EarlyOff: Using House Cooling Rates To Save Energy

Carl Ellis
Microsoft Research, UK and
Lancaster University, UK
carl.ellis@comp.lancs.ac.uk

James Scott
Microsoft Research, UK
jws@microsoft.com

Mike Hazas
Lancaster University, UK
hazas@comp.lancs.ac.uk

John Krumm
Microsoft Research, US
jckrumm@microsoft.com

Abstract

Home heating systems often have a significant thermal inertia, as homes stay warm after the heating is turned off for significant periods of time. We present the EarlyOff concept, whereby home heating can be predictively turned off in advance of occupants' departure, using this inertia to keep the house warm while saving energy. We use a previously gathered data set of real-time heating, gas, and occupancy readings from five houses and conduct a data-driven analysis of potential energy savings. Using an "oracle" predicting actual departure events, we show an upper bound savings of 4–12% of the gas used over the whole study period by applying EarlyOff. Using a real predictor which makes use of historical occupancy probabilities, we show savings of 1–8% of gas use.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Algorithms, Experimentation, Management

Keywords

Heating, Control, Occupancy, Departure, Predictive

1 Introduction

In many regions of the world, home heating consumes more energy than any other household end-use [1]. While many houses have programmable thermostats which allow users to schedule heating only for occupied periods, studies have found that many are badly programmed or not programmed at all [2]. Of those which are programmed, other studies show that people are poor predictors of their own schedule [3]. As a result, there have been a number of efforts aimed at automatic control of heating systems based on predicting future occupancy, and heating up in time for that occupancy, including methods such as neural nets [4], reactivity based on GPS work-to-home travel times [5], deep setbacks and reactivity to occupancy [6], and occupancy-predictive for per-room control [7].

In this paper we explore a complementary technique which we call *EarlyOff*. We observe that a house often stays warm for significant periods after the heating is turned off, due to thermal capacitance in the heating infrastructure, and the retention of warmth by insulated indoor environments. While prior automatic systems have focused on heating to set point by the time an occupancy period begins, we evaluate the potential savings from deliberately turning the heating off while the house is occupied, in anticipation of departure.

In this feasibility study we show that EarlyOff could ideally save 4–12% of the total heating energy with perfect prediction of departures and of cooling behaviour. We introduce a new departure prediction algorithm that we call "BigDrop" and show achievable savings of 1–8%, though we caution that in some cases these savings may be partially lost in subsequent heat events since the house would start from a cooler temperature.

2 Evaluating the Potential for EarlyOff

We evaluated the potential savings due to EarlyOff using a previously gathered data set for the PreHeat project [7]. The data set spans 5 occupied domestic houses, 2 in Cambridge, UK (UK1, UK2) and 3 in the Seattle metropolitan area, US (US1, US2, US3), and 3 types of heating system (wall-mounted and underfloor convection heating in the UK, and forced air heating in the US), for an average of 61 days per house in January–March 2011. This data set includes the following: heat on/off times, occupancy, and temperature, all at sub-minute intervals. In UK houses gas use at sub-minute intervals was also included, while the US houses included per-day gas use.

In the PreHeat study, the aim was to evaluate how well an occupancy prediction algorithm performed as a heating system controller. The study had two con-

Buildsys'12, November 6, 2012, Toronto, ON, Canada.

ditions: a scheduled condition similar to a 7-day programmable thermostat as a control, and a predictive algorithm which ran on alternating days in the study. For the data-driven analysis required to assess the feasibility of EarlyOff, we used the actual occupancy behavior in the houses to predict departure, and the actual cooling behavior of the houses to extrapolate the temperature behaviour when using EarlyOff. Since only predictive days give cooling events near actual departures, we ignored the scheduled days as cooling events often did not correspond to a departure. This meant that the useful data was between 20-33 days per house, median 27 days, total 132 days.

EarlyOff can also allow new opportunity for energy savings, by allowing the house to cool *just below* the set point in the run-up to departure—right before departure the house's occupants are likely clothed and active, thus requiring less heating for comfort. For the purposes of this research, we chose to allow a space to become up to 1°C colder than the set point in the run up to a departure. 1°C is a widely-used threshold for evaluating heating control systems; it is used for example to calculate MissTime [6].

3 Oracle and BigDrop

To test the potential for EarlyOff, we evaluated two occupancy predictors. The first, "Oracle", simply uses the actual departure events, which is not implementable in a live system without perfect knowledge of the future. However, it allows us to quantify the upper bound for the possible gas savings and estimate how well a highly accurate predictor would perform.

The second algorithm, "BigDrop", predicts departure events based on historical occupancy data, and works as follows. The first week's data is used purely for training. After that, for each day, all the data up to but not including that day was used to calculate a single departure prediction event for that day. We started with a prior algorithm [3], but with 15 minute time quanta to obtain occupancy prediction into the future. We augmented this algorithm in order to transform the occupancy predictions into discrete departure events. We noticed that departure could be predicted by looking for *significant drops in the occupancy probabilities*. For each day, we found the standard deviation of the difference between occupancy probabilities in successive timeslots, and looked for drops in occupancy probability of at least 2.5x this standard deviation. Since there were often multiple significant drops in adjacent 15 minute timeslots, we grouped these and then used the middle time point of any group for the actual departure prediction time. We also limited BigDrop to outputting a single prediction per day by looking for the largest drop in probability of occupancy between 7am and 2pm. This allowed us to reduce false positive predictions which were more likely later in the day.

4 Modeling Cooling Behaviour

We used the following data-driven method to estimate the amount of gas that could be saved due to EarlyOff. For a given departure prediction, if the actual unoccupied period was less than half an hour, we reported zero gas savings since the house would have to start heating again for the future arrival. We used the actual temperature trace of the house after the heating was switched off (i.e. the real cooling behaviour), and extrapolated this backwards in time to simulate if we had turned the heating off early, as in Figure 1. We used this to find the earliest heat-off time that did not occur MissTime. To avoid over-predicting EarlyOff savings, the house must reach the set point at least once during the heating period, so EarlyOff does not affect the initial heating phase but only the "steady state" phase of gas consumed to keep the house at a set point.

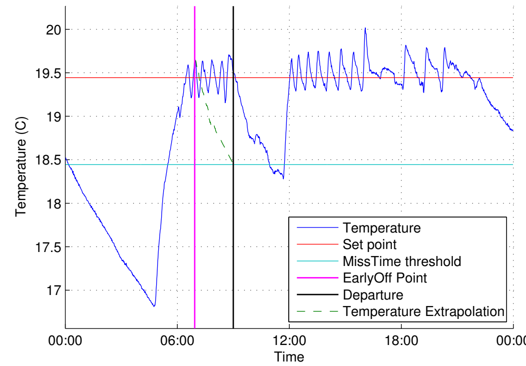

Figure 1: Example of EarlyOff in operation

To calculate gas savings, we summed the gas used between the actual heat off time and the EarlyOff heat off time: this is the amount of gas that would be saved. In US houses, this was done by taking the gas used on that particular day and dividing this by the number of minutes that the furnace was on for that particular day. This is valid as the furnaces in the US houses utilize a fast reacting control system in which they are either full-on or full-off. In UK houses, we had sub-minute-interval gas use data. However, UK houses controlled heating on a per-room basis. We extrapolated the cooling for each room individually, which led to each room having its own EarlyOff duration, but we conservatively used the shortest duration as the whole-house EarlyOff duration for the gas saving estimation.

5 Results

We first examine the Oracle results, to estimate the ideal savings available with EarlyOff. The results are summarized in Table 1 which shows for each house the proportion of days where EarlyOff could be used and the estimated gas savings as a fraction of the total gas used throughout the whole data set period (this includes days where no EarlyOff occurred).

As the table shows, in all houses a significant number of days' behavior could benefit from EarlyOff, with 22–96% of days having a EarlyOff event. We quantified the total gas saved due to EarlyOff, and found that this was from 4% to 12% of the total gas used for the study

House	Oracle			BigDrop						
	EO days	Mean EO duration	Gas saved	EO days with pred.	Predictions					Gas saved
					Correct	Late	Early	False(occ)	False(unocc)	
UK1	22%	179 min	4%	63%	6%	24%	0%	59%	12%	1%
UK2	80%	91 min	5%	75%	40%	13%	0%	13%	33%	1%
US1	52%	95 min	4%	81%	15%	22%	15%	48%	0%	2%
US2	96%	60 min	12%	93%	44%	32%	8%	16%	0%	8%
US3	56%	73 min	6%	92%	22%	30%	9%	35%	4%	3%

Table 1: EarlyOff (EO) performance, including gas savings as a percentage of total gas used on all days

period. However, we caution that some of these savings would be lost in reality, since if EarlyOff had triggered cooling early, then the subsequent heating period would have to start from a slightly colder temperature, so it would consume slightly more gas (though never as much as was saved). The amount of gas savings that are subsequently lost is not easy to estimate; this would involve predicting the temperature over the next heating period. However, the potential savings we have found is nonetheless a very encouraging result.

5.1 BigDrop performance

Next, we evaluated BigDrop's performance in predicting departure events and how these translate into gas savings. Table 1 summarises how often predictions occur (63–93% of days) and categorises them according to their relationship to ground truth. The *Correct* predictions were ± 15 minutes from the actual departure time (i.e. as close as our quantisation of time permitted). *Late* predictions were later than an actual departure, but within the EarlyOff duration that our extrapolation found for the departure (i.e. some gas would have been saved). *Early* predictions were up to 30 minutes before the actual departure - these would save gas but would entail a short MissTime. False predictions fell in one of two categories - occupied and unoccupied. When the house was occupied, these would incur MissTime since the house would cool while occupied. We envisage a system where if the departure was not soon after the predicted time, the heating would of course resume. No gas savings are calculated for these periods. EarlyOff predictions during unoccupied periods are simply discarded, since they have no ill effect.

The Correct/Early/Late predictions represent gas saving opportunities, and together account for between 30% to 84% of predictions. BigDrop performs better in some houses than in others, with UK1 being hardest to predict and US2 having occupancy most predictable by this algorithm. The gas savings achieved by the BigDrop algorithm are 1–8%; the algorithm manages between 25% and 66% of the ideal gas savings from the Oracle. It is encouraging that BigDrop is capable of achieving a reasonable fraction of the ideal savings, though future work may further improve on this.

6 Conclusions

We have described EarlyOff, a new method of saving domestic heating energy by turning off the heating early, taking advantage of the fact that a domicile stays warm for some time after the heating turns off. This is complementary to existing methods in the literature, which focus on the start of the heating period. We conducted a data-driven analysis of EarlyOff's potential energy savings in five houses. In these houses, we show that EarlyOff could save between 4% and 12% of the total heating energy with ideal prediction of departures and of cooling behaviour. With our "BigDrop" departure algorithm, this drops to a 1–8% saving. BigDrop achieves between 25% and 66% of the ideal gas savings in the houses studied.

We caution that some of these savings will be lost when the house next heats due to the longer cooling period. The cooling behavior we used for this evaluation is based on actual cooling behavior, and for a real world implementation to be achieved this must also be predicted. Clearly, another element of future work is to mitigate the MissTime caused by false departure predictions. One approach would be to display the predicted departure time in the house, and provide occupants with an easy way to indicate if this was incorrect. Nonetheless, we have shown that EarlyOff can increase energy savings for domestic heating, which is a major contributor to energy use.

Acknowledgement. Lancaster University was supported by EPSRC grants EP/I00033X/1 and EP/G008523/1.

7 References

[1] L. Pérez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394–398, 2008.

[2] A. Meier, C. Aragon, T. Peffer, D. Perry, and M. Pritoni. Usability of residential thermostats: Preliminary investigations. *Building and Environment*, 46(10):1891–1898, 2011.

[3] J. Krumm and A.J. Brush. Learning time-based presence probabilities. In *Proc. Pervasive*, pages 79–96. Springer, 2011.

[4] M. C. Mozer, L. Vidmar, and R. H. Dodier. The neurothermostat: Predictive optimal control of residential heating systems. *Advances in Neural Information Processing Systems*, 9:953–959, 1997.

[5] M. Gupta, S. Intille, and K. Larson. Adding gps-control to traditional thermostats: An exploration of potential energy savings and design challenges. In *Proc. Pervasive*. Springer, 2009.

[6] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat : Using occupancy sensors to save energy in homes. In *Proc. SenSys*, pages 211–224. ACM, 2010.

[7] J. Scott, A.J. Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. Preheat: controlling home heating using occupancy prediction. In *Proc. UbiComp*. ACM, 2011.