleting “noise” rectangles. This abstracts from individual rectangles to a consistent collection, representing timer settings and the times they change. It also detects instances where settings differ depending on the day of the week, and timer drift.

This sounds straightforward, but there are many challenges.

1. The “signal” (consumption of the pool pump) is often swamped by “noise” (other devices), which is strongly correlated. When a household wakes up, it causes a fairly uniform increase in power consumption, which resembles a device on a timer. Noise is reduced by taking a “rolling minimum” over both days and consecutive samples within a day, and by preferentially discarding rectangles that resemble the pattern of other devices, such as the morning routine or heaters on winter evenings.
2. Turn-on events typically occur within a half-hour sample, and so rather than having an abrupt change from “off” to “on”, there is a sample between the fully-off and fully-on samples that has a fraction of the pump power. This makes it difficult to distinguish steps from gradual ramps.

The remaining challenges arise from complications in the pool pump use itself, and are informative about customer behaviour.

3. Many timers have a “weekend” setting, allowing the turn-on and turn-off times to differ for two consecutive days of the week. Many customers use this function but do not correctly set the date of the timer, so that the alternative settings are not on the weekend.
4. Most pumps use constant power when they are on, but many have multiple operating powers. If the power is changed without the timer settings changing, then the variance of the power jumps is high, which sometimes cause pumps to be considered unreliable.
5. Timers drift, and so the edges to be detected are not horizontal. This can be mistaken for multiple timer settings.

Given these challenges, how can we know that the devices were pool pumps and not some other device? This will be addressed again in the results section, but for now we make two observations.

First, many other devices on timers can be ruled out. Hot water and electric vehicles may be turned on by a timer, but their duration of operation is controlled by a thermostat or state of charge. Some heaters run on timers, but we exclude devices not used in summer. Water features and security lights may operate on timers, but consume much less power than a pool pump. This provides confidence that the identified devices are indeed pool pumps.

Second, surveys say that 8% of households in Melbourne have pools and the above procedure flagged about half of that. Some may have pumps permanently on, or controlled manually, but we have many more false negatives than false positives. This again suggests that most of these are indeed pool pumps.

### **3. Quantifying care**

To estimate the willingness of a customer to relinquish control of the pool pump to a DR program, it is useful to estimate how much they care about the timing of their pumps. This can be based on features suggesting careful adjustment of the pump’s timer settings. The more attentive a pump owner is, the less likely they are to be suitable for a program whose goal is to garner flexibility for the grid. This is intended as an adjunct to questionnaires, perhaps to indicate which customers need not be included.

We estimate the “care” as a weighted sum of several features. Section 3.1 describes some features that may indicate care, and Section 3.2 considers ways of assigning them weights.

### 3.1. Features

Clearly the customer on the right of Figure 1 puts more effort into selecting the time of operation of the pump than that on the left. This suggests the following list of features that indicate care. The names in brackets are used in figure legends, and the numbers will be explained in Section 3.2. ‘Log’ indicates that the actual feature was the logarithm of the number described.

**Number of times the timer settings change** (‘# timer changes’, 5, log). A change is defined as one timer setting ending and a different timer setting starting within three days. This indicates that effort was put into ensuring specific times of operation, and so indicates care.

**Average number of operations per day** (‘opns per day’, 4). Most pool pump timers can turn the pump on once or twice per day. Turning on twice causes wear on the pump and effort to set the timer, and so choosing that option suggests greater care about the times of operation. For example, the owners may want the pump to operate only outside of certain time periods due to its noise.

**Expected value of setting duration** (‘365 – rms(days)’, 3). The longer the time settings remain constant, the less the customer appears to care about the setting. We consider the question “On a day chosen uniformly at random, what is the expected number of days on either side that the settings remain unchanged?” This is proportional to the mean square length of time between setting changes. To avoid skew, we use the root-mean-square (RMS) for this feature, and subtract it from the total number of days, since a high RMS value suggests low care.

**Max drift** (‘-drift’, 3). Customers who allow the times to drift far from their initial times appear to care less about the precise timing than those who either continually reset the timer or use accurate timers. We use as a feature  $(-|t_{end,r} - t_{start,r}|)$  where  $t_{start,r}$  and  $t_{end,r}$  are the turn-on times at the first and last day of rectangle  $r$ , and the subtraction is modulo 24 hours.

**Different settings on different days of the week** (‘alt days’, 3). This feature is 1 if some days of the week use different timer settings, 0 otherwise.

**Number of “rectangles”** (‘num rectangles’, 3, log). A pump can be on multiple times per day. For each on period, if the setting is constant for a run of days, the result appears as a “rectangle” (see Figure 1). This feature is the number of such rectangles.

**Was there a change of settings without changing the duration?** (‘duration same’, 3). Filtration is often reduced over winter, which doesn’t signify care about the time of operation. However, a change to the settings that does not change the number of hours of operation does indicate care about the specific times of operation.

**Days on** (‘Days used’, 1). Filtration need not occur every day. This feature records the number of days per year that the pump is actually used. A high value suggests that the user cares more about adequate filtration.

**Change in daily number of operations** (‘# opns changes?’, 1). This is 1 if the number of times per day the pump operates changes. As with all changes of pattern, this suggests greater care about when the pump operates.

**Only one setting changes at a time** (‘single changes?’, 1) This is 1 if every time the timer settings change, either the on time or the off time remains unchanged. This suggests that the reason for the change was just to change the duration, rather than the time of day of operation.

Many of these features could be replaced by more informative ones, such as counting the number of “single changes”. However, several of these introduce undesirable correlations with the number of timer changes or the number of rectangles. The above forms have been chosen in an attempt to avoid such problems.

### 3.2. Care index

Given these features, it remains to assign weights to them to form our weighted average “care index”. If ground truth were available, this would simply require linear regression. However it is challenging without ground truth. Our approach here is to consider three very different ways of assigning weights, and study how different the resulting care indices are.

**Posited weights.** The first set of weights  $w$  are chosen manually, based on a heuristic understating of human behaviour. These are derived from the numbers listed after each feature in Section 3.1, by normalizing them to have unit  $l_1$  norm,  $w = \hat{w}/\|\hat{w}\|_1$ . Although we believe these are reasonable, there is no objective justification for them, which is why we compare them with two competing sets.

**PCA.** An independent way of selecting weights is principal component analysis (PCA). The principal component is the eigenvector corresponding to the largest eigenvalue of the covariance matrix of the features. However, PCA summarizes all structure in the data, rather than only how much they care about the timing of the pumps. The next metric will reflect both the domain knowledge of “posit” and objective structural features like “PCA”.

**Corr:** Intuitively, the weight given to each feature should reflect how well it correlates with the true care. The closest we have is the correlation with our best estimate of the care, namely the index using the posited weights. However, this is self-referential in that the weight posited for feature  $i$  affects the correlation. Instead, our third set of weights is calculated as follows. For each feature  $i$ , the care index is calculated with feature  $i$  omitted, giving  $s_i$ . The correlation between feature  $i$  and  $s_i$  is then calculated, and taken as a new (unnormalized) weight. These weights are again normalized to unit  $l_1$  norm.

## 4. Numerical Results

Since the features were chosen to increase as care increases, if the features genuinely point to an underlying phenomenon, we would expect most of the elements of the principal component to have the same sign. As seen in Figure 2, this is the case, which provides a degree of support for the choice of features. The same holds for the correlation-based weights.




Figure 2. The three sets of weights.

To determine which households are most likely to participate in demand response, the ranking induced by the care index is more important than the specific value of care index. Figure 3 compares the ordering assigned to customers by the manual weighting with those assigned by the two systematic sets of weights. This shows that the manually posited weights differ from the other two more the difference between those two. However, the overall agreement between the measures is quite good, with no more than 17% disagreement for any threshold.

This returns us to the question of whether the devices exhibiting a high care index are swimming pools at all. We obtain evidence that they are by the traditional method of using that hypothesis to make a prediction, and testing the prediction. Specifically, if many appliances other than pumps are in the

data, then the variance of the pump powers would be expected to be higher than if all appliances are indeed pumps. Figure 4 shows that the variance is not substantially higher for those with higher care indices. Figure 5 shows the use patterns for increasing care; the highest and lowest are in Figure 1.

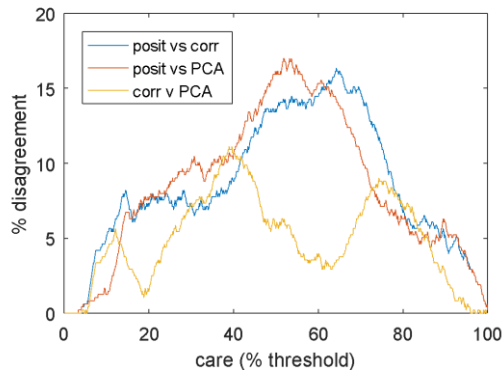


Figure 3. Fraction of customers whose rank is above a threshold by one set of weights and below by another.

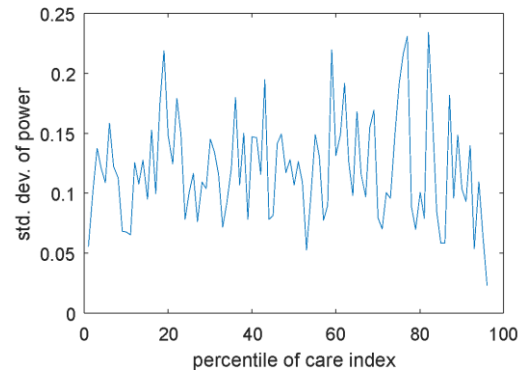


Figure 4. Standard deviations of powers within each percentile of customers




Figure 5. Energy patterns of 20th, 40th, 60th and 80th percentiles of care index. The patterns for households with lowest and highest care are shown in Figure 1.

## 5. Conclusion

Smart meters provide a useful means for preselecting which customers have large appliances such as pool pumps, and identifying which customers seem relatively indifferent to the times of operation of these devices. Such indifferent users are good candidates for targeting with proposals to participate in demand response schemes.

## 6. Acknowledgements

This work was funded in part by a seed grant from the Monash Energy Materials and Systems Institute. The authors thank Robert (Yu) Fu, Xuelin Situ and Zahraa Abdallah for assistance in manually correcting the estimated timer settings, and Abdulahi Khalif for collaboration in preliminary work on this project.

## 7. References

- [1] *Integrated System Plan Consultation*, Australian Energy Market Operator (AEMO), 2017,
- [2] *Clean Energy Australia Report 2018*. Clean Energy Council Australia, 2018.
- [3] V. V. Terzija, "Adaptive under-frequency load shedding based on the magnitude of the disturbance estimation," *IEEE Trans. on Power Systems*, vol. 21, no. 3, pp. 1260-1266, Aug. 2006.

- [4] D. Angeli and P.-A. Kountouriotis, "A stochastic approach to dynamic-demand refrigerator control," *IEEE Trans. Control Systems Technology*, vol. 20, no. 3, pp. 581–592, 2012.
- [5] E. Vrettos, F. Oldewurtel, and G. Andersson, "Robust energy constrained frequency reserves from aggregations of commercial buildings," *IEEE Trans. Power Systems*, vol. 31, no. 6, pp. 4272–4285, 2016.
- [6] O. Malik, P. Havel, "Active demand-side management system to facilitate integration of RES in low-voltage distribution networks" *IEEE Trans. Sustainable Energy*, vol. 5, no. 2, pp. 673–681, 2014.
- [7] M. S. Modarresi, L. Xie, C. Singh, "Reserves from controllable swimming pool pumps: reliability assessment and operational planning," *Hawaii Int. Conf. System Sciences*, 2018.
- [8] S. P. Meyn, P. Barooah, A. Busic, Y. Chen, and J. Ehren, "Ancillary service to the grid using intelligent deferrable loads," *IEEE Trans. Automatic Control*, vol. 60, no. 11, pp. 2847–2862, 2015.
- [9] B. Biegel, L. Hansen, P. Andersen, and J. Stoustrup, "Primary control by ON/OFF demand-side devices," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 2061–2071, Dec. 2013.
- [10] R. Couillet, S. Perlaza, H. Tembine, and M. Debbah, "Electrical vehicles in the smart grid: A mean field game analysis," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 6, pp. 1086–1096, Jul. 2012.
- [11] A. Molina-García, F. Bouffard, D. Kirschen, "Decentralized demand-side contribution to primary frequency control". *IEEE Trans. Power Syst.* vol. 26, no.1, pp. 411–419 Feb. 2011.
- [12] M. Aunedi ; P.A. Kountouriotis ; J. E. O. Calderon ; D. Angeli ; G. Strbac, "Economic and environmental benefits of dynamic demand in providing frequency regulation," *IEEE Trans. Smart Grid*, vol. 4, no. 4, Dec. 2013
- [13] J. Eto, J. Nelson-Hoffman, C. Torres, S. Hirth, R. Yinger, J. Kueck, B. Kirby, C. Bernier, R. Wright, A. Barat, and D. Watson. 2007. *Demand Response Spinning Reserve Demonstration*. LBNL-62761. Lawrence Berkeley National Laboratory, Berkeley, CA.
- [14] C. Vivekanathan, Y. Mishra, G. Ledwich, and F. Li, "Demand response for residential appliances via customer reward scheme," *IEEE Transactions on Smart Grid*, vol. 5, no. 2, 2014.
- [15] Design & Engineering Services Customer Service Business Unit Southern California Edison, *Pool pump demand response potential: Demand and run-time monitored data*, DR 07.01, Jun 2008. [Online]. Available: <https://www.etcc-ca.com/reports/pool-pump-demand-response-potential> (Accessed 27 Dec 2018).
- [16] World Health Organization, *Guidelines for safe recreational water environments*, 2006.