PreHeat: Controlling Home Heating Using Occupancy Prediction

James Scott[1], A.J. Bernheim Brush[2], John Krumm[2], Brian Meyers[2]
Mike Hazas[3], Steve Hodges[1], Nicolas Villar[1]

[1]Microsoft Research Cambridge
Cambridge, UK
{jws|shodges|nvillar}@microsoft.com

[2]Microsoft Research Redmond
Redmond, WA USA
{ajbrush|jckrumm|brianme}@microsoft.com

[3]Lancaster University
Lancaster, UK
hazas@comp.lancs.ac.uk

ABSTRACT
Home heating is a major factor in worldwide energy use. Our system, PreHeat, aims to more efficiently heat homes by using occupancy sensing and occupancy prediction to automatically control home heating. We deployed PreHeat in five homes, three in the US and two in the UK. In UK homes, we controlled heating on a per-room basis to enable further energy savings. We compared PreHeat's prediction algorithm with a static program over an average 61 days per house, alternating days between these conditions, and measuring actual gas consumption and occupancy. In UK homes PreHeat both saved gas and reduced MissTime (the time that the house was occupied but not warm). In US homes, PreHeat decreased MissTime by a factor of 6-12, while consuming a similar amount of gas. In summary, PreHeat enables more efficient heating while removing the need for users to program thermostat schedules.

Author Keywords
Energy, environment, home heating, sensing, prediction.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Algorithms, Experimentation, Human Factors.

INTRODUCTION
Home heating uses more energy than any other residential energy expenditure including air conditioning, water heating, and appliances [1]. This makes increasing the efficiency of home heating an important goal for saving money and reducing our ecological footprint. Although programmable thermostats provide the technology to reduce this problem, they are underutilized. Surveys have found that fewer than 50% of US households have programmable thermostats, and even worse, the US Environmental Protection Agency estimates that 30% or more of US households with programmable thermostats are not using their thermostat's programming feature, so they are not saving the 10%-30% claimed for such devices [2, 6].

Fundamentally, home heating is a trade-off between energy use and warmth. By leaving their thermostat set permanently to a warm temperature, households incur increased energy use costs; by using a programmed thermostat to only heat for some of the time, households can use less energy, but occupants may be cold if the program is wrong or while waiting for the house to heat.

Our home heating system, PreHeat, aims both to eliminate the need for manual user programming of a thermostat schedule and to improve the efficiency of home heating (i.e., improving the tradeoff achieved between energy use and time in which occupants are cold). PreHeat uses occupancy sensing and historical occupancy data to estimate the probability of future occupancy, allowing the home to be heated only when necessary.

We experimentally evaluated PreHeat in five family homes. In three US homes, we controlled whole-house forced air heating systems and used active RFID for occupancy sensing. In two UK homes, we controlled radiators and underfloor heating per room, and we sensed occupancy using motion sensors. During the study, we alternated days between using PreHeat and using a schedule, with an average 61 study days per house. We found that PreHeat did indeed achieve a better tradeoff between gas consumption and MissTime [8] – the amount of time someone was home and the house was not warm. US homes used about the same amount of gas, but had a 6x-12x reduction in MissTime. In the UK houses (with per-room control), improvements in both metrics were achieved.

We acknowledge that there are many complementary ways to save energy in home heating, including better insulation, new architectural standards, and persuasive/informative approaches. Our contribution focuses on better heating scheduling. By deploying our system into real homes, and by directly comparing it to the previous control systems, we provide evidence that our novel predictive heating algorithm can improve on the tradeoff between energy consumption and MissTime, and can do so without requiring homeowners to program complex occupancy schedules.

UbiComp'11, September 17–21, 2011, Beijing, China.

BACKGROUND AND RELATED WORK

In this paper, we focus on the control of two types of *central heating* systems common in Europe and North America. The first type is *hot water space heating*, in which a boiler heats water to 60°C or more and circulates it around the house. Heat in each room is distributed from floor-embedded pipes (*underfloor heating*) or by wall-mounted *radiators*. The second type, more common in North America, is *forced air heating*, in which a furnace heats air and circulates it through ductwork to rooms.

For both types of systems, the temperature is often controlled using a single thermostat located in a central area of the home such as a living room or hallway. The thermostat takes temperature readings and switches the central heating on as needed to maintain a user-defined *setpoint* temperature. Systems with radiators can also be fitted with a thermostatic valve at each radiator, which can be manually set from "1" to "5" to define a separate target temperature for each room. In some houses, particularly those with underfloor heating and some forced-air systems, the system control is split into *zones*, each of which has its own thermostat.

Modern thermostats are typically programmable, allowing the user to specify a *schedule* of times during the day or week when the home should be heated to the setpoint. Outside of these times, the thermostat may use a lower *setback* temperature to reduce the heating energy required, or, as is common in the UK, simply be turned off.

More sophisticated thermostats exist: some are controllable via the Internet; some measure the heat-up latency of the house and compensate the schedule; and others are reactive to occupancy as sensed by entryway or motion sensors. These are expensive, not commonly installed, and as others have reported [3, 8], the latter two tend to heat unnecessarily, actually increasing the energy required.

Researchers have studied the use of domestic heating systems for some time. In 1978, Sonderegger observed that the total energy required by a home's heating system is heavily determined by the particular people living there, dwarfing the effects of insulation or heating infrastructure [9]. Other work has aimed to reduce heating energy using feedback, e.g. via increased frequency and detail of utility bills [10], or real-time energy displays [4]. Feedback and more direct interventions such as "persuasive" interfaces are complementary to our work: while PreHeat may lower energy use by heating less during unoccupied times, a persuasive display may "nudge" an occupant to lower the setpoint and wear a sweater.

Closest to our work are approaches that have aimed to improve control beyond what off-the-shelf products currently offer. Mozer et al. describe a "Neurothermostat" which utilizes a hybrid occupancy predictor, making use of an available daily schedule and a neural network which was trained on five consecutive months of occupancy data [7]. Although real occupancy data from one house is used in the analysis, a first order approximation of the heating system and house is used to model temperature and energy consumption. Mozer et al. show that the Neurothermostat results in a lower unified cost, where energy and occupant "comfort" are expressed as a combined figure, in dollars. We believe that the equivalency of comfort and energy is problematic (the relationship between setpoint deviation and dollars is nonlinear, depends upon the individual, and changes over time), so we prefer to characterize separately the deviation from setpoint temperature at occupied times, and the energy required.

Gupta et al. propose using live data from mobile phones or in-vehicle GPS devices to control home heating and cooling [3]. Their method works by ensuring the home can always be brought to the setpoint in the time it would take the person to travel home from the current location. Using a simulation based on look-up tables extrapolated from three days of temperature readings, they compute the savings possible in four houses based on two months of GPS data. While their scheme does ensure the house is always at setpoint upon arrival, it typically results in lower savings than programmable or manual thermostats.

Lu et al. formulate a hidden Markov model to predict occupancy and control HVAC systems [8]. They collected occupancy data in eight US households for one to two weeks. Using leave-one-out cross-validation to train and test the HMM, they evaluate their approach's MissTime (i.e. total occupied time not at setpoint) and energy savings for each day in a week using the US Dept. of Energy's EnergyPlus simulator. Their forward-looking approach is designed to control specialized two-stage HVAC systems and employs a second "deep" setback (10°C/50°F). By contrast, we show how our system performs in real households using single setback temperatures and single-stage, gas-fired equipment commonly deployed today.

Our occupancy prediction algorithm is itself an improvement over the above approaches: it results in more favorable trade-offs between miss time and energy than the GPS-reactive system [3]. Krumm and Brush [5] also presented an occupancy prediction algorithm that gives probabilities of occupancy at different times of day. However, this algorithm computes a representative Sunday, Monday, etc. for each day of the week, without being able to respond to changing occupancy patterns as PreHeat does.

More broadly, and unlike any of the above work, we evaluate the performance of PreHeat in situ in five houses, using custom embedded sensing and real-time control of the central heating. Previous evaluations have relied upon simulations [3, 7, 8]. While industry standard simulators such as EnergyPlus may go beyond simple first-order approximations, we question the efficacy of simulators for characterizing either (1) the daily energy requirement of heating infrastructure – savings of smaller than 10% can be important, so daily errors of even 3 kWh can muddy analysis; or (2) the deviation from setpoint at occupied

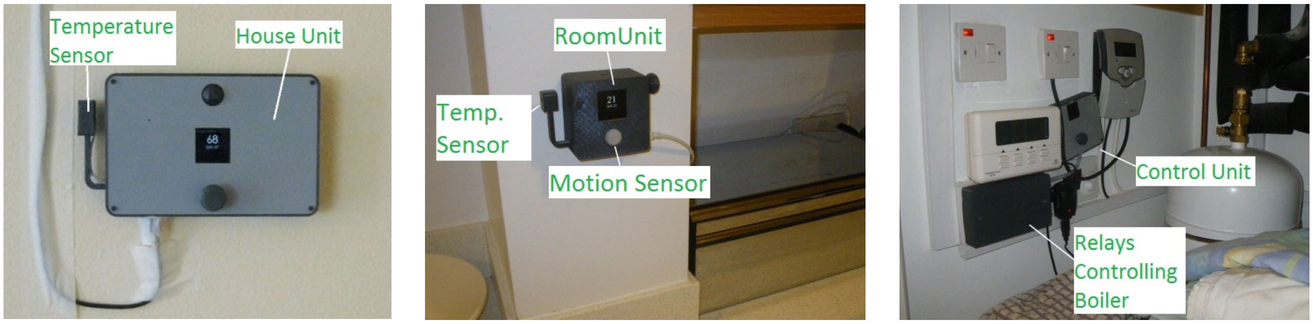

(a) (b) (c)

Figure 1: Photographs of (a) House Unit installed in US house replacing thermostat on wall, (b) Room Unit in UK room, (c) Control Unit beside old house thermostat in UK, wired into boiler control circuitry.

times (i.e. MissTime) – this requires simulator granularity finer than five minutes, better than 1°C, and for individual rooms (when using per-room prediction and control). In contrast, our deployment allows us to characterize PreHeat using real temperatures, gas readings, and occupancy sensors.

EXPERIMENTAL HEATING SYSTEM

We implemented an experimental system, designed to facilitate comparison between PreHeat and existing heating controls. We deployed this system in five homes: three in Seattle, USA (referred to as US1, US2, US3) and two in Cambridge, UK (UK1, UK2) during the winter from January – April, 2011. We chose to deploy in homes of project researchers or our colleaagues, so people knew how to cope if any problems occurred. All of the homes were family homes with two adults and one or more children – ranging from one toddler (UK1) to three school-age children (US1). All US homes used forced air heating, UK1 used both radiators (4 of 10 independently heatable rooms) and underfloor heating (8 of 10 rooms, 2 with radiators too), and UK2 used radiators in all 14 rooms.

Hardware

We built custom hardware, shown in Figure 1, to control the heating systems. In US houses, we replaced the existing thermostat for whole-house control with a "House Unit" built using our in-house-developed prototyping platform, Microsoft .NET Gadgeteer. The core of this unit is a mainboard based on GHI Electronics' Embedded Master module with ARM7 CPU, running custom software in C#. Additionally there were a number of peripheral modules. The House Unit measured temperature using a Sensiron SHT15 sensor (accuracy +/- 0.1°C), and used an Avago ASSR-1611 solid state relay to control the furnace through the existing wiring. It also included a rotary encoder, RGB light level sensor, a 128x128 OLED display and a passive infra-red motion sensor (Panasonic PIR-AMN34111J, 10m max range and 110 degree beam angle).

In UK houses, we controlled the heating on a per-room basis by installing "Room Units" in each space with independent heating. These were similar to "House Units", but instead of a relay had an 868MHz transmitter (TX868-785 from elv.de) to control wireless radiator valves (HHFHT-8V from HouseHeat) which replaced the existing thermostatic radiator valves on the radiators in each room.

In the UK we also deployed "Control Units" that did not have any sensors, but included either Avago ASSR-1611 relays for controlling per-room underfloor heating valves or Omron G6D relays providing 240V control of the house boiler. Turning on heating in any individual room involved activating the boiler (if it was not already active for another room's heating) and per-room valves – either on radiators or for the underfloor heating. These units also controlled the hot water heating for sinks, showers, etc., using a static schedule reflecting the previous hot water schedule.

All hardware was powered through USB mains power adaptors, aside from the battery-powered wireless radiator valves. Power optimization is possible (c.f. battery powered wireless security sensors), but is not our focus. All units had ZigBee radio modules (the "XBee ZB") for wireless mesh communication with a central server PC in each house, running custom software described in the next section.

To evaluate PreHeat, we logged actual gas usage. In the UK, we deployed gas meter readers based on the RFXPulse sensor from RFXCom. In the US, we used daily meter readings provided by the local energy company website.

Occupancy sensing

In US houses, we sensed when occupants were home by using an RFID receiver plugged into the server and placing a small Active RFID tag (RF8315T-s manufactured by Ananiah Electronics) on the house keys of each adult using the house, including the nanny for US3. The option was provided to each household to have RFID tags for kids, but none of the houses accepted this offer, since the kids were in general not expected to be home alone. Visitors were not provided with RFID tags. Each tag sent its identity to the receiver every 5s when in range. The server was placed such that the 8m nominal range included the whole "front hall" area of the house, where we asked participants to leave their keys. We also asked participants to take their keys whenever they left the house (compliance results are described in the Deployment section).

In UK houses, where we could control heating at the room level, we added motion sensors that could determine per room occupancy. These detected motion by any occupant: adults, children or visitors – note that neither UK house had pets. This allowed our predictive system in the UK houses to heat for and predict arrival of any type of occupant, not just those with RFID tags. UK occupants also had RFID tags deployed which were only used during system evaluation, and not for prediction.

Our occupancy sensors (RFID and motion) generated events at discrete points in time. From this data, we derived periods of occupancy by filling in gaps between two sensed events within a certain time difference (2 minutes for RFID, 5 minutes for motion sensing during the day, and 30 minutes for motion sensing during pre-defined sleep hours).

Software

The PC we deployed in each house ran our custom software which managed the network of devices, collected sensor data and ran the heating algorithm. The heating algorithm was provided with a collection of preset parameters. These comprised (for each day of the week): Sleep and Wake times, a Sleep setpoint temperature to use between those times, and Occupied[1] and Away setpoints. In the UK, each room could have its own Occupied temperature. For our experiments, we deployed different heating algorithms, and we switched between algorithms at the preset Sleep time (which all houses chose to be the same time on all seven days of the week). The experimental system was designed so that the only difference between conditions was the algorithm's choice of which setpoint to use at a particular time. Our study used the following algorithms:

Scheduled: Equivalent to a seven-day programmable thermostat, it used preconfigured times for Leave and Return for each of the seven week days. The Away setpoint is used between Leave and Return times, and the Occupied setpoint is used otherwise (other than during Sleep hours).

AlwaysOn: Forced the use of the Occupied setpoint all the time including between Sleep and Wake.

PreHeat: Our prediction algorithm which chooses between Occupied and Away depending on current and predicted occupancy – more details later in this section.

Heating in Anticipation of Occupancy

If a space (whole house in US or room in UK) was deemed occupied by the heating algorithm, this determined the setpoint to use. However, even if a space was not occupied, it may require heating in order to be warm for future need. To achieve this, the system looked ahead up to three hours into the future. If a higher setpoint was expected by the heating algorithm, then the system evaluated the current temperature, target temperature, look ahead duration, and the HeatRate, which is the warming rate of the house in degrees per hour, to decide whether it needed to start heating at the present time. For example, with a HeatRate of 3°C/hour, space temperature of 20°C, and a requirement for 22°C in one hour's time, no heating is yet needed. Twenty minutes later, if the space temperature remained 20°C, then the system would activate the heat in order to warm up in time for the future (predicted or scheduled) occupancy.

User Interface and Overrides

We provided a user interface for occupants on the House Unit/Room Units themselves. This showed the current space temperature and the setpoint. Just as with normal heating controls, we provided a menu option on the units that gave occupants a way to override the system in case they were too cold. We also used the light sensor to automatically dim the screen when the space was dark, to make our system "bedroom friendly".

System Reliability

Since our system ran in real homes, we conducted extensive in-situ testing and implemented a number of features to promote reliability, including the ability to remotely log in to the servers, automatic emails if problems occurred, an auto-restart program to recover from software crashes, etc.

We also implemented a "failsafe mode." In the rare case that the unit was unable to contact the server, while it attempted to recover communication it would act as a local thermostat using the Occupied setpoint. As we later detail, the system achieved an average 99.8% uptime in the study.

PreHeat Occupancy Prediction Algorithm

The PreHeat prediction algorithm works in two ways. First, it uses occupancy-reactive heating; when a space is occupied, it uses the Occupied setpoint (or Sleep setpoint at night). Second, when a space is not occupied, it predicts when it will next be occupied by matching the occupancy data from the current day against historical occupancy data.

We represent space occupancy as a binary vector for each day, where each element represents occupancy in a 15-minute interval, as shown in Figure 2. In the UK such spaces are individual rooms, in the US we use a single whole-house space – no distinction is made by the algorithm. The vector element is 1 if there is any occupancy during the interval or 0 otherwise. As a day progresses, we maintain a partial occupancy vector from midnight up to the current time. To predict future occupancy, we use this partial occupancy vector to find similar days in the past. Specifically, we compute the Hamming distance between the current partial day and the corresponding parts of all the past occupancy vectors. (The Hamming distance simply counts the number of unequal corresponding binary vector elements.) We then pick the K nearest past days for making the prediction. Based on initial experiments, we found K=5 proved to be a good choice for high prediction accuracy.

[1] Although our system supported having different morning and evening setpoints (mirroring many US thermostats), all of our households chose the same value for these, so we refer to a single Occupied setpoint to simplify discussion.

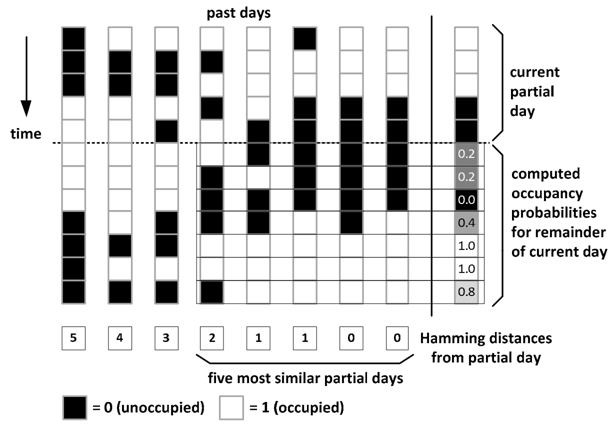

Figure 2: PreHeat prediction algorithm. Each vertical set of blocks represents one day of occupancy split into 15-minute periods. Given a partially observed current day, the algorithm finds the five best matches in the past and averages the remainder of those matched days to compute probabilities for future occupancy.

The predicted occupancy probability for a future time is simply the mean of the corresponding occupancy values in the K nearest past days.

There are two slight variations of this basic algorithm that we implemented to improve accuracy. The first variation distinguishes between weekdays and weekends. Weekday predictions are only computed from past weekdays, and similarly for weekends. This helps accuracy, because households often have quite different occupancy patterns on the two types of days. The second variation augments the beginning of each occupancy vector with four hours of occupancy data from the previous day. This gives predictions near the beginning of the day some extra basis on which to compare to previous days. We also pad the end of each occupancy vector with four hours of occupancy data from the following day. This avoids complexities in making predictions that span midnight.

While in this algorithm we treated every historical day equally for prediction purposes, we could instead automatically adapt to changing occupancy schedules by preferring more recent days when matching.

One advantage of an algorithm like ours is that it gives occupancy probabilities. This gives us the freedom to set a probability threshold for declaring when a space will be occupied. For our experiments, we set this threshold to a neutral value of ½. But, a more environmentally-conscious user may set the threshold higher to ensure that the system is more confident of future occupancy before heating. In contrast, a user more concerned about being warm would choose a lower probability threshold. We demonstrate the effect of this tradeoff in the results section below.

In testing, we found that PreHeat did not perform well in per-room situations for a certain class of rooms – those which are only occupied for very short periods of time and at random intervals, e.g. bathrooms and hallways. These spaces tended to be predicted as “never occupied”. We therefore decided to use whole-house occupancy (determined by using the mathematical union of signals from all motion sensors) to predict and heat these rooms. Intuitively, whenever you are home it is possible that you might use the bathroom. In addition to these “whole house spaces”, this left 6 individually controlled rooms out of 10 for UK1 and 8 rooms out of 10 for UK2.

DEPLOYMENT

Deployment allowed us to evaluate the heating algorithms using real weather conditions, central heating systems, and human occupancy behavior. Our three phase deployment lasted three to four months in each house, with an average of 61 days per house in Phase 2 (see Table 1), the main comparison between PreHeat and Scheduled heating.

Phase 0: Debug/Acclimatization (≥ 7 days)

In this phase we installed our system running the Scheduled algorithm. We showed adult members of the household how to use the UI and attached RFID tags to keys. To get the most accurate schedule we used the household’s current thermostat settings as a baseline and asked household members to update the program if necessary. For per-room heating in the UK we asked participants to provide “Occupied” setpoint temperatures for each room. To model typical systems in the UK which simply turn off between heating periods rather than use a setback, we did not ask for Sleep or Away temperatures, but used a low value of 5°C, which never triggered any heating during our study.

The goal of this phase was to make sure our system was running smoothly in the home and to give people time to adjust their setpoint temperatures or scheduled times if desired. Although changes were permitted in later phases, no household requested to make such a change. We used the average observed heating rate from this phase to set the HeatRate parameter for each house (ranging from 1.5°C/hour in UK1 to 3°C/hour in US2). Phase 0 lasted a minimum of a week and was longest in US1, US2 and UK1 where we did initial in-situ testing.

Phase 1: Initial Data Collection (14 days)

PreHeat’s prediction algorithm requires some occupancy history data to work. We therefore ran a 14 day phase without Prediction, which gave the subsequent prediction phase enough historical data to work accurately. In a real deployment of PreHeat, the system might bootstrap by simply heating during all non-sleep hours until it determined that its predictions were accurate enough to start using Away. In our post hoc analysis, however, we found that prediction accuracy was adequate after just one or two days of occupancy data. Given that prior research [3, 8] has used AlwaysOn as a baseline, during Phase 1 we decided to alternate between the AlwaysOn and Scheduled conditions.

Phase 2: PreHeat vs. Scheduled (48-72 days, 61 average)

This comparison was our main focus and comprised the majority of days in the study (see Table 1). In order to balance any effect weather or household schedule changes may have on our study, we alternated each day between two conditions: using PreHeat's prediction algorithm and the Scheduled algorithm. Using neighborhood weather data provided to each US house by their utility company and from nearby public weather stations in the UK, we determined that, for all houses, the average outdoor temperature for PreHeat days differed from Scheduled days by less than 0.3°C.

Our primary metrics for comparing PreHeat and Scheduled were gas consumption and the MissTime metric used by Lu et. al [8]. They defined MissTime as the total time in minutes that the home is occupied but the temperature was more than 1°C below the Occupied setpoint. We used the same definition, but applied it per-space (room in the UK, house in the US). For gas consumption, we used automated meter readings of the volume of gas used as described earlier. Because the system was deployed in our own houses, we purposefully avoided any subjective metrics.

Households US2, US3, and UK2 each went on vacation for about a week during the study. Since at such times it is usual to do some exceptional thermostat programming (e.g. turn it off), we collected a list of such days and excluded them from Phase 2. This is primarily because vacation times are "unfair" to the Scheduled condition which heats regardless. We took advantage of vacation days to run additional AlwaysOn days in these houses. We continued to collect occupancy data during vacations, as PreHeat is robust to having these periods in the historical data set.

Reliability During Deployment

During the study period we had a small number of technical problems which caused us to remove three days from the analysis (N.B. In Table 1 and elsewhere, we have excluded these days already). In UK1, an electrical fuse trip caused two days to be excluded. In UK2, the XBee radio module attached to the PC "hung" and required manual power cycling, causing one day to be excluded.

Of the study days used for analysis, the system was in a not-fully-functional state for an average of three minutes per day (uptime 99.8%). House UK2 had the most problems, averaging eight minutes per day where units were in "failsafe" mode. Most problems were due to XBee, and the system recovered automatically.

Ensuring Occupancy Correctness

Because our MissTime metric is only as valid as the occupancy data used, we took extensive steps to ensure accuracy. Each server directed an email to a house occupant and a project member every morning with a summary of the previous day's per-space occupancy. The recipients looked for any discrepancies, consulting family members if necessary. The documented errors included both "social" failures due to forgetting to carry an RFID tag, and "technical" failures where some aspect of the sensing failed. We also ran data integrity checks on the 37,000 occupancy records generated over the course of the study to find failure cases: In the US, we examined days where RFID tags disappeared many times (a radio range failure). In the UK, we used the RFID data (which was not used for heating prediction purposes) to highlight times when the house was occupied but no motion sensors were active.

Using the documented errors and integrity checks we found a set of 57 issues that could affect the analysis and corrected them in a "ground truth" occupancy table. In the US, there were 17 corrections, of which 14 were "social", and 12 of those occurred in US1. This household had keyless entry so the RFIDs were not as convenient. In the UK, 21/40 corrections were due to occupants being asleep outside the Sleep/Wake period (and thus difficult to detect using motion). In future work, we hope to address sleep detection. Issues with a motion sensor in UK2 not completely covering a room resulted in 18/40 corrections.

We compared the MissTime metric using the original table against the ground truth table and found that only a single instance actually altered the result set (this impacted UK2 - the corrected data is reported in this paper). In the remaining instances, the house was warm enough to prevent the error from impacting the occupants. It is important to note that the PreHeat algorithm did NOT use corrected data; it predicted based on live data that included any errors. This is more representative of a real-life deployment.

PREHEAT VS SCHEDULED RESULTS

Figure 3 and Table 1 summarize how the measured gas consumption and MissTime metrics differ between the Scheduled and PreHeat conditions in each house.

In the UK, PreHeat performed better than Scheduled on both metrics. PreHeat saved energy by not only selecting better heating times, but also by heating rooms different amounts – more details to follow in the Per-Room Heating section. Gas usage decreased from Scheduled by 18% in UK1 and 8% in UK2. In both conditions, the UK houses had comparatively little MissTime, but PreHeat also succeeded in decreasing this by 38% and 60%.

In the US houses, MissTime improved by a large factor (84%, 88% and 92% reductions, a factor of 6-12 decrease!).

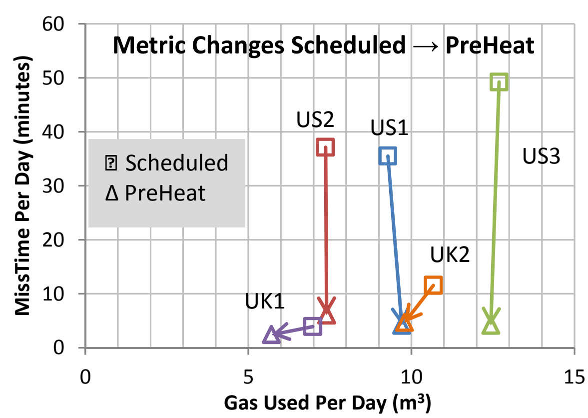

Figure 3: PreHeat improves the trade-off between gas used and MissTime compared to Scheduled

	US1	US2	US3	UK1	UK2
Phase 2 Days (Alternating PreHeat and Scheduled)	72	64	58	62	48
Average Daily Occupancy	89%	69%	67%	73%	68%
Savings in MissTime PreHeat vs. Scheduled	88%	84%	92%	38%	60%
Savings in Gas Used PreHeat vs. Scheduled	-5% (worse)	-1%	2%	18% (better)	8%
Savings in Gas Used PreHeat vs. AlwaysOnModel	3%	17%	10%	27%	35%

Table 1: PreHeat decreased MissTime in all houses. Gas Used decreased in PreHeat in the UK homes with per room heating and was equivalent or slightly higher for the three US homes.

This is further shown by Figure 4 which illustrates that MissTime happens unevenly through the study – over half the days had no MissTime, but one Scheduled day in five (in US houses) had over one hour. PreHeat, unlike a static schedule, is able to dynamically heat day-by-day to more closely match the occupancy. To give one example, in US2, PreHeat heated more on weekends and less on weekdays than Scheduled, while being adaptive to instances of absence on weekends or occupancy on weekdays. PreHeat was able to achieve these large improvements in MissTime while using a similar amount of gas – saving slightly in US3 and using slightly more in US1 and US2. In the Occupancy Prediction section, we discuss how PreHeat could instead be tuned to favor energy savings at the expense of a smaller improvement in MissTime.

Many More Overrides in Scheduled

Our system's user interface allowed occupants to override the current action. We classify overrides into two categories: "time" overrides when the system was not heating and the user overrode it to heat to the normal Occupied temperature, and "temperature" overrides when a non-standard temperature was set. During the study, there were 23 time overrides, 21 of which occurred in the Scheduled condition. The "time" overrides in Scheduled happened when people were home on vacation, sick, or otherwise home at times when Scheduled was not heating the house and they felt cold. The overrides caused more energy to be used, but without these overrides, the MissTime metric would have been (even) higher. The occupancy-reactive element of PreHeat reduces the need for overrides. There were two "time" overrides in the PreHeat condition, which were due to occupancy sensing failures – a guest in US2 who did not have an RFID tag, and a resident of US1 left her RFID tag in the car. The small number of "time" overrides in PreHeat supports our claim that PreHeat requires less programming of a heating schedule by occupants (since overriding is a form of programming).

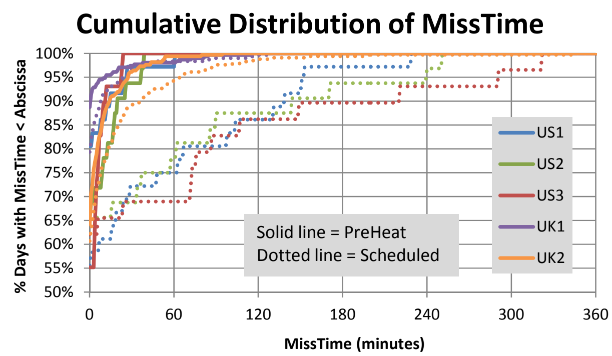

Figure 4: Occurrences of long MissTimes are much reduced with PreHeat, particularly for US houses. (Note Y-axis intercept at 50%)

There were nine "temperature" overrides during Phase 2. These types of overrides can occur in either condition if an occupant wanted a deviation from the normal setpoint. In US3, participants overrode the temperature to be 1-3 degrees higher than the setpoint seven times: three in Scheduled and four in PreHeat. In UK2, an occupant overrode the temperature to be lower than the setpoint twice, both in PreHeat, because it was sunny and he didn't feel the heat needed to be on.

Per-Room Heating

Although we did not study the effect of per-room heating through a direct comparison, we can still get a measure of how well PreHeat tailored its behavior to individual room occupancy patterns.

Figure 5 shows UK1's six per-room-controlled spaces (the other four spaces used whole-house occupancy for heating as previously described), illustrating average occupancy in each space and the amount of time that the Occupied setpoint was used in both PreHeat and Scheduled conditions. We see that Scheduled heats around the same amount in each space (variation is due to turning on a little earlier or later in each space depending on how cold it gets when away/asleep). PreHeat tracks occupancy fairly well.

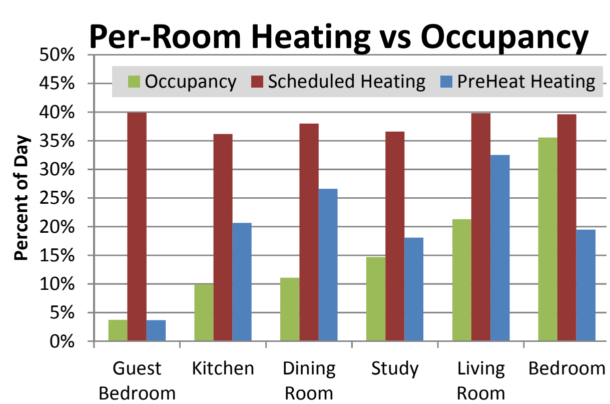

Figure 5: PreHeat heats higher-occupied rooms more, while Scheduled does not adapt per-room.

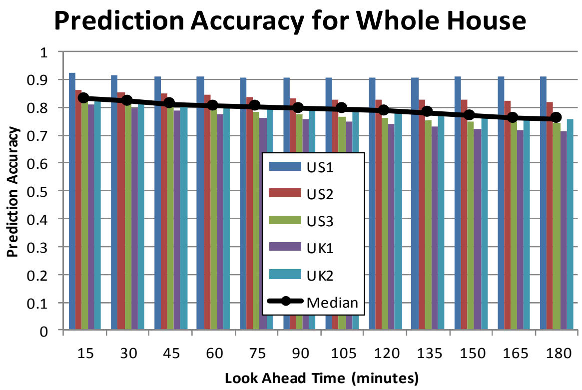

Figure 6: Prediction accuracy varied only slightly with the look ahead time.

Note that the rightmost room is the bedroom where all occupants sleep, so this space needs less heating than occupancy might suggest.

OCCUPANCY PREDICTION

Occupancy prediction is a key part of PreHeat. This section explores how accurate the predictions are compared to actual occupancy.

Figure 6 shows the overall prediction accuracy for different look ahead times for the five houses. We compute accuracy for a given look ahead time as simply the number of correct occupancy predictions (either true or false) divided by the number of attempted predictions. The graph reveals that prediction accuracy for all houses was generally high. The median line shows that prediction accuracy only slightly decreases for longer look ahead times.

Since the system needs advance notice of the need for heat in order for the space to reach the desired setpoint in time, we should evaluate the prediction based on how well it can achieve this goal. Examining each daytime heating instance during Phase 2, we find that 91% of the time, the system needed 90 minutes or less advance notice. For the remainder of our prediction assessments, we evaluate based on this 90-minute look ahead time.

Figure 7 shows the prediction accuracies for 90 minutes into the future. For comparison, it also shows the prediction accuracy of the manually programed Scheduled condition, which was worse by a median 10 percentage points. We suspect that the improvement from using PreHeat would be much greater for the general population, many of whom do not keep up to date programs on their thermostats or do not have programmable thermostats [3].

As one might expect, prediction accuracy varied with the time of day. Accuracy is highest during sleep time and drops during the day when occupancy is naturally less predictable. Because prediction accuracy is high during sleep times, we have eliminated sleep times from all our prediction accuracy assessments. Thus, all our prediction

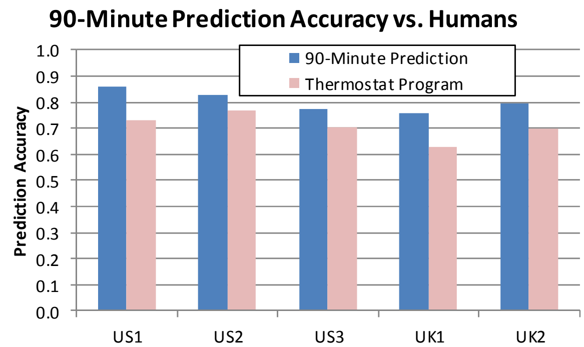

Figure 7: Our prediction algorithm was consistently more accurate than carefully programmed thermostats.

assessments, including those in Figure 6 and Figure 7, pertain only to non-sleep times, eliminating the uninteresting boost in accuracy we would otherwise get.

We also evaluated how prediction accuracy varied with the day of the week. As expected, weekend days are worse, with Sunday being the better of the two. We attribute this to the fact that our households generally had parents with regular working hours and children with regular school hours during the week.

Our prediction algorithm computes occupancy probabilities that are then subjected to a threshold to make a concrete occupancy prediction. For our study, this threshold was set to a neutral value of ½. The effect of this threshold is shown in the receiver operating characteristic (ROC) curve in Figure 8. For different values of the probability threshold, these curves show the tradeoff between false positives (mistakenly predicting positive occupancy when the home would actually be empty) and true positives (correctly predicting positive occupancy). Our selected threshold value is indicated by a dot on each ROC curve.

The ideal operating point is in the upper left corner, where the false positive rate is 0.0 and the true positive rate is 1.0.

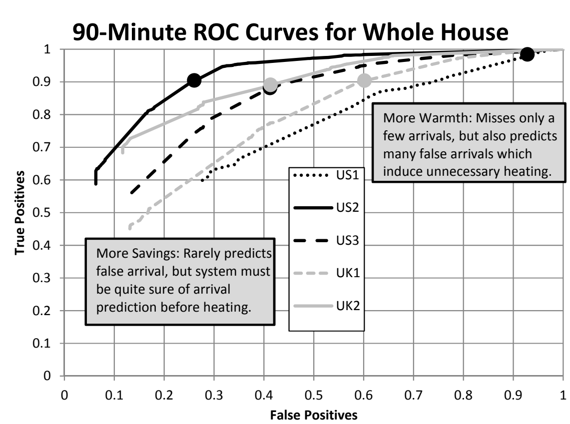

Figure 8: These ROC curves demonstrate the tradeoff in prediction errors due to adjusting the probability threshold.

In reality, as with almost all detection problems, increasing the true positive rate comes at the expense of increasing the false positive rate. An environmentally-conscious user, who is more concerned about energy rather than warmth, would operate with a high probability threshold, which is toward the left side of the ROC curve. This would give fewer mistaken heating events at the expense of missing times when the heat should actually be turned on. A more warmth-sensitive user would operate with a lower probability threshold, which would cause heating more often during times of both occupancy and non-occupancy.

Our study results show that PreHeat improved MissTime in all houses by substantial amounts. This indicates that our selected threshold was generally closer to the "warmth" end of the curves than the "savings" end, as verified in Figure 8. The operating point for US1 is biased far toward warmth. This is because this house was almost continuously occupied, resulting in prediction probabilities that were generally high (including sleep time, US1 was occupied 89% of the day vs. a median of 69% for the other houses). Even with a relatively high false positive rate of 93%, the house was occupied so much that its overall 90-minute prediction accuracy was highest of all the houses at 86%.

Finally, we also looked at prediction on a per-person basis, by applying our algorithm to each individual RFID tag using the tag's unique ID. Running our algorithm in this way showed that our prediction accuracy is even higher, with a median accuracy of 97% compared to a whole-household median of 80%. This means our algorithm could be used to accommodate different temperature preferences when only one of the home's occupants is present.

DISCUSSION

We now discuss how PreHeat compares with AlwaysOn and to occupancy-reactive heating, and share some observations from our experiences of living with PreHeat.

Comparison with Other Heating Algorithms

Our main study directly compared PreHeat with Scheduled based on actual measurements. We now discuss how PreHeat compares to other algorithms.

AlwaysOn

Prior research [e.g. 3, 8] frequently uses the energy necessary to keep the house at a permanent setpoint (AlwaysOn) as a baseline for comparison. Given its past inclusion, we also provide this comparison baseline; however, we believe the comparison between PreHeat and Scheduled to be more meaningful.

Since there were not sufficient days in a single winter to run three conditions, we elected to only gather AlwaysOn data during Phase 1 and vacations. We then used linear regression on this data to model the gas consumed given the average outside daily temperature. The linear regression showed average daily temperature was a very good predictor of gas used for the US houses, where more than 92% of the variance in the data was accounted for in the linear regression equation (all $R^2 > 0.92$). Modeling did not

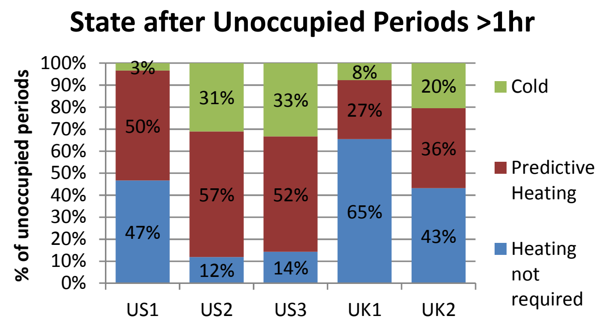

Figure 9: Predictive heating was important in all houses. While occupancy-reactive heating would have sufficed in 36% of cases (average across households) where heating was not required, predictive heating was used to achieve warmth on return for an average of 45% of cases.

work as well in the UK, with $R^2 = 0.69$ in UK1 and $R^2 = 0.78$ in UK2. However, the models still give us some ability to compare PreHeat to AlwaysOn.

For each PreHeat day we used the average daily temperature as input to the linear regression equation. We show the comparison between the model's result and the actual gas used by PreHeat in Table 1. As expected, PreHeat uses less gas than the AlwaysOn model (ranging from a 3% to a 35% decrease). Although the difference in US1 is quite small (3%) we attribute this to the fact that the house was only empty 11% of the day which does not provide much opportunity for savings.

Occupancy-reactive

The PreHeat system incorporates both occupancy-reactive and predictive heating. We wished to evaluate how often just occupancy-reactive heating would have been sufficient in our study, and thereby get a measure of the value of the predictive element. We therefore looked for instances when someone entered a space that had been unoccupied for more than 60 minutes on PreHeat days. As Figure 9 shows, occupancy-reactive heating would have worked 36% of the time (average for all houses) without any MissTime – i.e., the space was still warm enough when someone returned. However, in 45% of the instances PreHeat heated the unoccupied space prior to the person's return and so occupancy-reactive heating would have caused MissTime to occur. Thus we can conclude that predictive heating plays a very significant part in the PreHeat system.

Living With PreHeat

We quantitatively compared the Scheduled and PreHeat algorithms because of the potential for bias when using our own system. However, based on our experience with living with PreHeat, we report a few qualitative examples of how PreHeat gracefully adapted to our homes.

PreHeat Better Handles Weekend Chaos: During the setup phase, we heard that weekend schedules varied more than weekday schedules, and that it was difficult to program the thermostat for weekends. Households adopted one of two strategies for their (previous) thermostat programs, either heating all day on the weekends (US1, US3) or using a shorter away duration than their weekday schedule (US2, UK1, UK2). These strategies were carried over into the study's Scheduled condition. In either case, PreHeat better handled weekend heating without requiring manual effort by household members. For households with a weekend setback, the house stayed warm when they were home on PreHeat days without an override. Households that left the heating on all day had the potential to save energy on weekend days if they left home.

PreHeat Supports More Complicated Occupancy Patterns: Programmable thermostats typically allow people to schedule one away period per day. In contrast, PreHeat can predict multiple away periods. For example, by the end of the study, PreHeat was correctly predicting that US2 would come home and then leave again for a regular Friday evening appointment, and similarly for UK1 on Tuesday evenings. More valuable was PreHeat's ability to predict on a per-room basis in the UK. Even if households had the ability to program on a per room basis, it is unlikely they would make the effort to maintain up-to-date programs in each room separately. We also observed that PreHeat handled occupancy patterns that re-occurred, but not on a weekly basis. For example, PreHeat smoothly handled the fact that a member of US2 arrived home early roughly every other Tuesday afternoon.

PreHeat Adapts to Changing Schedules: In US3, the Nanny who stayed at home after school with the kids took another job during the study. As is likely typical after this type of change, it did not occur to the occupants of US3 that they should change their heating schedule. However, PreHeat adapted over time and began correctly predicting later arrival times at the home.

CONCLUDING REMARKS

By predicting future occupancy from historical data and current occupancy, the PreHeat system provides a better trade-off between energy use and MissTime (the amount of time an occupied space is cold) than a thermostat program, and does so without requiring a user to program an occupancy schedule (which past research has shown that many users fail to do). We evaluated PreHeat using a real deployment in five family homes during winter 2011, alternating days between PreHeat and scheduled heating to provide a direct comparison with measured gas consumption and MissTime.

In three US homes, we found that PreHeat reduced MissTime by a factor of 6-12 while using around the same amount of gas. Across two UK homes, PreHeat halved MissTime and also reduced gas usage by 8% and 18% - this is because UK homes used PreHeat on a per-room basis, so it was able to make additional savings by heating rooms adaptively at different times of day, again without requiring any programming of per-room schedules.

Our research suggests several interesting directions for future work. We want to investigate improvements to the PreHeat algorithm, e.g. using other sources of data such as location from phones, or further exploring per-person predictions. Sleep detection would assist with automatically determining times for setpoint changes. A more sophisticated heating model could also enable more savings, e.g. by predicting departures and "pre cooling" since houses stay warm for some time. Finally, we plan to explore exposing the high-level tradeoffs PreHeat offers between the likelihood of being warm when arriving home unexpectedly and energy consumption to users so they can customize based on their personal preferences.

ACKNOWLEDGEMENTS

Thanks to Kori and Jeff Quinn, and to our families, for being our guinea pigs!

REFERENCES

1. EIA, USE.I.A. Table 2.5 "Household Energy Consumption and Expenditures by End Use and Energy Source, Selected Years 1978 – 2005." Retrieved in April 2011 at http://www.eia.doe.gov/emeu/aer/consump.html.
2. EPA, USE.P.A., Summary of Research Findings From the Programmable Thermostat Market. Available from: http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/thermostats/Summary.pdf
3. Gupta, M., S.S. Intille, and K. Larson. "Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges." *Proc. of Pervasive,* 2009.
4. Hargreaves, T., Nye, M., and Burgess, J. "Making energy visible: A qualitative field study of how householders interact with feedback from smart energy monitors." *Energy Policy* 38:10, pp. 6111-6119, 2010.
5. Krumm, J. and Brush, A. "Learning Time-Based Presence Probabilities", *Proc. of Pervasive 2011*, 2011.
6. Meier, A., Aragon, C., Hurwitz, B., Mujumdar, D., Peffer, T., Perry, D., Pritoni, M. "How People Actually Use Thermostats", *Proc. of ACEEE 2010.*
7. Mozer, M.C., Vidmar, L., and Dodier, R.M. "The Neurothermostat: PreHeat Optimal Control of Residential Heating Systems." *Advances in Neural Information Processing Systems* 9, pp. 953-959, 1997.
8. Lu, J, Sookoor, T., Srinivasan V., Gao, G, Holben, B., Stankovic, J., Field, E., Whitehouse, K. "The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes." *Proc. of SenSys,* 2010.
9. Sonderegger, R. "Movers and stayers: The resident's contribution to variation across houses in energy consumption for space heating." *Energy and Buildings* 1:3, pp. 313-324, 1978.
10. Wilhite, H. and Ling, R. "Measured energy savings from a more informative energy bill." *Energy and Buildings* 22:2, pp. 145-155, 1995.