Smart heating control with occupancy prediction: How much can one save?

Wilhelm Kleiminger
Institute for Pervasive Computing
ETH Zurich
8092 Zurich, Switzerland
wilhelmk@ethz.ch

Silvia Santini
Wireless Sensor Networks Lab
TU Darmstadt
64283 Darmstadt, Germany
santinis@wsn.tu-darmstadt.de

Friedemann Mattern
Institute for Pervasive Computing
ETH Zurich
8092 Zurich, Switzerland
mattern@ethz.ch

UbiComp '14, September 13 – 17 2014, Seattle, WA, USA.

ACM 978-1-4503-3047-3/14/09$15.00.
http://dx.doi.org/10.1145/2638728.2641555

Abstract
Research results on smart heating systems based on occupancy prediction are often difficult to reproduce and to compare. Evaluating the performance of these systems through simulation or real experiments requires defining suitable scenarios and setting a large number of parameters. As different authors rely on different scenarios and parameter settings, comparing the reported performance results is often infeasible. In this paper, we argue that overcoming this problem is crucial to bring research on smart heating systems a step forward. We outline the main factors influencing the performance of such systems and we show how these factors can be integrated by proposing a simple yet thorough evaluation methodology for smart heating systems. Using parameters synthesised from real-world occupancy and weather data, we describe how this methodology can be used to establish performance bounds of smart heating systems.

Introduction
Interest in smart heating systems has recently received much attention in both industry and academia [3, 10, 11, 13, 14]. Traditional thermostats can be made smart by adding the capability to automatically control the temperature setpoint depending on the current and predicted presence of household occupants. When the occupants leave the building and the household becomes

unoccupied, the system lowers the temperature to save energy. In order to instruct the thermostat to heat up the building in time, occupancy prediction algorithms are used to estimate the arrival times of the occupants. A large number of such occupancy prediction algorithms have been proposed in the literature. Some relate the arrival time of the occupants to their departure time or exploit weekday similarities [11, 13]. Others use past mobility traces to predict when the home may be reached based on the current location of the occupants [3, 10]. Existing results presume smart heating control systems that rely on occupancy prediction algorithms can achieve energy savings of 8% or even 28% with respect to traditional, timer-based heating control systems [11, 13]. However, the use of different evaluation scenarios and parameters makes it difficult to compare results obtained by different authors. Also, a thorough description of the evaluation setup of an approach is often too verbose to be included in a research publication. The lack of such a description, however, makes it infeasible for other authors to reproduce previously achieved results.

In this paper, we argue that enabling the comparability and reproducibility of research results on smart heating control is crucial to build and improve upon state-of-the-art approaches. In the following, we outline a generic methodology to evaluate the performances of smart heating control systems. We illustrate how this methodology can be applied in practice by determining the potential energy savings of a smart heating system for two fictitious households in the Lausanne area in Switzerland.

Methodology

The amount of energy spent for heating a building is determined by several factors, including the level of insulation (transmission losses), the building's orientation and number of windows (solar gains), its exposure and tightness (ventilation losses), heat gains due to occupants and appliances (internal gains) and of course the outside temperature. Whenever a household is occupied, the heating system is required to maintain its indoor temperature at a pre-specified *comfort* level (e.g. $20\,^{\circ}\mathrm{C}$) and match the sum of the transmission and ventilation losses minus the sum of solar and internal gains. A smart heating system can save energy by letting the indoor temperature of the household drop to a *setback* level while the building is not occupied. The actual savings achievable using this approach are mainly influenced by four factors: (1) the actual *occupancy* of the building; (2) the *weather conditions*; (3) the *characteristics of the building* and (4) the *control strategy* used to set the target temperature of the thermostat.

Occupancy

To assess the performance of smart heating systems based on occupancy prediction in a realistic manner, the availability of actual occupancy data is critical. As a practical example, we refer in this paper to actual occupancy data extracted from the Lausanne Data Collection Campaign (LDCC) data set [6]. In a previous study [7] we have shown how to derive *occupancy schedules* from this data set. An occupancy schedule is a binary vector containing ones and zeros denoting occupied or unoccupied time intervals of a household.

Weather

Weather conditions significantly influence the behaviour of any heating system. When the temperature outside the building is lower than the indoor temperature, the building loses heat to the environment. The higher the difference between the indoor and outdoor temperature, the more heat is lost and must be compensated for by the heating system. A larger difference between indoor and outdoor temperature also implies that the *preheat time* (i.e. the time needed to heat the house to a desired *comfort* temperature starting from a lower *setback* temperature) increases. However, this temperature difference is not the only factor influencing the amount of energy that is spent by the heating system. The heat transferred to the building through solar radiation, the so-called *solar gains*, must also be taken into account. Such is the importance of solar gains that the field of passive solar building design focuses on utilising solar gains in winter and avoiding those gains in summer to optimise heating and air-conditioning [12]. Considering adequate *weather scenarios* that allow to model the effect of actual outside temperatures and solar radiation is thus essential to evaluate the performance of smart heating systems.

Why using weather scenarios?
A possible solution to include realistic weather scenarios in a simulation consists of using traces of actual weather data. Ideally, this data should be retrieved for the same location and period for which actual occupancy data of a target building is available. However, this approach is often infeasible for two reasons. First, this would prevent us from using any occupancy data collected in summer which severely reduces the available occupancy data.[1] Second, two identical occupancy schedules on different days can result in fundamentally different heating costs if their associated weather conditions differ. Thus, as savings are somewhat "left to chance", we cannot draw conclusions about possible savings *in general*. We must therefore decouple the weather data from the occupancy data. However, simulating the days in the occupancy dataset with all possible combination of days in the weather data is too expensive computationally. Thus, rather than picking a few weeks in the winter months to fit the occupancy schedules, we use historical weather data to derive *characteristic* weather scenarios. This allows us to combine multiple days with the same weather characteristics to generate a smaller number of discrete, characteristic weather scenarios. In the remainder of this section we show an exemplary calculation of the characteristic weather scenarios for Pully, Switzerland[2].

Temperature thresholding

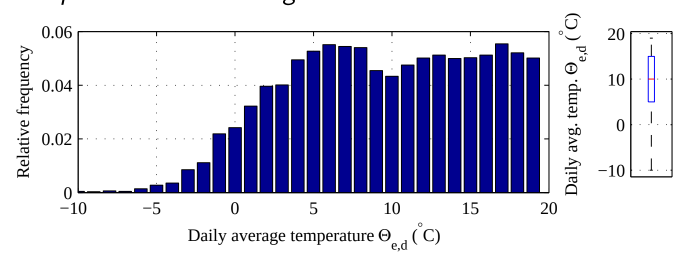

Figure 1: Distribution of daily average temperature $\Theta_{e,d}$ for $\Theta_{e,d} \leq 20\,^{\circ}\mathrm{C}$ in Pully, Switzerland (Jan 1994 to Jan 2014). Right subfigure shows median and quartiles.

Figure 1 illustrates the distribution of the daily average temperatures in Pully over the last 20 years[3]. It shows the relative frequency (i.e. the empirical probability) of observing a particular average temperature on any given day. To generate the weather scenarios, we only consider days with an average outside temperature $\Theta_{e,d}$ less than $20\,^{\circ}\mathrm{C}$, as no heating is necessary if the outside temperature reaches the comfort temperature.

[1] Note, that working with characteristic weather scenarios we implicitly assume that no correlation between the occupancy schedule of a household and the current weather conditions exists.

[2] We chose the weather station at Pully near Lausanne in Switzerland to match the occupancy data from the LDCC dataset.

[3] The weather data has been obtained from MeteoSwiss.

The right part of Figure 1 shows that the median temperature in the observed sample is $10\,^{\circ}\mathrm{C}$ (i.e. 50% of the days requiring heating have temperatures below $10\,^{\circ}\mathrm{C}$), while the lower quartile is $5\,^{\circ}\mathrm{C}$ (i.e. 25% of the days requiring heating have temperatures below $5\,^{\circ}\mathrm{C}$). The minimum value observed is $-10\,^{\circ}\mathrm{C}$.

Table 1: The four temperature scenarios.

Scenario	Range
Very low temperature	$-6\,^{\circ}\mathrm{C} \leq \Theta_{e,d} \leq -4\,^{\circ}\mathrm{C}$
Freezing temperature	$-1\,^{\circ}\mathrm{C} \leq \Theta_{e,d} \leq 1\,^{\circ}\mathrm{C}$
Low temperature	$4\,^{\circ}\mathrm{C} \leq \Theta_{e,d} \leq 6\,^{\circ}\mathrm{C}$
Moderate temperature	$9\,^{\circ}\mathrm{C} \leq \Theta_{e,d} \leq 11\,^{\circ}\mathrm{C}$

An important parameter for the design of heating systems is the norm outside temperature. The norm outside temperature is the lowest two-day average temperature which was measured at least 10 times over a period of 20 years. It is used to determine the worst-case temperature scenario for the dimensioning of the heating system [2]. For Pully, we determined the norm outside temperature as $-6\,^{\circ}\mathrm{C}$. Using a range of $\pm 1\,^{\circ}\mathrm{C}$ around the norm outside temperature ($-6\,^{\circ}\mathrm{C}$), the freezing point ($0\,^{\circ}\mathrm{C}$), the lower quartile ($5\,^{\circ}\mathrm{C}$) and the median ($10\,^{\circ}\mathrm{C}$) temperatures, we defined four characteristic temperature scenarios for Pully shown in Table 1. For further details on the definition of the weather scenarios, the interested reader is referred to [9].

Table 2: All 8 weather scenarios. For each of the 8 scenarios, the table shows the daily average temperature $\Theta_{e,d}$ and the daily average of the global radiation I_{avg} for reference.

	$\Theta_{e,d}$ ($^{\circ}\mathrm{C}$)		I_{avg} (W/m^2)	
Scenario	**clear**	**cloudy**	**clear**	**cloudy**
Very low temperature	-5.4	-4.7	142.9	35.5
Freezing temperature	0.1	0.0	137.5	30.2
Low temperature	5.1	5.1	148.5	26.1
Moderate temperature	10.1	10.0	180.7	29.7

Note, that by choosing the bounds for the *very low temperature* scenario to be above rather than around the norm outside temperature, we seek to avoid a rare scenario where our modelled heating system cannot fully heat up the building during parts of the day.

Solar gain thresholding

The amount of solar gains is dependent on a number of factors including the cloud cover, the location of the building, the number, type and size of windows as well as their orientation and shading such as curtains. The incident solar radiation causing the internal gains may be divided into two categories: *Direct* and *diffuse* radiation. Direct radiation is caused by line-of-sight rays from the sun, while diffuse radiation is light reflected from the surroundings. Here we focus on the direct radiation. Many weather stations provide only the *global* (direct $+$ diffuse) radiation on a horizontal surface. To calculate the direct solar gain through the windows at a given instant of time, we first obtain the position of the sun relative to our location. We then partition the global radiation into direct and diffuse radiation using the Reindl* method [4]. Finally, we transform the direct radiation from the horizontal to a vertical plane and thence to the different orientations of the windows. We define days with an average global radiation above 100 W/m^2 as *clear*. Days with an average global radiation below 50 W/m^2 are considered *cloudy*.

Weather scenarios

Table 2 shows statistics of the 8 weather scenarios. Due to limited availability of the global radiation data, the scenarios are built using data from January 1, 2005 to January 1, 2014. Figure 2 provides an example of the resulting data for the freezing temperature, clear sky scenario. $I_{dir,East}$, $I_{dir,South}$ and $I_{dir,West}$ denote the direct radiation on the east, south and west walls in

W/m^2, respectively. Θ_e is the outside temperature in $^\circ\mathrm{C}$. The figure shows that while the south wall experiences the highest influx of radiation around noon, the highest cumulative gains occur just before 3 p.m. when the sun is in a south-west position. The outside temperature is reaching its maximum value shortly after 4 p.m..

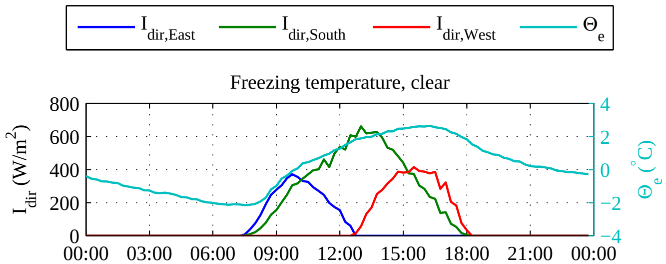

Figure 2: A weather scenario: Freezing temperature, clear sky.

Building characteristics

At the beginning of this paper we have highlighted that the energy spent to heat a building is determined by the building's transmission losses, solar gains, ventilation losses and internal gains. These heat losses and gains link the ability of the building to store and retain energy to the weather conditions discussed in the previous section. In the following we will discuss the factors needed to sufficiently describe the building characteristics using two fictitious but typical and rather different properties.

Building geometry

The two building configurations used in this study are a studio flat and a house[4]. For both we assumed building parts with low U-values (good insulation), following recent legislatorial guidelines [1]. The transmission losses can then be derived from the wall and ceiling area together with the appropriate U-values. The studio flat has an area of 52 m^2. The house has an area of 176 m^2. All windows are sized 2 m^2. The height of the rooms is 2.5 m in both cases. The doors are 2.8 m^2. The flat has one window facing east, and 3 windows facing south. The house has two east-facing windows, four windows on the south side, two to the west and two windows facing north.

Ventilation losses

Buildings need to be ventilated frequently for hygienic purposes. In addition, they lose heat through drought (caused by cracks and small openings in the building envelope). These losses must be matched by the heating system. We have calculated the ventilation losses according to the simplified method of the DIN EN 12831 standard which uses the maximum value of the hygienic and natural ventilation to denote the ventilation losses [2].

Internal gains

When humans are present in the building, additional heat is generated through the operation of appliances (e.g. television set, dryer and stove) and the *metabolic heat* from the occupants themselves. An average person produces 125 W of heat. We model the flat to be occupied by 2 people and the house by 3 people. Since most occupancy prediction algorithms algorithms consider binary occupancy, we assume all occupants are present whenever the building is occupied (i.e. 250 W and 375 W whenever the flat or house are occupied, respectively).

Dimensioning the heating system

In order to appropriately dimension the heating infrastructure (e.g. radiators and boilers) in a building, the DIN EN 12831 standard allows for the calculation of the design heating load [2]. The design heating load is the amount of heat that needs to be supplied to a building to

[4]Further details regarding the building configurations can be obtained from [9].

keep $\Theta_{comfort}$ even when the outside temperature is very low. We thus make sure to neither over- nor underestimate the performance of the heating system.

Simulation model

We use the 5 resistance 1 capacitance (5R1C) model from the ISO 13790 standard to simulate the energy consumption of the heating system. In the 5R1C model, the transient thermal conduction between the building and its surroundings is modelled analogously to an RC circuit. The ISO 13790 standard [5] was mandated by the EU Directive 2002/91/EC on the energy performance of buildings (EPBD) which required a "a common methodology for calculating the integrated energy performance of buildings". Today, this model has been widely adopted for building simulations in Europe. We have adapted the 5R1C model to calculate the indoor air temperature and heat input at intervals of 15 minutes.

Table 3: Controller Parameters

Symbol	Units	Description
$\Theta_{comfort}$	$^\circ\mathrm{C}$	Comfort / set-point temperature ($20\,^\circ\mathrm{C}$)
$\Theta_{setback}$	$^\circ\mathrm{C}$	Setback temperature ($10\,^\circ\mathrm{C}$)
t	/	Current 15-minute time interval
S	$[\{1,0\}]$	Actual occupancy schedule
P_t	$[[\{1,0\}]]$	Predicted occupancy schedule at interval t

Predictive controller

In order to act upon the predictions made by occupancy prediction algorithms, we must translate their predicted binary occupancy schedules (i.e. predictions about the future occupancy of the building) into an actual heating schedule containing setpoint temperatures for each time of the day. Our suggested approach does so by simulating the response of the indoor temperature if we forego heating at the current (unoccupied) interval. If the system is still able to heat up in the remaining intervals until the (predicted) arrival of the occupants, we decide to forgo heating at this interval. Table 3 shows the parameters used in the controller.

Algorithm 1 Control algorithm

procedure CONTROLLER
    $t \leftarrow$ Current time interval
    $S \leftarrow$ Actual occupancy schedule
    $P_t \leftarrow$ Predicted occupancy schedule at interval t
    Reactive policy:
    if $S_t = 1$ **then** */* Occupied */*
        $\Theta_{set} \leftarrow \Theta_{comfort}$
    else
        Predictive policy:
        $n_{horizon} \leftarrow nextOccupied(P_t)$
        $\Theta_{air,noheat} \leftarrow iso13790_{5R1C}(\Theta_{setback}, \Theta_{m,t-1}, \dots)$
        $n_{preheat} \leftarrow$ Preheat time from $\Theta_{air,noheat}$ to $\Theta_{comfort}$
        if $n_{horizon} \geq n_{preheat}$ **then**
            $\Theta_{set} \leftarrow \Theta_{setback}$
        else
            $\Theta_{set} \leftarrow \Theta_{comfort}$
    $\Theta_{m,t}, \Phi_t, \Theta_{air,t} \leftarrow iso13790_{5R1C}(\Theta_{set}, \Theta_{m,t-1}, \dots)$.
    $t \leftarrow t + 1$
    goto *Reactive policy*.

Algorithm 1 shows the predictive controller used to alternate between the setpoint $\Theta_{comfort}$ and setback $\Theta_{setback}$ temperatures. For each 15-minute time interval t, the controller looks at the current occupancy S_t of the household as given by the occupancy schedule S at time t. If the household is currently occupied, we must keep the setpoint temperature and therefore set $\Theta_{set} = \Theta_{comfort}$. If the household is not occupied at time t, we use the predictive policy. The predictive policy first finds the number of unoccupied intervals until the next occupied interval $n_{horizon}$. This is the maximum time available to heat the building to $\Theta_{comfort}$. The next step is to apply the ISO 13790 model of the building and to compute the indoor air temperature at the next time interval if we only applied enough heat to keep $\Theta_{setback}$ (i.e. we forgo heating) at the current interval. Based on this temperature $\Theta_{air,noheat}$, we compute the number of intervals needed to heat to the comfort temperature – $n_{preheat}$. If $n_{preheat} \leq n_{horizon}$, the remaining time is not sufficient to heat up the building on time and we must heat at the current time t. Using this controller, a simple reactive policy (REA) is obtained if $\forall t_1, t_2 : P_{t_1,t_2} = 0$. Similarly, an always-on policy (AO) is obtained if $\forall t : S_t = 1$. We call an optimal predictive controller with perfect knowledge of the future occupancy OPT.

Energy savings: Simulation results

In the remainder of this paper we will show simulation results for our two fictitious example dwellings in Lausanne, Switzerland. We report the results of a reactive policy (REA) and optimal policy (OPT) as defined in the previous section. REA and OPT define the boundary cases for any occupancy prediction algorithm. REA achieves the highest savings as it waits with increasing the setpoint temperature until the building is occupied. However, due to the time lag between changing the setpoint and the building reaching the new temperature, REA also incurs a large loss of comfort for the occupants. The optimal controller OPT, on the other hand, starts heating at exactly the right moment for the building to reach $\Theta_{comfort}$ upon the arrival of the occupants. We compare the two algorithms using a third scheme: *always-on* (AO). AO simply keeps $\Theta_{comfort}$ throughout the day, without resorting to a setback value. While in reality, a statically scheduled setback regimen is in place in many households (especially during the night), using AO as the baseline has several advantages. First, it rather under- than overestimates possible savings (reducing the total daily consumption by using a low setback temperature during the night increases the percentage savings for the whole day). Secondly, and more importantly, any static setback regimen is bound to be arbitrary in terms of length and setback value. Results with different static regimens are thus hardly comparable.

In order to evaluate the performance of predictive heating controllers, we use the concept of *efficiency gain*. Efficiency gain is the percentage savings of a predictive controller using a specific prediction algorithm over the AO controller. More specifically, it is expressed as $\frac{Q_{AO}-Q_{pred}}{Q_{AO}}$, where Q_{AO} denotes the energy spent by the AO controller and Q_{pred} the energy spent by the predictive controller, respectively. We have computed the annualised efficiency gain by weighting our weather scenarios (cf. Table 2) according to the temperature distribution derived previously (cf. Figure 1).

Table 4: Annual efficiency gains.

	Efficiency gain (%)			
Control strategy	OPT		REA	
Building scenario	clear	cloudy	clear	cloudy
Flat	10	10	11	12
House	8	8	9	11

Table 4 presents the annualised savings obtainable from using the OPT and REA controllers. The results show that for our two model buildings an efficiency gain without comfort loss between 8% and 10% may be obtained with optimal occupancy prediction (averaged over all occupancy schedules derived in [7]). A higher efficiency gain necessitates a loss of comfort as shown by the REA controller. Its efficiency gain is between 9% and 12%.

Conclusions and outlook

In this paper we have shown a methodology to calculate the annual savings for a smart heating system. The results for our exemplary scenario in Lausanne, Switzerland show that, with occupancy prediction, the energy spent may be reduced by an average of 9%. This result is based on a thorough simulation using the ISO 13790 standard and real weather data. Our predictive controller can be easily extended to support further occupancy prediction algorithms. In a forthcoming paper [8], we will report on achievable energy savings with different occupancy prediction algorithms and analyse the savings potential for other building types and climate zones.

Acknowledgements This work has been partially supported by the Collaborative Research Center 1053 funded by the German Research Foundation and by the LOEWE Priority Program Cocoon funded by the LOEWE research initiative of the state of Hesse, Germany.

References

[1] Verordnung über energiesparenden Wärmeschutz und energiesparende Anlagentechnik bei Gebäuden in der Fassung der Bekanntmachung vom 24. Juli 2007 (BGBl I S. 1519), zuletzt geändert am 18. November 2013 (BGBl I S. 3951–3990).

[2] DIN. Heating systems in buildings - method for calculation of the design heat load. DIN EN 12831-03:2008, Deutsches Institut für Normung, Berlin, Germany, 2008.

[3] Gupta, M., Intille, S. S., and Larson, K. Adding GPS-control to traditional thermostats: An exploration of potential energy savings and design challenges. In *Proc. Pervasive'09, Nara, Japan*, Springer (May 2009), 1–18.

[4] Helbig, N. *Application of the radiosity approach to the radiation balance in complex terrain*. PhD thesis, University of Zurich, Zurich, Switzerland, 2009.

[5] ISO. Energy performance of buildings - calculation of energy use for space heating and cooling. ISO 13790-1:2008, International Organization for Standardization, Geneva, Switzerland, 2008.

[6] Kiukkonen, N., Blom, J., Dousse, O., Gatica-Perez, D., and Laurila, J. Towards rich mobile phone datasets: Lausanne data collection campaign. In *Proc. ICPS'10, Berlin, Germany*, ACM (July 2010).

[7] Kleiminger, W., Beckel, C., Dey, A., and Santini, S. Using unlabeled Wi-Fi scan data to discover occupancy patterns of private households. In *Proc. SenSys'13, Rome, Italy*, ACM (2013).

[8] Kleiminger, W., Mattern, F., and Santini, S. Predicting household occupancy for smart heating control: A comparative performance analysis of state-of-the-art approaches. Submitted for publication, July 2014.

[9] Kleiminger, W., Mattern, F., and Santini, S. Simulating the energy savings potential in domestic heating scenarios in Switzerland. Tech. rep., ETH Zurich, Department of Computer Science, July 2014.

[10] Krumm, J., and Brush, A. J. Learning time-based presence probabilities. In *Proc. Pervasive'11, San Francisco, CA, USA*, IEEE (June 2011), 79–96.

[11] Lu, J., Sookoor, T., Srinivasan, V., Gao, G., Holben, B., Stankovic, J., Field, E., and Whitehouse, K. The Smart Thermostat: Using occupancy sensors to save energy in homes. In *Proc. SenSys'10, Zurich, Switzerland*, ACM (November 2010), 211–224.

[12] Norton, B. *Harnessing Solar Heat*. Springer, 2014.

[13] Scott, J., Brush, A. J., Krumm, J., Meyers, B., Hazas, M., Hodges, S., and Villar, N. Preheat: Controlling home heating using occupancy prediction. In *Proc. UbiComp'11, Beijing, PRC*, ACM (2011), 281–290.

[14] Winkler, R., and Wakabayashi, D. Google to buy Nest Labs for $3.2 billion. In *The Wall Street Journal* (Jan 13, 2014).