

Adding GPS-Control to Traditional Thermostats: An Exploration of Potential Energy Savings and Design Challenges

Manu Gupta, Stephen S. Intille, and Kent Larson

House_n, Massachusetts Institute of Technology
Cambridge, MA 02142 USA
{manug, intille, kll}@mit.edu

Abstract. Although manual and programmable home thermostats can save energy when used properly, studies have shown that over 40% of U.S. homes may not use energy-saving temperature setbacks when homes are unoccupied. We propose a system for augmenting these thermostats using just-in-time heating and cooling based on travel-to-home distance obtained from location-aware mobile phones. Analyzing GPS travel data from 8 participants (8-12 weeks each) and heating and cooling characteristics from 5 homes, we report results of running computer simulations estimating potential energy savings from such a device. Using a GPS-enabled thermostat might lead to savings of as much as 7% for some households that do not regularly use the temperature setback afforded by manual and programmable thermostats. Significantly, these savings could be obtained without requiring any change in occupant behavior or comfort level, and the technology could be implemented affordably by exploiting the ubiquity of mobile phones. Additional savings may be possible with modest context-sensitive prompting. We report on design considerations identified during a pilot test of a fully-functional implementation of the system.

1 Introduction

With only 5% of the world's population, the U.S. uses 25% of the world's energy [1]. The U.S residential sector is responsible for 21% of the total U.S energy consumption, and heating and cooling accounts for 46% of the total energy consumed in U.S residential buildings. Overall, 9% of total U.S energy consumption is expended on residential heating and cooling [2, 3]. Forty-nine percent of homes in the U.S are unoccupied during the day, and it is estimated that in 53% of U.S homes the temperature (T) is *not* lowered during the daytime when no one is at home in winters (conversely, in 46% the T is not raised in summers) [4]. Even in the 30% of the U.S homes that have programmable thermostats (P-Therms), as many as 44% may not use daytime setbacks to save energy [4]. As **Table 1** shows, as many as 55 million U.S. households – some with manual thermostats (M-Therms) and some with P-Therms – may not change their T settings when no one is home.

Although per capita consumption of energy is much lower in other countries [1], a significant amount of energy is likely being wasted heating and cooling unoccupied environments in many industrialized countries because common thermostats do not

Table 1. Thermostat usage statistics in the U.S (summarized from [4])

(In millions)	Total homes in the U.S.	Estimated no. of homes using setback when away	Estimated no. of homes not using setback when away
Manual Thermostat	62.16	21.7	40.46
Programmable Thermostat	33.3	18.7	14.60
Total	95.46	40.4	55.06

adapt to variable occupancy schedules and because people have difficulty setting and optimizing P-Therms [5]. The challenge, therefore, is to create a system to augment existing thermostats so that regardless of what the home occupants do, the thermostat (1) saves energy, (2) requires non-burdensome user input and no reliance on memory, and (3) doesn't sacrifice comfort, where we define comfort as ensuring that the home is always at a desirable temperature upon return. Additionally, a thermostat needs to be inexpensive to use and install.

We describe a concept for augmenting existing thermostats with a just-in-time heating and cooling mode that is controlled using travel-to-home time computed from location-aware mobile phones. Although existing P-Therms can save substantial amounts of energy when used effectively [6], we show, via a set of simulations using real travel data and home heating and cooling characteristics, that the proposed just-in-time system augmentation might provide energy savings for the substantial number of people *who do not use M-Therms or P-Therms optimally*. The system that we propose does not require users to program occupancy schedules. In fact, *no change in behavior on the part of the home occupants from what they currently do is necessary*. We focus on standalone housing and commuting patterns common in the northern U.S. and leave the question of how these results might generalize to other climates, housing types, and lifestyles for future work.

2 Prior Work

Pervasive computing systems that can infer context clearly offer potential for energy saving. Harris et al. [7], for example, argue that context-aware power management (CAPM) could use multi-modal sensor data to optimally control the standby states of home devices to optimize energy use, reducing so-called vampire power consumption [8]. They conclude that to optimally save energy, in addition to predicting what someone is currently doing, a system should predict what someone is *about* to do. Reliable detection of intentionality to control appliance energy use indoors is a difficult problem that is the subject of ongoing research [9]. Nonetheless, Harle and Hopper [10] showed that even without such prediction, in one office building using location of occupants would have permitted energy expended on lighting and "fast-response" electrical systems to be reduced by 50%.

Although inefficient use of electrical devices can be a substantial source of energy waste in a home or office, others have instead focused on improving home thermostats to lower heating, ventilation and air condition (HVAC) costs. A thermostat balances two competing factors: energy savings and air temperature/humidity comfort levels. There are three common types [11].

Manual thermostats (M-Therms) can be the most energy efficient option. People who set the T very low in winter when they leave the house and then turn up the T when they return achieve maximal energy savings but with significant discomfort upon return to the home. Avoiding that discomfort may be one reason that over 65% of people with M-Therms do not use setbacks when they are away from their homes in winter [4].

Programmable thermostats (P-Therms) automatically regulate the T according to a pre-scheduled program. P-Therms do not adapt to variable occupancy schedules – if schedules change, the user must remember to re-program the system in advance, and reprogramming is often tricky with current interface designs. The lack of responsiveness and difficulty of programming may be one reason that over 43% of people with P-Therms do not use daily setbacks when away in winter [4].

So-called intelligent thermostats have “adaptive recovery control,” so that rather than starting and stopping based on timers, they set the T when away to ensure that given typical heating/cooling patterns, the home will reach the comfort T at the right time. These thermostats may also learn the T preferences of the user for different contexts [12] and use occupancy sensors to infer occupancy patterns [13, 14]. Others use light levels to change the T settings in the house [15] or control the air velocity and direction [16]. Some even have persuasive elements, such as informing users about the minimum T settings that can produce the desired comfort level [12, 17]. When these systems imperfectly infer behavior patterns, however, they optimize savings at the expense of comfort, and they typically require complex sensor installations to be retrofit into the home.

Unfortunately, all of these thermostats are often misused. An estimated 25-50% of U.S. households operate the thermostat as an on/off switch rather than a T controller [18]. A common misconception is that the more one changes the T dial, the faster the thermostat will make T change [19, 20]. Also, it has been shown that P-Therms do not save as much energy as predicted [5, 21, 22], most likely because they are difficult to use [5, 23]. Clearly, it is important that the thermostat interface be made as simple as possible.

3 Opportunity

The key idea advocated here is to *augment* current thermostats with the ability to control heating and cooling using travel time, as determined automatically via GPS-enabled mobile phones that will become commonplace.¹ When the thermostat is not being used regularly in setback mode, the thermostat should switch to this “just-in-time” travel-to-home-time mode. In this mode, the thermostat system communicates with the GPS-enabled mobile phones of the residents. Based on the location of the residents as determined by each resident’s mobile phone and free geo-location mapping services, travel-to-home time is continuously estimated. The thermostat uses travel time of the home occupants, inside and outside T, and heating/cooling characteristics of the home to dynamically control the thermostat so that energy savings are

¹ In this work we use GPS data and the terminology GPS thermostat (GPS-Therm), but phones may also use multiple methods to determine location (e.g., cell tower triangulation or beacons).

maximized without sacrificing comfort. By setting T as a function of the fastest possible return time of the closest resident (and the other factors mentioned above), the system ensures that the home will always be comfortable on return.

The system has the following characteristics: (1) it requires no thermostat programming from the user, (2) it adapts to irregular schedules, (3) it ensures that the user always returns to a comfortable house, (4) it creates opportunities for motivating additional savings using context-sensitive prompting, and (5) and it does not require installation of a complex new sensor system in the home. As we will discuss, in its most straightforward implementation, it does not save more energy than a M-Therm or P-Therm that is regularly used with daily setbacks, but it will save energy overall for the general population because so many people fail to use their thermostats properly. The concept, therefore, is to layer the GPS thermostat (GPS-Therm) capability on top of existing thermostats, so that the GPS system engages only when users are not using a more efficient setback strategy.

This work is inspired by solutions for controlling appliance use in the home or office based on indoor location [9, 10, 24], but the proposed system does not require an extensive sensor or distributed appliance control network to be installed in the home to achieve savings. We make only the following assumptions: (1) that mobile phone adoption trends continue so that in many households everyone who travels alone will have a phone, (2) that within a few years nearly all new phones will have location-finding and Internet data transfer capabilities, and (3) that many homes will have Internet access and home wireless networks. For households where these assumptions hold, we discuss the energy saving potential of the system.

4 Measuring Potential Energy Savings

In this section we describe the results of an exploratory simulation we conducted to better understand the extent to which a GPS-controlled thermostat system might save energy.

4.1 Data Collection

We recruited 8 people living in 4 different homes in the greater metropolitan area of Boston, Massachusetts using mailing lists, flyers, internet advertisement and word of mouth. None of the subjects had any affiliation with the research team. Each person worked outside of his or her home and had a separate car that was used as the main mode of transport. Each house had a heating system that was not shared with any neighboring residences (see **Table 2**). The study was approved by our human subjects review board.

To gather realistic data on travel patterns, between March and June, 2008 a Track-Stick Pro GPS logger [25] was installed on the dashboard of each vehicle of each member of each house. The logger was plugged into the cigarette lighter socket. These loggers were left for up to 3 mo, recording position of the vehicle each minute whenever it was operating. Data were recovered from each logger every 4 weeks.

To measure the heating and cooling properties of the homes, two T and humidity loggers (EL-USB-2, EL-USB-1) [26] and Logit LCV electrical current loggers [27]

were installed in each house for 3 d of measurement. One T and humidity logger was placed near the thermostat and another outside of the house on the north facing wall or window, where the sensor was not in contact with direct sunlight. The current logger was attached to the current-carrying wire from the thermostat to the HVAC system to gather information about when the heating system was activated.

During the 3 days that the sensors were installed, each home had its P-Therm re-programmed by the investigators as follows: set to 50°F (10°C) from 9 AM to 4:30 PM and set to a comfortable T, typically between 67-69°F (19.4-20.5°C), other times.

In addition to the participant homes, data were collected from an unoccupied, newly constructed and well insulated (R-25) 1100 ft² condominium (control house) in December, 2007. A typical day for the type of people who participated in the study would be a person leaving at 7 AM for an 8 AM to 6PM workday with arrival home at 7 PM. This routine was scheduled on the P-Therm for 3 wk, with a comfort T of 69°F (20.5°C) and a setback T of 45°F (7.2°C), the Energy Star recommended settings [28].

Table 2. Participant house details. At the time of the experiments, the cost of natural gas was \$1/therm. The cost of oil was \$1.20/liter. Hot water heating systems used radiators. All the homes had programmable thermostats.

House	Heating fuel	Heating system	Capacity (Btu/h)	Insulation	Commute travel time	Days of data	Vehicles
1	Gas	Forced air	100k	Low	35 min-50 min	75	1
2	Gas	Hot water	130k	Medium	10 min-15 min	75	2
3	Oil	Hot water	133k	Medium	7 min-10 min	63	2
4	Oil	Hot water	154k	High	20 min-35 min	63	2
Control	Gas	Forced air	100k	Very high	Simulated	90	0

4.2 Evaluation

Software was written to simulate the functionality of manual, programmable, and GPS-controlled thermostats. The simulator, which uses the same algorithm later described when discussing a real-time, fully-functional prototype in Section 6, requires the following for input at each point throughout the day at 1 min intervals: (1) indoor T, (2) outdoor T, (3) latitude/longitude coordinate for each occupant's phone (if available), and (4) heating/cooling tables for the home. Additional information can be provided to the simulator when modeling various conditions (e.g., minimum allowable T in the home, occupant schedules, a T setting for P-Therm and M-Therm, and heating system type).

Heating/cooling tables were created for each house using the T profile data collected over 3 d. A heating table was created for heat gain (i.e. the time it takes for the house to heat up by 1°F (-17.2°C) from each starting T given an outdoor T with the heating system running at full capacity), and a cooling table was created for heat loss. The 3 d of data typically do not span the entire range of outdoor T for winters for the region (0°F (-17.7°C) to 60°F (15.5°C)). Therefore, values not directly observed were estimated from the 3 d of data and the energy transfer equation for a building [29].

At each point in time throughout the day when a longitude and latitude coordinate is available, to simulate operation of the GPS-Therm, the simulation software sends the location and the occupant's home coordinate to the MapQuest web service [30]

and obtains estimated travel-to-home time. The MapQuest web service uses a proprietary algorithm for calculating travel time and distance that appears to use road type, speed limit, and distance, but not local traffic conditions. An outdoor T file was created for the greater Boston area for the entire duration of the study by accessing two online weather archive databases.

Each minute the simulator outputs the target T based on the travel-to-home time, estimated indoor T, and HVAC on/off cycle duration. It also outputs the simulated indoor T and the HVAC cycle state (on/off). Using either the pre-programmed times for leaving for work and returning or the time detected when someone returned home, the system also simulates the operation of the M-Therm and P-Therm.

4.3 Results with Common Travel Patterns for Daily Workers

First we discuss simulator results using the control home and simulated travel patterns where people commute every weekday with average commute times of 15 min, 26 min (Boston’s mean commute time) [31], and 90 min. We assume that the comfort T of the home is set to the Energy Star recommended setting of 69°F (20.5°C) [28]. We compare four scenarios. The first is the *baseline*, where the thermostat is set to the comfort T at all times of day. As indicated in Section 1, many people [4] with both P-Therms and M-Therms do not use setbacks at all. The second is *manual setback*, where the T is lowered manually upon leaving the house and raised manually upon returning. The third is *programmable setback*, where schedules are programmed for lowering and raising the setpoint each day based on standard work patterns. In programmable mode, we assume that the system starts heating 30 min before the return time and maintains a target T (setback) of 60°F (15.5°C) during the day when the home is unoccupied. Finally, the last case is the *GPS-thermostat*, where the target T is set as a function of travel-to-home time. All of the savings reported are calculated with respect to the baseline condition.

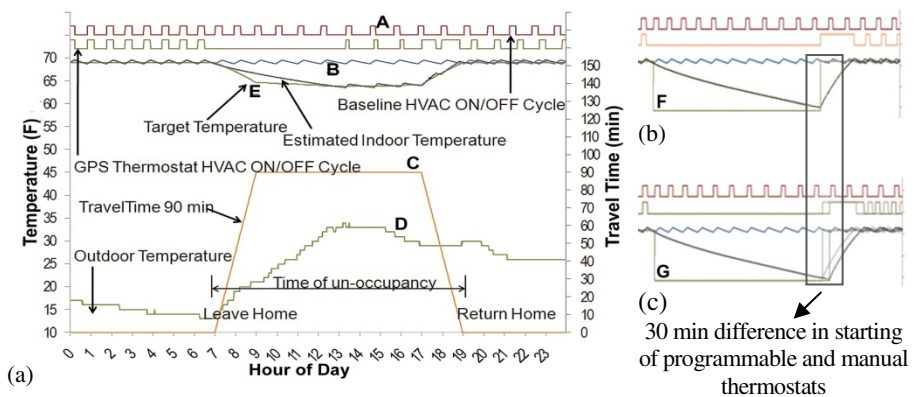


Fig. 1. A typical day with 90 min simulated travel time and simulated results with different thermostat types, as explained in the text

Figure 1(a) shows the simulation results for a 24 h period assuming a long travel time to and from work of 90 min each way. Approximately 1.7 million people in the U.S. commute for 90 min or more each way [26]. Line A shows when the simulation software estimates that the heater will cycle on and off in the baseline condition, with up being an on-cycle. Line B is the baseline T, showing how the target T will oscillate around 69°F (20.5°C), assuming the HVAC triggers when the T drops 0.3°F (0.17°C) below the target (at 68.7°F (20.4°C)) and runs until the T exceeds the target by 0.5°F (0.28°C) above the target (or 69.5°F (20.8°C)). Line C shows the 90 min travel time, leaving at 7 AM and returning at 7 PM, assuming an 8:30-5:30 workday. Line D shows the outdoor T fluctuation, which does influence cycle times (a small change can be seen in lines A and B from morning to mid-day in cycle length). Line E is the target T determined based on travel distance. This value is determined using the estimated heating/cooling parameters of the home at various indoor and outdoor T. It drops as the travel time increases and plateaus at the lowest T that will allow the home to heat back up in time to achieve the comfort T given the travel time. In **Figure 1(b)**, line F shows the target T for the P-Therm simulation. In **Figure 1(c)**, line G shows the target T for the M-Therm simulation.

First we compare the just-in-time GPS-Therm directly to M-Therms and P-Therms. In addition to presenting results for daily savings when the devices are properly used, we present results in terms of “expected energy savings” and “expected monetary savings.” Expected savings is equal to the estimated savings multiplied by expected compliance of use of the particular thermostat type. The expected compliance of a M-Therm assumes 35% [4] of users will use manual setbacks when leaving the house. The expected compliance of a P-Therm assumes 56% of users will have it programmed to use setbacks. Although the GPS-Therm requires no action on the part of the user, we assume that 10% of the time the user may forget to take the phone, lack a GPS connection, be out of mobile phone coverage, or discharge the phone battery (resulting in a relatively high estimated compliance of 90%).

Table 3 shows the simulation results for the three thermostat types for different travel times. The savings using the GPS-Therm increase as the travel time increases, whereas the savings from P-Therm and M-Therm are constant and independent of travel time. The expected savings from the GPS-Therm begins to exceed the expected savings from the P-Therm and M-Therm when travel time reaches 60+ min. The simulations show what we expect to be true: that maximal savings can be achieved by simply turning off a heating or cooling system whenever someone leaves home and sacrificing comfort on return. Using travel time provides less substantial savings than P-Therm for people with short commutes and predictable work schedules, because if someone works near home the system does not allow the T to drop very far to ensure that the home can return to comfort quickly enough no matter when a person leaves work for home. Therefore, even though the GPS-thermostat is likely to have a much higher compliance than M-Therms and P-Therms, the simulations suggest advising against generally replacing P-Therms with GPS-Therms. Instead, the travel-time control should activate only in those situations where the system detects that manual or programmatic setbacks will not be in use.

Table 3. Simulation results on daily savings on workdays of different thermostats for different travel times. The data are for the unoccupied, well-insulated control home (see Table 2), assuming an 8:30AM-5:30PM job. The expected savings are adjusted by expected compliance rates.

Type (travel time)	Savings each workday (%)	Savings each workday (\$)	Expected compliance (%)	Expected savings (%)	Expected savings (\$)
Manual thermostat	24.7	2.16	35	8.65	0.756
Prog. thermostat	21.74	1.9	56.2	12.22	1.068
GPS therm (15 min)	5	0.38	90	4.5	0.342
GPS therm (26 min)	7.6	0.57	90	6.84	0.513
GPS therm (45 min)	11.75	0.88	90	10.57	0.792
GPS therm (60 min)	13.82	1.04	90	12.44	0.936
GPS therm (90 min)	17.05	1.28	90	15.35	1.152

Table 4. The expected savings for the control home when the GPS-Therm mode *augments* M-Therms and P-Therms for a 26 min commute and 90 min commute

Type (commute length)	Compliance (%)	Expected savings per workday (%)	Expected savings per workday (\$)
Manual thermostat (26 min)	35	8.65	0.76
Manual therm augmented with GPS therm (26 min)	100	13.59	1.13
Programmable thermostat (26 min)	56.2	12.22	1.07
Programmable therm augmented with GPS therm (26 min)	100	15.55	1.32
Manual thermostat (90 min)	35	8.65	0.76
Manual therm augmented with GPS therm (90 min)	100	19.73	1.59
Programmable thermostat (90 min)	56.2	12.22	1.07
Programmable therm augmented with GPS therm	100	19.69	1.63

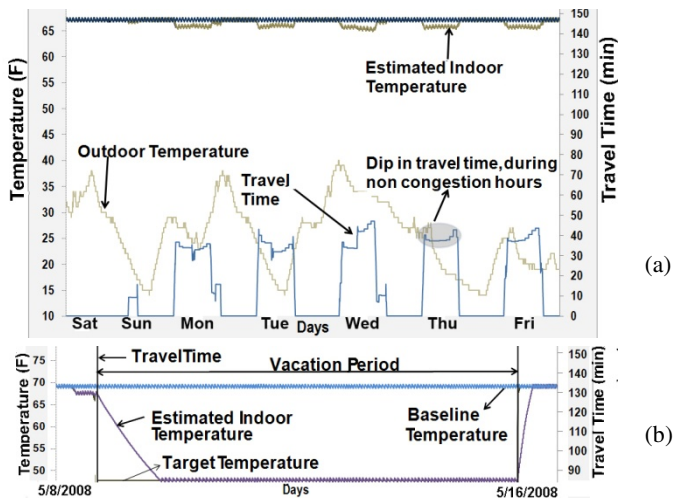


Fig. 2. (a) One week of real travel time data of a participant (House#1) and GPS-Therm simulation. (b) GPS-Therm simulation for a vacation period of participant of House#2

A key benefit of the GPS-Therm concept is that, unlike existing technologies, it adapts to changes in behavior *without requiring any behavior change on the part of the end user*. For the over 50% [4] of households not performing manual setbacks or using programmable thermostats properly, the GPS-Therm could provide a low-burden energy-saving option.

Table 4 shows “GPS add-on” expected savings. Here, expected savings are estimated using the programmable thermostat model for the 56.2% of the programmable thermostat owners who program it, but for the remaining 43.8% who do not program it, we assume that the thermostat defaults to the GPS model. Similarly, we assume that 35% of households that use the M-therms use setbacks but that the remaining 65% would default to the GPS system. For the control home with a 8:30 AM-5:30 PM job and commute of 26 min, the GPS-Therm add-on system would improve the overall expected performance of M-Therms by 4.9% and the overall expected performance of P-Therms by 3.3%. For a commute of 90 min, the savings jump to 11.1% and 7.5% respectively.

4.4 Simulation Using Real Travel Time Data

In the simulation above we assume that the person commuting always leaves and returns at the same time each day. To better evaluate potential savings, we used the real travel data from our 8 participants obtained from the GPS devices in their vehicles. Each house had two participants, so the simulator always used the minimum travel-to-home time of the two. The drive time given by the MapQuest server does not take into account the traffic congestions and delays in commute time during the different hours of the day. Therefore we increased travel times at each hour of the day proportionally to the traffic congestion index for Boston.

Figure 2 shows the travel data and GPS-Therm simulation for a typical week of one participant from House#1. The travel pattern of the participant is fairly regular throughout the work week with a small trip during a weekend. **Figure 3** shows the travelling patterns of the participants in House#2. The simulation algorithm uses min travel time because the system must always be able to condition the environment in time to reach comfort conditions for the closest person. For some homes this puts an upper bound on savings at a short (e.g., 7 min) commute. The figure areas that are




Fig. 3. Travel patterns for entire duration of the study for one household and the minimum combined travel time. V's mark vacation periods.

marked with “V” are the time periods where participants went for a vacation. If only one person in a house went for the vacation, the travelling time of the person staying at the home is considered. Travel time during vacations was manually entered into the dataset in cases where participants left their cars at home. For example, participants in House#2 went to London which is 8 h by airplane. Therefore we manually entered 8 h of travel time during the vacation.

Table 5 shows the simulation results for the entire study duration for all the four participating houses for P-Therms and GPS-Therms. The house details are given in **Table 2**. The baseline for calculating the energy and monetary savings is the cost of maintaining the comfort T throughout the study using no setbacks.² For the simulation, the P-Therm for all the participants is scheduled from 7 AM to 6:30 PM every day (and, to be conservative, including weekends). This interval was selected because all of the participants leave and arrive home at about this time. The target T (setback) for the P-Therm was set to 60°F (15.5°C). The effectiveness of the GPS-Therm fully depends on the travel patterns of the individuals and their home heating characteristics. For House#3, the GPS-Therm performs only 0.3% better than the baseline and a P-Therm with daily setback settings would be most effective. House#2 has larger savings because the house occupants had 2 vacations during the study totaling 12 days when the GPS-Therm automatically selects a very low target T that is sustained for the duration of the trips – a situation where the GPS-Therm excels, as shown in **Figure 2(b)**. Overall, however, when only considering savings and not return comfort, M-Therms or P-Therms are clearly preferable over the GPS-Therm when they are used properly. The savings for P-Therms are higher than that of M-Therms, because the occupants spent more time in the home than the programmable settings assumed (especially on weekends).

We know, however that compliance rates are low and that use of the GPS system can increase overall *expected* performance without requiring complicated programming or sacrificing comfort if augmented on top of M-Therms or P-Therms that are not being used. **Table 5** (bottom) shows the expected savings that might have been achieved in that case in the larger population for similar homes as the test homes.³ Savings range from 4.9% to 9.4% for GPS-augmented M-Therm to 7.9% to 12.2% for GPS-augmented P-Therm.

² Our baseline condition assumes that some of the 64% of U.S. manual thermostat owners who do not set them back regularly do not do so even when leaving for vacation, either due to lack of understanding, concern about plants, pets, or pipes, not desiring to return to a uncomfortable home, or simply forgetting to do so.

³ The households selected all have dual commuters and therefore the results represent savings that might be achievable for only that type of household. We fully expect, for example, that savings in households with stay-at-home parents and young children might be substantially less, because setbacks are often not appropriate. The GPS-Therm might actually be most convenient in those homes, however, where the occupants have highly variable travel time schedules that are rarely known in advance, and where the occupants may be less willing to tolerate a house that is uncomfortable upon return.

Table 5. Simulated energy and monetary savings for the entire duration of the study for all the participating houses, including energy and monetary savings when using the GPS-Therm to augment manual and programmable thermostats, based on actual commuting patterns and expected compliance rates. CT = comfort temp. LT = lower temp.

Thermostat	Savings for study duration % (\$)			
	House#1 - 75 days CT = 67, LT = 65	House#2 - 75 days CT = 69, LT = 67	House#3 - 60 days CT = 69, LT = 67	House#4 - 60 days CT = 69, LT = 67
Programmable	19.4% (\$168)	14.1% (\$102.80)	17.1% (\$387.50)	13.5% (\$244.30)
Manual	18.0% (\$165.90)	13.7% (\$97.30)	14.6% (\$332.90)	12.6% (\$225.70)
GPS	2.9% (\$25.50)	7.1% (\$49.70)	0.3% (\$7.50)	0.8% (\$15.50)
Expected savings for study duration % (\$)				
Manual defaults to GPS	8.2% (\$74.60)	9.4% (\$66.30)	5.3% (\$121.40)	4.9% (\$89.10)
Programmable defaults to GPS	12.2% (\$105.60)	11.0% (\$79.60)	9.7% (\$221.10)	7.9% (\$144.10)

4.5 Simulation Using Just-in-Time Questions

Using the GPS data, it is possible to improve the GPS-Therm mode by creating a system that benefits from modest user feedback without requiring the user to proactively remember to change the thermostat or predict schedules far in advance. Suppose when the user is detected to be away from home and not moving (i.e., just arrived at work), the system prompts with a simple question on the phone.

To estimate the savings this small interruption might enable, three additional (winter time) scenarios were simulated for the control house.

- Return at lower T: A user agrees to return at a T slightly lower than his or her comfort T. On the user's return, the home will be at a lower T, but the house will continue to heat up until it reaches the comfort T (see **Figure 4** (a)). In this scenario the system will have a lower target T and more energy savings during the day as compared to the T set automatically by the GPS-Therm.
- Specify a time to return home: If the user decides to return back at a specific time, the GPS-Therm will ignore the GPS data and operate like an intelligent thermostat that calculates the lowest possible target T (setback) and heating start time so that when the user returns, the house is at comfort T, resulting in a lower target T and more savings.
- Return at lower T and specify a return time: If the user agrees to return at a lower T and also specifies a return time, maximal savings are achieved (see **Figure 4** (b)).

Occasional questions presented on a mobile phone and *only asked when someone is away from the home* may be an effective way to gather energy-saving information with only modest burden. A properly-programmed P-Therm may achieve similar savings, but remembering to change schedules in advance when someone has a variable schedule may be a challenging task. To evaluate potential savings from a small amount of user input, we simulated expected savings results for two cases for the control house where questions are answered on 2 and 3 workdays on the mobile phone when the system detects that someone has left home. **Table 6** shows the results for the scenarios described above with a travel time each way of 26 min. In short, answering just 2 prompts per week can boost workweek savings by up to 3.6%.

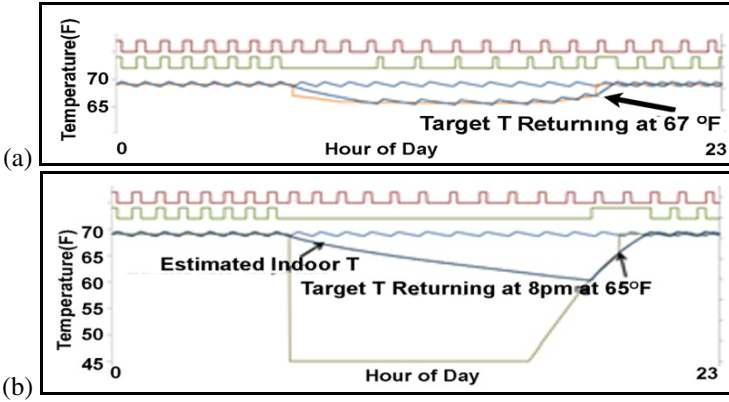


Fig. 4. 24h simulation for workday travel time of 26 min with GPS-Therm where user responded to a prompt to (a) return at lower T and (b) return at a lower T at a specified time

Table 6. Simulated savings for a work week for prompting scenarios with a travel time of 26 min

	Num prompts answered/wk	Expected savings(%)	Expected savings(\$)
Manual augmented with GPS	0	13.59	1.13
Manual augmented with GPS + Return time	2	16.91	1.37
Manual augmented with GPS + Return time	3	18.62	1.50
Manual augmented with GPS + Lower T	2	14.54	1.20
Manual augmented with GPS + Lower T	3	15.03	1.23
Manual augmented with GPS + Return time + Lower T	2	17.21	1.41
Manual augmented with GPS + Return time + Lower T	3	18.89	1.53
Programmable augmented with GPS	0	15.55	1.32
Programmable augmented with GPS + Return time	2	17.79	1.49
Programmable augmented with GPS + Return time	3	18.92	1.57
Programmable augmented with GPS + Lower T	2	16.19	1.37
Programmable augmented with GPS + Lower T	3	16.52	1.40
Programmable augmented with GPS + Return time + Lower T	2	17.95	1.51
Programmable augmented with GPS + Return time + Lower T	3	19.21	1.61

5 Design Observations from a Real-Time Implementation

To begin to assess the practical feasibility and usability issues that might arise with the proposed GPS-based travel-time mode, a fully-functional prototype was implemented in a participant’s house for 2 wk. Due to this study being conducted in warmer months, the system controlled air conditioning rather than heating. We report on some observations from this pilot deployment.

5.1 System Design

The back-end of the GPS-Therm prototype system is a client server model using TCP. The client is the GPS-enabled mobile phone (Motorola 9Qh Global) and the server is a laptop computer that was placed at the participant’s house near the location of the thermostat. The server receives the GPS coordinates from the client (via GPRS) once per

minute. It contacts the MapQuest web service to get the travel time and distance from the home of the participant. The server then contacts Yahoo weather web service [32] to get the outside T of that area. The algorithm calculates the target T based on the travel time and outside T using the heating/cooling tables measured in the home (using the process described in Section 4.2). Finally the server sends the new target T over a serial connection to the computer-controlled thermostat (RCS TR40 [33]), and in reply the server gets the confirmation and the current room T. The client, in reply to sending the GPS coordinates, receives the travel time, distance, current home T, and energy saving information related to the intervention questions. In some cases, the phone prompts for information, and the responses are sent to the server as well.

The prototype GPS-Therm system is divided into two interfaces: one on the laptop located in the house and the second on the mobile phone. The interface on the phone is minimalist. Nothing is displayed except when the phone detects that the phone user is over two-minutes (drive time) away from home after just having been there. In that circumstance, the phone beeps and a question is displayed on the phone's screen, which remains until the user has a chance to respond – typically on arrival at a destination. The prompt asks the user if he or she is willing return home to a 1-2°F warmer house and a return time (if known). In each case, the interface provides the user with the information on the savings expected when additional data are entered. If arrival at a slightly warmer house is selected, the interface indicates how much time it will take the warmer house to reach the desired comfort T upon return home. The prompts are easily ignored – the user is not forced to answer the questions.

The laptop interface in the prototype system is intended to simulate a replacement thermostat wall interface that would have a small digital display. It provides the user with system status information and full manual control. Unlike most current thermostats, it always displays what the system is currently trying to do. As shown in **Figure 5**, the system displays (A) the current home T, (B & C) the current state of the system, (D) why that is the state, and (E) what the user should do if he or she wants to change the T settings. The interface provides the energy and monetary savings since the system was installed (F). It also provides control to manually change the comfort settings (L). It rewards the participant for making energy savings decisions by

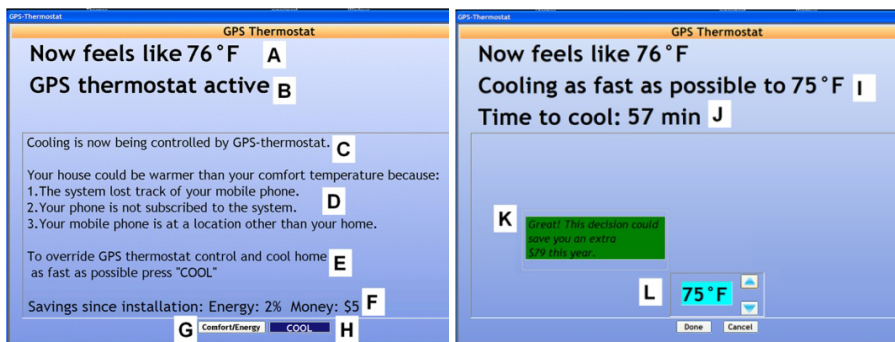


Fig. 5. Screenshots of the laptop thermostat interface (simulating a wall thermostat LCD)

showing monetary tradeoffs in real-time (K) and how the system is going to react to the change (I&J). There is a “COOL” button that can override the GPS control and resume the manual control of the cooling system (H). **Figure 6** (left) shows a typical scenario that may occur when the participant returns to a warmer home and system is not able to track his or her mobile phone.

5.2 Participant

A working professional (realtor and musician) living in a single family house near Boston with a central cooling system was recruited. The study protocol was approved by our IRB and the participant had no affiliation with the researchers.

The participant used his car as the main mode of transport during the 14 d study and lived alone in a stand-alone house with a cooling area of 3000 ft² (278.7 m²). At the time of the experiment, the house was approximately 40 years old, and the cooling system was approximately 23 years old. The cooling capacity of the compressor was 60,000 BTU/h and the whole system was controlled by one M-Therm. The participant had a pet but he mentioned that he left his pet in the basement whenever he was away from the house, and the basement was not included in the area cooled by the air conditioning system. The comfort T of the participant was 74°F (23.3°C). The travel pattern of the participant was irregular because he sometimes worked from home.

5.3 Experimental Setup

A professional electrician installed the computer-controlled thermostat (RCS TR40) in the participant’s home. The laptop computer (server) was kept underneath it on a table with the screen clearly visible. The participant had a broadband wireless Internet connection, which was configured so that the client (mobile phone) and server could communicate. The participant was given a GPS-enabled mobile phone (Motorola 9Qh) running the software continuously and was asked to use it as his personal phone (moving his SIM card). The participant was told to recharge the phone every night. A GPS logger was also installed in the participant’s car. The remaining procedures were the same as those described in Section 4.3, although adapted for cooling rather than heating.

5.4 Evaluation

The same procedures as described in Section 5.2 were used to estimate savings that the GPS-Therm add-on could have achieved for the 2 wk period of the pilot study. The baseline used was the comfort T of the participant (74°F (23.3°C)). The primary goal of the pilot experiment was to gather qualitative feedback on the practicality and usability of a fully-functional system.

5.4.1 Energy Saving Estimations

Under normal circumstances, this participant said “*I never change my setback temperature.*” Based on the participant’s home cooling characteristics and his travel patterns as obtained from the mobile phone, our simulations showed that by running the GPS-Therm prototype for two weeks he therefore saved an additional 3.4% and \$2.70. The impact of using the phone’s GPS versus a GPS logger in the car was




Fig. 6. Travel patterns of the participant for 2 wk. Highlighted portions show the difference between travel data from the phone’s built-in GPS and GPS logger in the car, which are modest.

Table 7. Expected savings across similar homes and commute patterns for two weeks

Thermostat	Expected Savings (%)	Expected Savings 2 wks (\$)
Manual	1.05	0.84
Manual augmented with GPS	3.3	2.60
Programmable	5.84	4.64
Programmable augmented with GPS	7.3	5.80

minimal. As indicated in **Figure 6**, an additional savings of only 0.2% would have been achieved had the phone GPS worked equivalently to the Trackstick Pro in the car.

As in previous simulations, we used this participant’s travel patterns and his home’s temperature response characteristics to estimate the expected savings across others with homes and travel patterns like his, assuming the M-Therm and P-Therm compliance rates. **Table 7** shows the expected savings possible for the two week period. Although programmable systems would save the most energy,

5.4.2 Responding to Prompts

During the 14 d study, the participant answered 8 of 24 prompts he received on occasions when he left his home, and he entered some information that led to energy savings in each case. During 3 of his trips, for example, he responded positively to the “return at warmer temperature” and the “specify a return time” questions. Due to the novelty effect, we hesitate to make generalizations about the question-answering. However, we *can* measure the energy saving impact, and each answer he gave saved the participant an average of 0.32%, or \$0.25, in energy when the thermostat was in GPS mode. For some people, this amount earned may offset the burden of the interruption.

6 Discussion

The results from our participants using both simulation and the functional system suggest that a GPS/travel-time thermostat mode could save substantial amounts of energy. Here we list issues that may need consideration if such a system were to be implemented at scale.

Issue 1: Aesthetic concerns and impact on design. Most thermostats are wall mounted at eye-level, far from either an electrical outlet or an Internet outlet. Home occupants may be unwilling to introduce unsightly cables running to the device or invest in costly, messy renovation. To encourage widespread adoption of a thermostat with a GPS mode, the system must (1) be powered by the 24v line running to every thermostat, and (2) connect to the Internet via a wireless link. The extra energy consumed by the continuous wireless link (e.g., 802.11) must be factored into the potential savings. A typical wireless router consumes approximately 0.3KWh/d, which in the Boston area would cost \$0.06/d and \$22/year, far less than the GPS-mode savings for many people.

Issue 2: Cost. In our prototype, the server was implemented on a laptop, but mobile phones with data connections are sufficient to run the simulations. By leveraging the phones that people will buy anyway, the only new functionality that must be added to a thermostat is a wireless link so the thermostat can be set and read remotely. In homes similar to our participant houses where people do not use setbacks, the technology could pay for itself quickly. On average, winter month savings would range from approximately \$45 to \$120 in our area. Considering *only* winter months, this leads to estimated payback times from 12 mo for House#2 to 36 mo for House#3. Without knowing how much he could have saved, the participant using the prototype reported that he would pay up to \$300 for the device that he tested, an amount adequate to cover the cost of the device.

Issue 3: Unanticipated consequences in behavior. Our participant using the fully-functional system mentioned that his work travelling patterns are irregular and that he did not use setbacks, but that the GPS mode was beneficial because, "I don't have to remember to change my temperature settings before leaving." However, if long term he relied on GPS mode instead of using setbacks for his regularly scheduled trips, he would waste energy. If the system detects someone is relying on GPS mode, the interface should gradually introduce interactive prompting with the goal of encouraging use of the P-therm or gathering information every day that allows the system to operate at the same efficiency. The unintended consequence of someone who otherwise would have adjusted the thermostat instead relying on the GPS mode must be discouraged.

Issue 4: Time away: inference or prompting? Due to the (long) lag time of heating and cooling a home, maximal savings can only be obtained when return time is known or accurately inferred. Otherwise, many people during a normal workweek will work so close to home that the interior T cannot drop/raise sufficiently fast to accrue major savings because return time is so fast. An alternative is to infer typical return-time behavior automatically from prior behavior. Inferring intentionality may be valuable for energy savings [7], but it is also prone to error. An alternative that may lead to a more predictable user experience would be to simply recognize one particular behavior – when someone has stopped traveling -- and then present an easy-to-ignore prompt on the phone. The burden of the interruption could be softened by presenting real-time, tailored information about the savings that the interruption enables. Our simulations suggest that the daily value of answering a question for many homes could be as high as \$.25. The participant, "liked the persuasive status messages, because they were showing how much I can save," but even a single extra reward message screen generated the comment that, "the message was adding extra

information to an already cluttered day.” Is \$.25 savings enough to offset the inconvenience of the interruption? Determining prompting sustainability requires longer-term studies with functional systems.

Issue 5: Explanation of thermostat state and behavior. The GPS-Therm prototype always indicated the current T, target T, and time to reach target T, unlike most thermostats on the market today, to discourage the user from falling into the common thinking that a higher setting leads to faster T change [19]. The participant using the pilot implementation did comment that, “the laptop interface gave me information about what the system was doing and why, which I cannot see on my thermostat.” Unfortunately, adding a GPS/travel-time mode creates new special case conditions that can be challenging to convey. For example, if a person leaves his or her phone at work or the phone battery dies, the home will be quite uncomfortable on return; the reason (that the system thinks the occupant is not home) must be conveyed to the user. Handling these special cases must be accomplished without losing the desirable simplicity of the GPS-Therm mode. Here too, longer term studies are warranted.

Issue 6: Incremental prompting frequency tailored based on prior answers. Our pilot interface could put the user in a position where a prompt for savings or return time information must be declined on a regular basis for reasons beyond the user’s control, thereby repeatedly creating disconcerting feelings of cognitive dissonance if the person wants to save energy. With respect to the prompts, the participant commented that, “it made me feel good if I make energy saving decisions, and feel bad if I did not.” To avoid creating negative feelings, the system could be improved by tailoring prompting rates based on previous frequency of positive responses.

Issue 7: Temperature vs. comfort. Our participants were accustomed to thinking in terms of “temperature” rather than comfort. The participant using the prototype, for example, reported about the minimal phone interface that intentionally did not show temperature that, “I did not find the GPS thermostat phone interface informative [because] it was not showing me the temperature of my house.” He further added that the phone should display current T, target T, and energy savings, and the ability to change the T or return time decision at any point of time. However, particularly in summer, the humidity can have an impact on T, as can other factors, such as what one is wearing. Ideally an interface would guide the user to focus on comfort *instead* of T, allowing more fluctuation in T (and therefore more energy savings).

Issue 8: Phone limitations. GPS lock times on the phone ranged from 2-5 min on sunny days with open skies to 15 min on cloudy days in urban areas. Fortunately, the ongoing massive industry investment in improving phone location-based services will only further improve the performance of the GPS-Therm mode. Despite current limitations, however, as mentioned in Section 5.1, the phone still performed adequately to produce savings.

Another consideration for practical deployment is battery life. Using the phone data transfer scheme described in Section 6.1 resulted in a battery life on the MotoQ 9h global of only 10-12 h on a charge. To improve the battery life, we subsequently devised a simple scheme for pilot testing where the GPS switches on every 3 min and remains on for 2 min to get a lock, which enabled 24 h performance.

Other issues. As we ran these experiments, we identified several other areas where our prototype could be improved that we mention briefly: (1) implementing an algorithm which dynamically updates the lookup tables for the T profile of the house

based on the current weather conditions like outside T, humidity, wind velocity, and sunshine, (2) implementing algorithms that take into account not only outside T but also forecasted outdoor T, so that a sudden change in weather condition does not create a situation where the GPS-Therm cannot catch up in time to ensure the home is at a comfortable T when the occupants return, (3) detecting “driving” from the GPS data in order to minimize the chance of an ill-timed prompt, (4) modification of the question-prompting to handle multiple participants in the house, including exchange of information between family members as they make decisions that might impact each others comfort, and (5) controlling the hot water heater. We have not considered night setbacks in our simulations, which may also be amenable to this type of interactive control.

7 Conclusion

In summary, in this pilot work we have prototyped, tested and evaluated a GPS-controlled thermostat system. We have shown through simulations that such a system is capable of saving as high as 7% on HVAC energy use in some households. This is less than savings obtained from optimal use of M-Therms or P-Therms, but the GPS mode we propose has the potential to save energy for the more than 50% [4] of the U.S homes that do not change their T settings when there is no one in the house, and the system could be easily and affordably installed in many homes. More work is needed to fully evaluate potential savings and feasibility and usability of the user interface and interactive prompting components and to explore differences in climate, living environments, and lifestyles in other parts of the world.

Acknowledgements

This work was funded by the MIT House_n Consortium. The authors thank the reviewers for helpful comments, Telespial Systems for Trackstick donations, and our participants.

References

- [1] CIA, The World Fact Book (cited 2008 08/15/2008), <https://www.cia.gov/library/publications/the-world-factbook/>
- [2] U.S. DOE Annual Energy Review (cited 2008 08/15/2008), <http://www.eia.doe.gov/fuelelectric.html>
- [3] U.S. DOE Building Energy Data Book (cited 2008 08/15/2008), <http://buildingsdatabook.eren.doe.gov/docs/1.2.3.pdf>
- [4] U.S. DOE Residential Energy Consumption Survey (cited 2008 08/15/2008), http://www.eia.doe.gov/emeu/recs/recs2005/hc2005_tables/detailed_tables2005.html
- [5] Nevius, J., Pigg, S.: Programmable Thermostats that Go Berserk? Taking a Social Perspective on Space Heating in Wisconsin. In: Proc. ACEEE Summer Study on Energy Efficiency in Buildings, vol. 8, pp. 233–244 (2000)

- [6] Plourde, A.: Programmable Thermostats as Means of Generating Energy Savings: Some Pros and Cons. Canadian Building Energy End-Use Data and Analysis Centre, Technical Report CBEEEDAC 2003-RP-01 (2003)
- [7] Harris, C., Cahill, V.: Exploiting User Behaviour for Context-Aware Power Management. *Proc. Wireless And Mobile Computing, Networking And Comm.* 4, 122–130 (2005)
- [8] Ross, J.P., Meier, A.: Whole-House Measurements of Standby Power Consumption. In: *Proc. 2nd Int'l Conf. on Energy Efficiency in Household App. and Lighting*, vol. 108(13) (2000)
- [9] Roy, A., Bhaumik, S.K.D., Bhattacharya, A., Basu, K., Cook, D.J., Das, S.K.: Location Aware Resource Management in Smart Homes. In: *Proc. IEEE Int'l Conf. on Pervasive Computing and Communications*, pp. 481–488 (2003)
- [10] Harle, R.K., Hopper, A.: The Potential for Location-Aware Power Management. In: *Int'l Conf on Ubiquitous Computing*, pp. 302–311 (2008)
- [11] Thermostat History (cited 08/15/2008), <http://www.prothermostats.com/history.php>
- [12] Keyson, D.V., de Hoogh, M.P.A.J., Freudenthal, A., Vermeeren, A.P.O.S.: The Intelligent Thermostat: A Mixed-Initiative User Interface. In: *Proc. Conf. on Human Factors in Computing Systems*, pp. 59–60 (2000)
- [13] Fountain, M., Brager, G., Arens, E., Bauman, F., Benton, C.: Comfort Control for Short-Term Occupancy. *Energy and Buildings* 21, 1–3 (1994)
- [14] Mozer, M.C., Vidmar, L., Dodier, R.H.: The Neurothermostat, Predictive Optimal Control of Residential Heating Systems. In: *Advances in Neural Information Processing Systems*, vol. 9, pp. 953–959. MIT Press, Cambridge (1997)
- [15] Titus, E.: Advanced Retrofit: A Pilot Study in Maximum Residential Energy Efficiency. In: *Proc. ACEEE Summer Study on Energy Efficiency in Buildings*, vol. 1, pp. 239–245 (1996)
- [16] Fuji, H., Lutzenhiser, L.: Japanese Residential Air-Conditioning: Natural Cooling and Intelligent Systems. *Energy and Buildings* 18, 221–233 (1992)
- [17] Springer, D., Loisos, G., Rainer, L.: Non-Compressor Cooling Alternatives for Reducing Residential Peak Load. In: *Proc. ACEEE Summer Study on Energy Efficiency in Buildings*, vol. 1, pp. 319–330 (2000)
- [18] The Lawrence Berkeley National Laboratory. Thermostats and Comfort Controls (cited 2008 08/14/2008), <http://comfortcontrols.lbl.gov/>
- [19] Kempton, W.: Two Theories of Home Heat Control. *Cog. Science* 10, 75–90 (1986)
- [20] Gladhart, P., Weihl, J.: The Effects of Low Income Weatherization on Interior Temperature, Occupant Comfort and Household Management Behavior. In: *Proc. ACEEE Summer Study on Energy Efficiency in Buildings*, vol. 2, pp. 43–52 (1990)
- [21] Sachs, H.: Programmable Thermostats. American Council for an Energy Efficient Economy Technical Report (2004)
- [22] Analytics, R.: Validating the Impact of Programmable Thermostats: Final Report. RLV Analytics, Inc., Middletown, CT (2007)
- [23] Karjalainen, S., Koistinen, O.: User Problems with Individual Temperature Control in Offices. *Building and Environment* 42, 2880–2887 (2007)
- [24] Mozer, M.: The Neural Network House: An Environment that Adapts to its Inhabitants. In: *Proc. AAAI Symposium on Intelligent Environments*, pp. 110–114 (1998)
- [25] TrackStick (cited 2008 08/15/2008), <http://www.trackstick.com/>

- [26] LASCAR Logger (cited 2008 08/15/2008),
<http://www.lascarelectronics.com/>
- [27] LOGiT Current and Voltage Data Logger (cited 2008 08/15/2008),
<http://www.supco.com/LOGiT%20Data%20Loggers.htm>
- [28] U.S. Env. Protection Agency Energy Star Programmable Thermostats Specification (cited 2008 08/15/2008),
http://www.energystar.gov/index.cfm?c=revisions.thermostats_spec
- [29] Mull, T.E.: HVAC Principles and Applications Manual. McGraw-Hill Professional, New York (1997)
- [30] MapQuest Webservice (cited 2008 08/15/2008),
<http://developer.mapquest.com/>
- [31] U.S. Census Bureau (cited 2008 08/15/2008), <http://www.census.gov/>
- [32] Yahoo Weather Webservice (cited 2008 08/15/2008),
<http://weather.yahooapis.com/forecastrss?p=ZipCode>
- [33] Residential Control Systems (cited 2008 08/15/2008),
<http://www.resconsys.com/products/stats/index.htm>