

# Demand Response-Enabled Residential Thermostat Controls

*Xue Chen, Jaehwi Jang, and David Auslander, Mechanical Engineering Dept., UC Berkeley  
Therese Peffer and Edward Arens, Architecture Dept., UC Berkeley*

## ABSTRACT

A number of Demand Response (DR) technologies work by responding to variable electricity pricing, but have not yet been applied to control residential HVAC systems. An autonomous thermostat system, the Demand Response Electrical Appliance Manager (DREAM), provides possibilities to improve price-based demand responsiveness in residences.

Built on low-cost, low-power wireless technology, the system uses a disaggregated set of energy- and environmental sensors. Control strategies are implemented to optimize electricity cost and user's comfort. To perform the optimization, the system starts from default values and learns the dynamic behavior of a house and HVAC system. A graphic user interface provides easy interaction with the system. Computer simulation, lab tests and field tests have been used to validate the system infrastructure and control strategies. These tests indicate that the DREAM responds automatically to price signals with appropriate energy saving behavior. The system can reduce electricity consumption during peak price hours without significantly decreasing comfort.

## Background

The State of California is moving toward managing electricity use during periods of shortages, as caused by curtailment of supply or high demand. Measures to reduce the demand for electricity during such shortages are termed 'demand response' (DR). Demand response measures include load reduction, demand bidding, and variable price rates with some mechanism for utility customers to respond to those rates. Demand response measures have the effect of adding elasticity to the electricity market. It has been estimated that a mere 2.5% reduction in demand in response to shortages can reduce the price spikes by 24% (EPRI 2002).

As part of the effort to increase demand responsiveness, the California Energy Commission is currently constructing a new policy to require DR-enabled thermostats for new residential construction in California. We looked at the requirements of the new thermostat and problems with existing technology to develop a smart, adaptive, integrated demand-responsive residential control system.

## Introduction

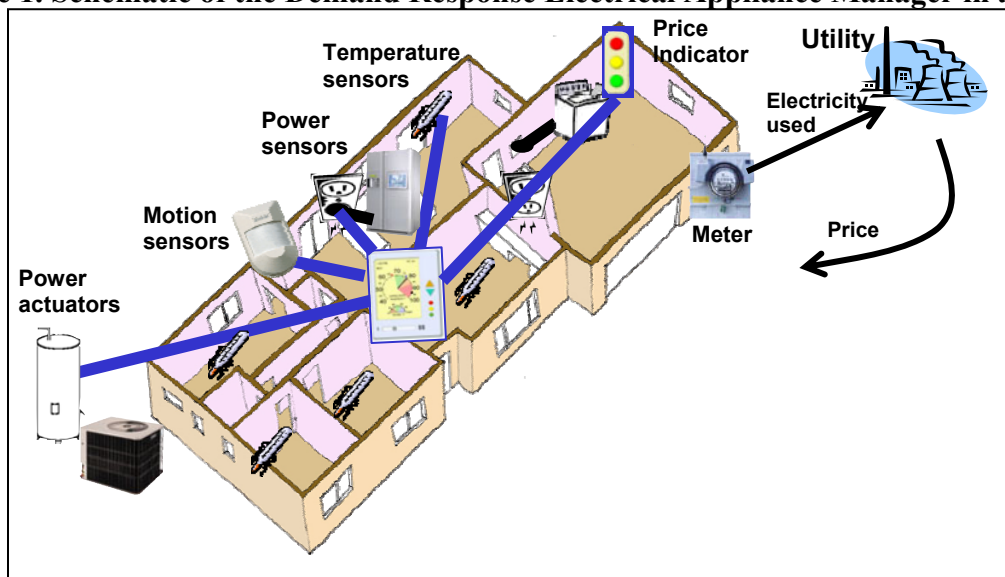
A number of issues need to be addressed before demand-response systems can be effectively deployed on a wide scale in the residential sector. The first is the infrastructure—the meters, communication, and responsive controls, such as a DR-enabled thermostat. Cost both to purchase and install and usability are other issues. A DR-enabled thermostat must automatically respond to price signals so that the homeowner is not forced to be a "day trader" in electricity. Homeowner's preferences for cost versus comfort must be easily transferred to the thermostat.

In order to achieve the goal of a ten-fold increase over existing thermostat functionality, we must understand the problems with the existing technology. The main issue is adoption: if

people do not accept the technology or use it as designed, then it will not achieve the energy savings objective. Approximately half of the houses in California that have thermostats have programmable thermostats (CEC 2004). However, a Carrier study estimated that 35% do not use the programming features, but put the thermostat in hold mode and operate manually (Archacki 2003). Adding DR functionality raises the level of complexity because a variable price schedule overlays the occupants' existing patterns of air conditioning and heating. Therefore, simplicity of operation is essential. The new DR systems have to behave autonomously based on effective initial defaults and machine learning. It needs to work right out of the box with no programming required in order for it to operate well. The user interface has to be intuitive. It should help people manage their energy use within a variable-price context (a new concept for most people).

A low-cost, autonomous, demand-responsive electrical appliance manager (DREAM) that exploits wireless sensor network technology, new smart control algorithms and a graphic user interface is being developed to address these issues. Figure 1 shows the DREAM concept schematic: 1) the system receives an electricity price from the utility that varies over time based on overall demand and supply; 2) the controller receives data from the wireless temperature, motion, and electrical current sensors, and controls appliances such as the air conditioner via wireless actuators; 3) the occupants determine their usage based on their economic, comfort, and convenience preferences. An intuitive interface receives user input as well as informs the occupants of the price and their current electricity consumption; and 4) the electricity usage is monitored frequently and relayed back to the utility. The DREAM consists of three parts: user interface, hardware, and control.

**Figure 1. Schematic of the Demand Response Electrical Appliance Manager in a Home**



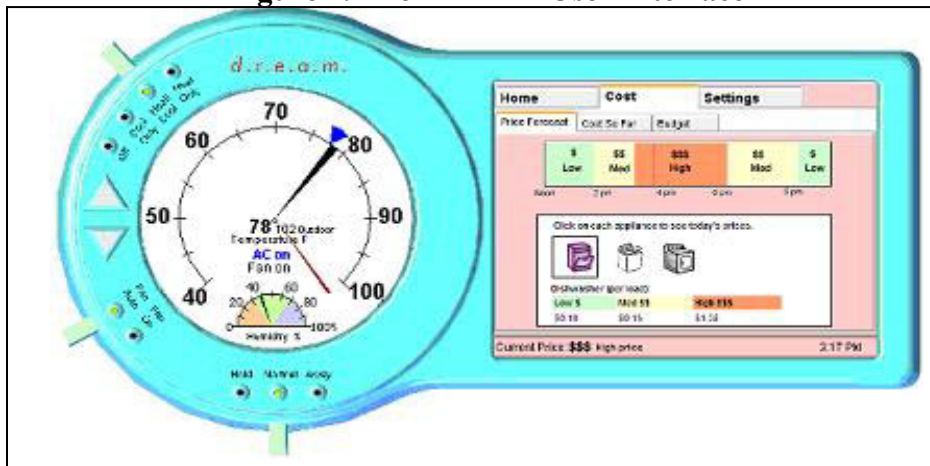
## System Design

### User Interface

The goal of the user interface is to display information to the user and allow input. The basis for the design is to help users understand the DR concept, learn how the system works, and

manage electricity use efficiently. Figure 2 shows the user interface of DREAM. The left of the interface was modeled after the Honeywell Round thermostat, which has great market penetration. The right side of the interface is designed as a touch screen “file folder” display, where users could see system messages sent by utility or cost information, electrical usage in total and for each appliance, and program schedule and temperature settings (Peffer et al. 2005).

**Figure 2. The DREAM User Interface**



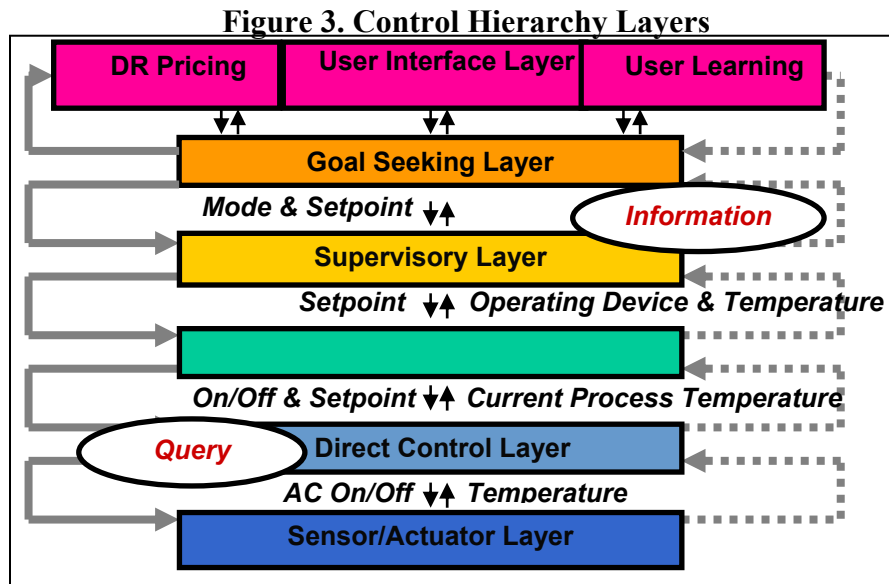
## Hardware

Coupled with the broader control functionality, there is a need for more energy and environmental information via sensing. Useful sensing metrics include temperatures throughout the house, outside weather conditions, occupant sensing, and power usage. These combine to allow for more targeted control and to be able to deliver predictable behavior and energy savings to the occupants. The system we are designing is thus far more information-rich than current thermostats and has extended command capability.

The system design, then, is driven by the need for distributed sensing and actuation in a system that can be implemented for a reasonable cost. Because the cost of wiring is usually the gating factor in systems requiring distributed sensing and actuation, a major enabling technology for this system is low-power, low-cost wireless communication. Each node in this wireless system, a so-called mote, has a low-power computer or microprocessor, a low-power radio transceiver, and multiple analog or digital input/output channels for sensing and actuation. Power is a major issue because an important design goal is ten years of operation without any needed scheduled maintenance, such as battery replacement. The system architecture uses a central controller and the wireless equivalent of a “star” network for connectivity to distributed motes. There is a base station connected to the central controller and one repeater strategically located elsewhere in the house. Motes are capable of the more general mesh networking but, thus far, that has not been necessary to implement this added level of overhead to improve communication range and reliability.

We replaced the household thermostat with a mote connected to a set of relays that actuated the air conditioner compressor, furnace, and blower fan. While many devices might be controlled under the DREAM scenario (lab tests coordinated the actuation of a ceiling fan in conjunction with air conditioning), the main actuation in our tests was this HVAC relay.

## Control Structure



The addition of demand responsiveness and whole house control to basic thermostat functionality leads to a control system with considerable complexity. In order to handle that complexity, we have adopted a layered design for the control system software (see Figure 3) (Auslander, Ridgely & Ringgenberg 2002). In a layered design, each layer (in theory) interacts only with the layers above and below it. This provides for modularization of function and semi-independent design of each layer. The lower part of the hierarchy describes basic control functions used to maintain temperature in the house and other functions (such as turning on other appliances). The lowest layer, the Sensor/Actuator Layer, maintains communication between the controller and the nodes with sensors and actuators. Thus, the function of the lowest levels is similar to a conventional thermostat—to manipulate the HVAC system for thermal control. The most complex layer is the Goal Seeking Layer. It receives price information and must make decisions about how to best compromise comfort and cost. In the middle layers, choices must be made as to how to meet the compromise decided on by the Goal Seeker. In many cases, there are choices as to how to achieve the goal; for example, for cooling, one might use a whole house fan, air conditioning, and/or ceiling fans. Embedded in the hierarchical structure, system functions are designed to improve demand responsiveness and optimal control performance.

## Control Algorithms

Our hypothesis is that if the new thermostat is autonomous, that is, could work well right out-of-the-box, then the technology would be more acceptable to the occupants of the house. Upon installation, the thermostat would immediately begin adapting control strategies to the specific HVAC system, house parameters, climate, and price. Towards that end, we have developed a learning algorithm to predict thermal behavior of house and HVAC system. We have also developed optimization strategies that look at the prediction, the data from the house and occupants, and electricity price in order to determine temperature setpoint.

## Optimization

The DREAM is designed to reduce electricity loads during hours of peak demand and minimize thermal discomfort caused by such reduction. The mechanism to drive the system is the DR signals—presumably price—set by utility companies. From the users’ point of view, the goal of the DREAM system is to maintain the users’ comfort for the minimum electrical energy cost. Therefore, how well the system compromises cost and thermal comfort is the key of DREAM.

The decision-making process takes three steps to determine a goal temperature. The first step is to choose a cooling/heating mode based on season and temperature trend. The second step is to adapt different control strategies to deal with temperature requirements varying with seasons, occupancy status, and price changes. The final step is to decide the temperature setpoint by minimizing a utility function consisting of electricity cost and thermal discomfort.

Four system states are defined for implementing different control strategies (Table 1). Although the objectives of each state are different, every state considers two factors: cost and occupant’s thermal comfort. These are somewhat competitive – less cost means less comfort and greater cost means more comfort. In each strategy, the trade-off between the two is evaluated for different time periods. For example, the normal state considers the steady state trade-off while the precool/preheat state focuses on the transient process during the period before and after a price increase. The last three strategies evaluate the temperature trend and estimated electricity consumption, which are obtained by learning the house characteristics.

**Table 1. Control Strategy Design**

Strategy Name	Objective
Normal	Optimize <i>cost vs. comfort</i> when there is no future price increase or predicted arrival/departure
Precool/Preheat	Save money by <i>shifting load</i> in anticipation of a future price increase period (pre-cool)
Departure/Arrival Preparation	<i>Save energy</i> if departure is predicted <i>Set comfort</i> temperature if arrival is predicted
Combination	Prepare for <i>both</i> future price increase and predicted departure/arrival

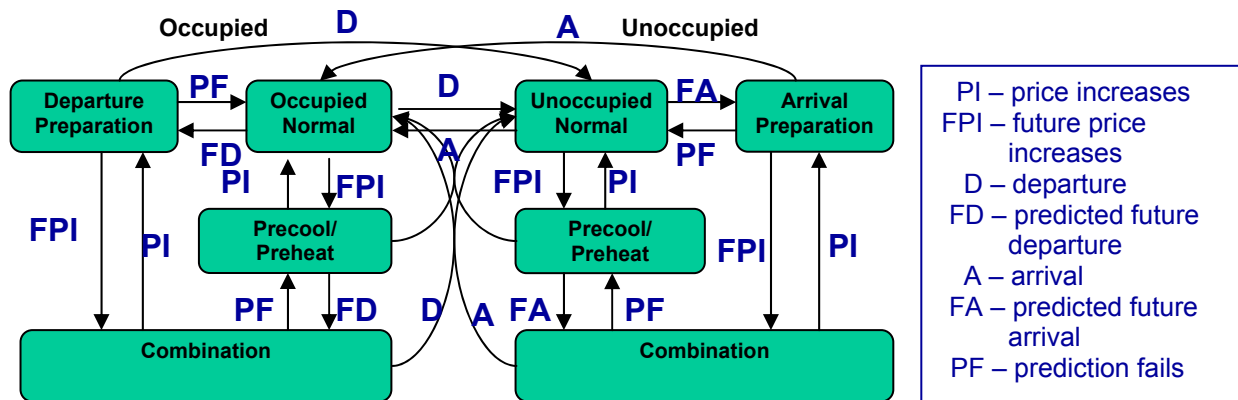
To evaluate residential thermal comfort, we used the Adaptive Comfort Standard (ACS), defined in ASHRAE Standard 55-2004 (ASHRAE 2004). It is designed for office buildings with natural ventilation, and adapts the comfort temperature range according to mean outdoor temperature. The comfort is denoted by the percentage of acceptance for a temperature range.

Since cost and comfort do not use the same unit, quantitative optimization required creating a common currency. We scaled both energy cost and comfort level in percentages, by dividing current values by the full range. Additionally we realize that users have different preferences of comfort and cost. For example, during high price periods, some people would pay more to maintain comfort while some may sacrifice comfort to save money. To customize the optimization, tools are needed to accept the users’ preference. The proposed candidate is called the “economics index”, as a user-specified term used in the utility function. Ranging from 0 to 1, it equals 1 when users would like to maintain 100% comfort without considering price. It is 0 when only minimum comfort is maintained; the users would like to keep the cost to no more than they would have spent if there were no increase in price. The default value 0.5 indicates a

common case: users are sensitive to price when price changes from medium to high so that setpoints are adjusted moderately. To help users understand the use of the economics index, the corresponding comfort level and cost changes are shown. Users' acceptance of this method is still under investigation.

The state/strategy transitions are driven by current and future events of price and occupancy. Figure 4 below shows an event-based state transition diagram. Every arrow indicates an event-enabling state change. For example, a future-price-increase event triggers the transition from the occupied normal state to the precool/preheat state. Predicted occupancy events are generated by users' daily schedule and future price events are delivered by the electrical utility.

**Figure 4. Event Based Strategy Transition**



## House Learning

While the controller's functions get quite complex, the controller must operate in an unmanaged environment and be easy and inexpensive to install. This is in contrast to controllers for commercial and industrial buildings, where a substantial amount of effort is spent on installation and tuning, and the system is professionally managed. For this reason, the "learning" function is among our key enabling technologies. Optimal control must operate well in a wide variety of physical environments—large houses, small houses, well-insulated, poorly insulated and so on—and adjust its operation to the local conditions. This adjustment must be entirely transparent to the occupants.

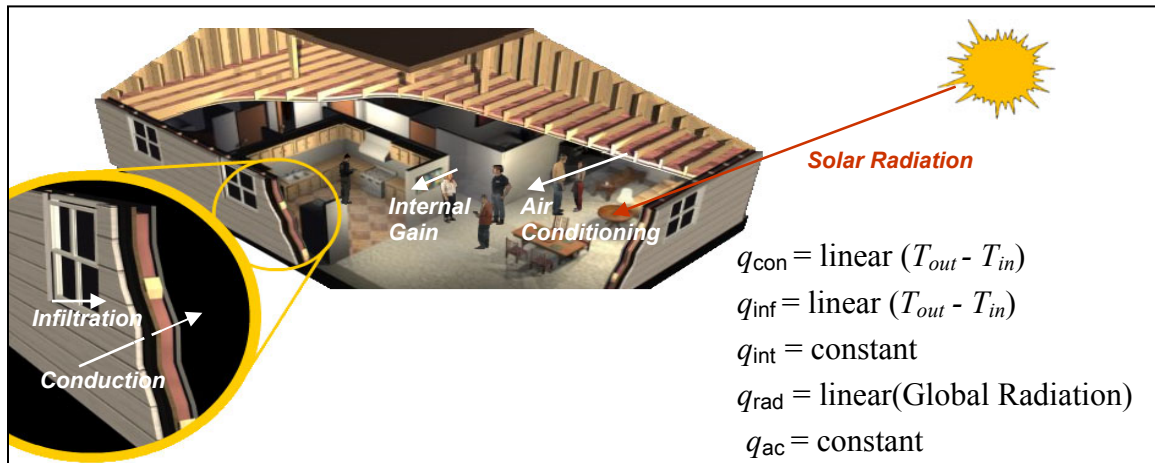
We proposed a first order time-invariant model to estimate and predict temperature trend and electricity consumption. As indicated in Figure 5, five sources of heat transfer that affect indoor temperature are considered: conduction, infiltration, internal gains, solar radiation, and air conditioning in summer or heater in winter. Conduction and infiltration are proportional to the temperature difference between outside and inside. So the corresponding heat flow is expressed as a linear function of temperature difference. Although internal gains due to people, lights, and equipment fluctuate daily, these influences are usually much less than the other sources of heat transfer and thus it is reasonable to assume this is constant. The temperature changes due to solar radiation depend on the size and orientation of windows as well as the structures around a house that block or reflect radiation; both are fixed. Then it is reasonable to linearly correlate temperature changes and radiation, which depends on time of day and time of year. Finally, we assume that the capacity of AC and heater remain constant. Let  $\alpha$  denote conduction and

infiltration rate,  $\beta$  denote internal gain,  $\gamma$  denote the dependence of radiation, and  $\delta$  denote the capacity of AC or heater. The following formula is proposed as the predictive house model:

$$\frac{VHC \cdot (T_{in}(t + \Delta t) - T_{in}(t))}{\Delta t} = \alpha \cdot (T_{out}(t) - T_{in}(t)) + \beta + \gamma \cdot (\text{radiation}) + \delta \cdot (\text{HVACstatus})$$

where VHC represents volumetric heat capacity.

**Figure 5. Five Sources of Heat Transfer for House Modeling**



Source: [www.ce.utexas.edu/bmcb/scenarios/heatingCooling.cfm](http://www.ce.utexas.edu/bmcb/scenarios/heatingCooling.cfm)

At the beginning of the system operation, a set of defaults meant to represent the majority of California houses provides a “reasonable” prediction. Through data acquisition and analysis, these defaults would be updated with learned parameters that represent a better description of the house and thus achieve better prediction. To tune the parameters for a specific house, data are clustered with respect to radiation conditions and air-conditioning status. Data obtained under no radiation and no air-conditioning (usually during night times) are used to tune  $\alpha$  and  $\beta$ . After the first two parameters are tuned, the next two parameters can be tuned with the data from radiation and air-conditioning effects in turn.

## Tests and Results

The choice of a hierarchical structure of the control software enabled seamless testing of the DREAM via various methods. We tested the functions and infrastructure of the DREAM system via a simulation tool, in a controlled setting at the university and a researcher’s home, and finally in the field. As a convenient and efficient testing method, computer simulation allowed us to easily change different aspects of the system in validating the design hypotheses and evaluating the performance of control strategies. The initial physical test of the hardware, software, and communication occurred in a controlled laboratory environment. Finally, testing in an occupied house helps demonstrate how the DREAM system works in the real world.

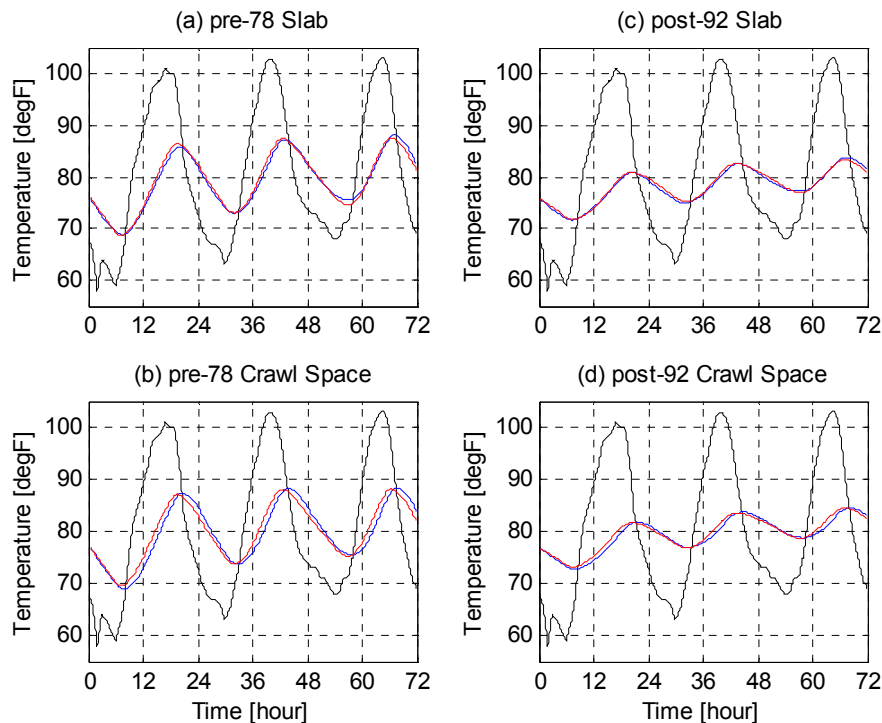
## Simulation

We built a simulation tool, named Multi-Zone Energy Simulation Tool (MZEST), based on a version of the California Non-Residential Engine to evaluate our control strategies. The

controller interfaces with MZEST in an iterative (5-minute time step) loop. The controller provides the on/off signals of the HVAC equipment and MZEST provides multi-zone temperatures based on house parameters that we specified. Additionally, a price generator was developed to simulate DR signals from the utility. The electricity rates were generated based on real-time electricity demand, which has strong correlation with outdoor weather conditions.

**Validating house learning with simulation.** We approximated the spectrum of California houses using four construction types. Two houses were modeled with relatively little insulation, representing houses built before 1978, one with a slab-on-grade foundation and the other a crawl space. The other two were modeled using the energy standards of 1992 (insulated envelope and double-paned windows), also one with slab-on-grade foundation and the other crawl space. These four types are constructed within MZEST and the energy consumption simulated. Each of the house models were set up using the learning algorithm proposed earlier. To validate the learning algorithm, we compared the simulated indoor temperature and the one predicted by the five-parameter house model. Figure 6 shows the results: the blue line represents indoor temperature generated by MZEST and the red line represents the prediction, while the black line is the outdoor temperature. Predictions are tracking the “actual” values very well for all cases. Optimal defaults have been chosen to minimize the total prediction error among the four different houses.

**Figure 6. Learning on Four Typical California Houses**

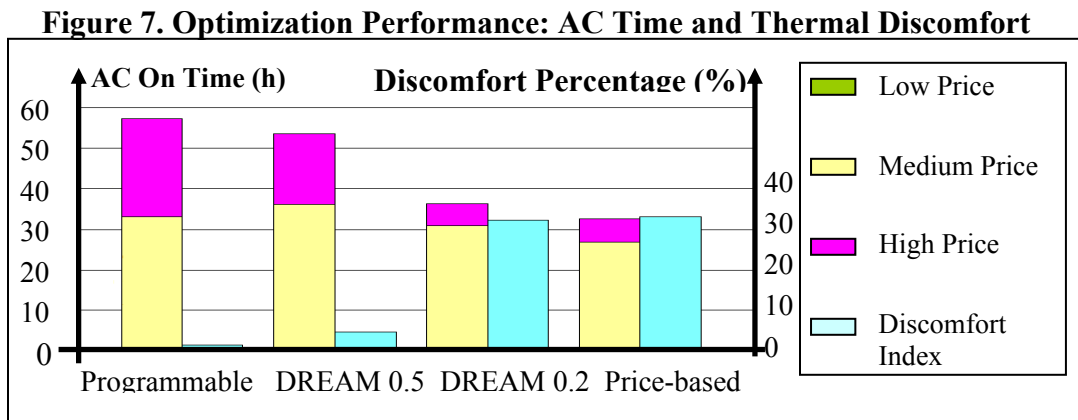


**Optimization on cost and comfort enable DR-response and adaptive control.** To evaluate the performance of the optimization strategy, we compared it with two typical strategies. One is the EnergyStar programmable thermostat’s default settings: daytime setpoint 25.5C and nighttime set-back setpoint 28C (EnergyStar 2006). It does not consider DR rates. The other strategy is a



price-based setting that adjusts the setpoints based on DR signals as well as time: daytime: 25.5C (low price); 26.5C (medium price); 28.5C (high price); nighttime: 28C. The DREAM optimization with economics index of 0.5 (default setting) was compared to the programmable thermostat; the DREAM optimization with economic index of 0.2 (low comfort) was compared to the price-based scenario. The house model with optimal default parameters was used. Using the MZEST, we ran the simulation using the pre-1978 house with crawl space in a Sacramento climate under thermostat control with these four settings. The metrics measured were AC on time at the different electricity rates and the user’s discomfort level based on the Adaptive Comfort Standard.

The results are presented in Figure 7. First, the left two results show that the DREAM with default economic index of 0.5 has fewer hours of air conditioning during high price periods and slightly more at medium price periods compared with programmable thermostat setting. Both offer acceptable thermal comfort since the discomfort indexes are small. This indicates that the DREAM successfully shifts the load from high-price period to medium-price period without significantly decreasing users’ comfort. The actual cost savings would depend on the difference between the high and medium prices. On the right, the price-based setting shows fewer hours of air conditioning use, but sacrifices users’ comfort. The DREAM with economics index of 0.2 has a similar performance. Thus we see that the DREAM responds automatically to price signals with appropriate energy saving behavior. Additionally, the design of the DREAM offers users choices on their economic and environmental preferences by adjusting a single variable—the economics index.



### Field Test

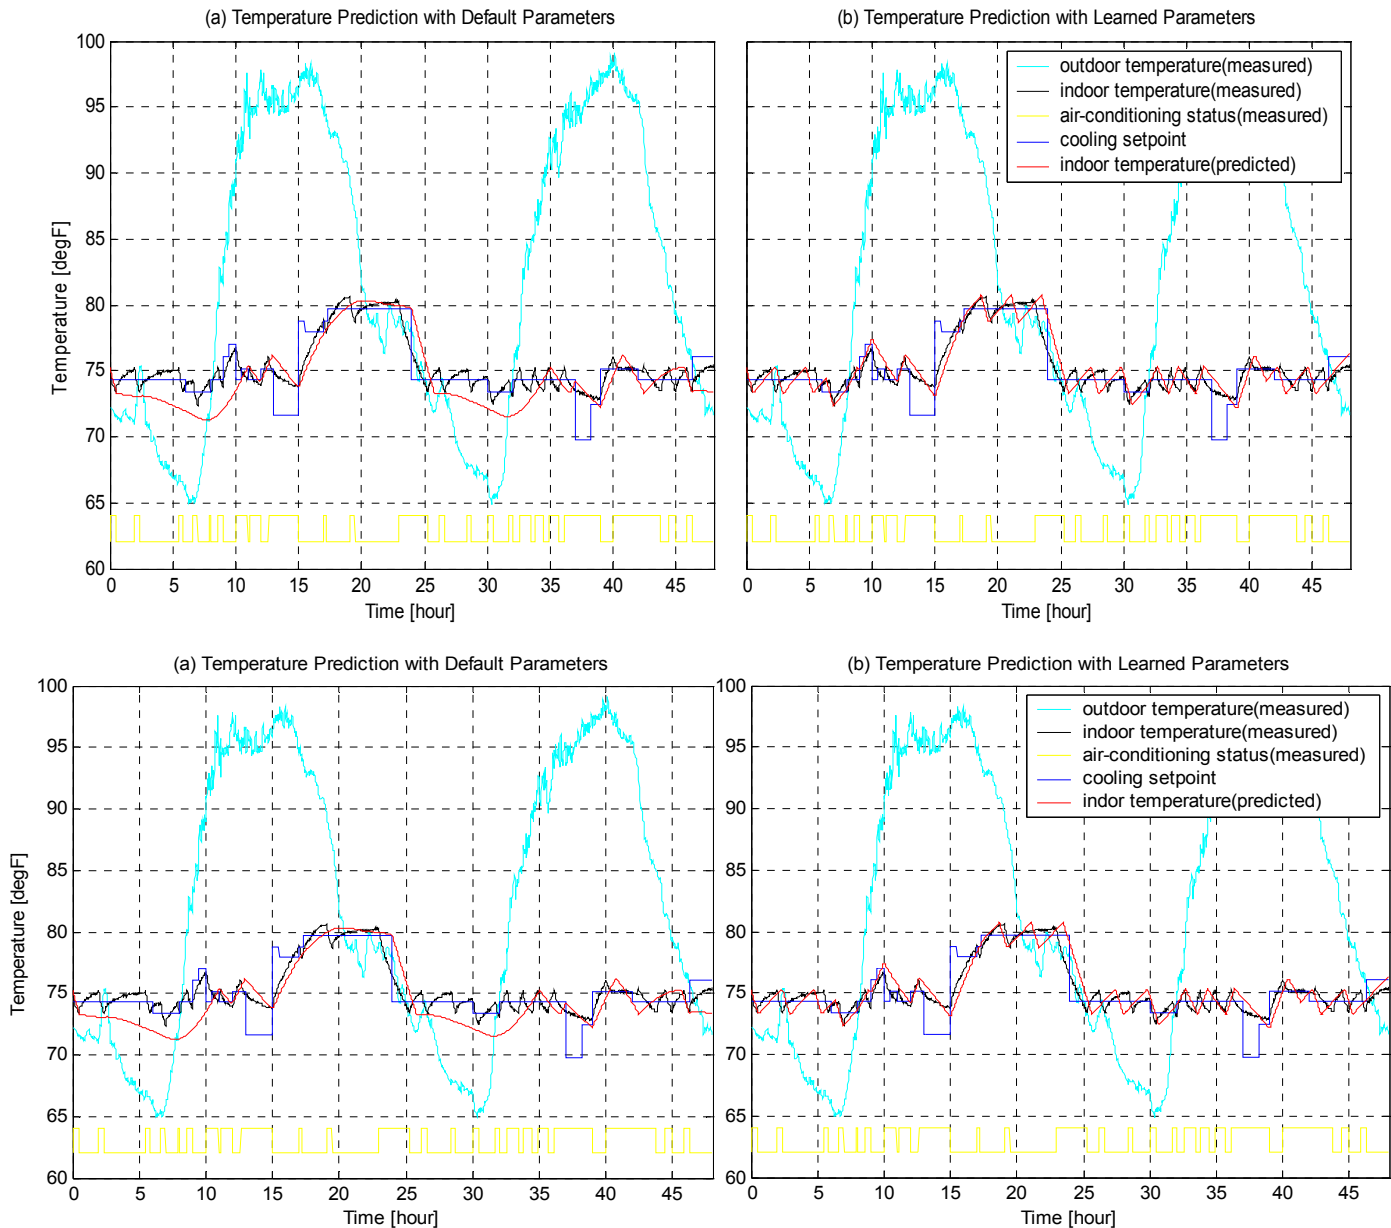
We tested the DREAM system in two single family occupied houses during summer 2007. The purpose was to test the functions of the system, to verify simulation results, and to get feedback from participants. The two houses were exposed to similar outdoor conditions, but the house structure, HVAC system and residents’ schedules were different. This diversity offered the opportunity to test the system under different conditions. The total time for the test was approximately six weeks. One test began two weeks earlier than the other, leaving time to fix problems if any occurred in test 1. Each test was divided into three phases as described in Table 2, while the communication reliability was continuously monitored.

**Table 2. Field Test Plan**

<b>Name</b>	<b>Length</b>	<b>Description</b>
System check out period	1 week	The DREAM monitors temperature, occupancy status, electricity use, and HVAC status under the control of original thermostat. The purpose is to learn participants' temperature preference at different time of a day and evaluate default house model and AC efficiency.
Mimicking period	2 days	The DREAM controls the HVAC system in the same manner as the original thermostat. This time is used to test the actuation functions and to train the occupants to interact with the DREAM interface.
Testing period	5 weeks	HVAC system is completely under control of the DREAM. Test focuses on optimization and house model learning: 1) Validate learned house model by comparing predicted indoor temperature with actual temperature. 2) Validate the strategy transition when price or occupancy status changes. 3) Compare the optimization performance (setpoint) under different values of economics index, using default house model or learned house model.

**House learning algorithm is promising.** Two sets of parameters are compared: the optimal default set and the parameters learned using one-month of measured data from house 1. The actual outdoor conditions (temperature, solar radiation), AC status and the initial indoor temperature were applied to the learned house model for the prediction. Close-loop control based on the given setpoint was applied to the house model. Figure 8 shows the comparison between measured indoor temperature (black line) with the predicted temperature (red line) for two consecutive days. Although the default model provided a good prediction (Figure 8(a)), we see significant improvement of the prediction quality after parameters were learned (Figure 8(b)).

**Figure 8: Indoor Temperature Prediction with/without Parameter Learning**



Antioch, California from Sep.1, 2007 – Sep.2, 2007

Prediction errors come from the assumptions we made when we set up the house model. Internal gain was initially modeled as constant, but actually varies by the number of residents and their activity. Occupants' behavior such as opening windows also changes the infiltration rate of the house. If this detailed information could be provided, the prediction would likely be more accurate.

**Strategies transitioned accurately.** Because the house was always occupied during these two days, state transitions were only triggered by price events. Future events—defined as those

occurring in two hours—were taken into account because the time frame is long enough to apply precooling strategies. Based on the price schedule, precooling was enabled at 6am and 1pm.

**Optimization generates reasonable setpoints based on price.** To validate the economics index, it was set to 0.3 for day 1 and 0.7 for day 2. Users are assumed to be comfortable (100%) at 74F. Setpoints were recalculated by optimization every half hour. Table 3 shows the average setpoints determined. In normal mode, setpoints were the same at low price, while slightly higher at medium price periods and much higher during high price period for day 1, due to the smaller economic index.

**Table 3. Average Setpoint at Different Prices for Two Days**

Economics Index	Low	Med	High
Day 1: 0.3	74F	76F	79F
Day 2: 0.7	74F	74F	75F

## Conclusion

We developed and tested a disaggregated thermostat that automatically responds to utility price signals. The Demand Response Electrical Appliance Manager (DREAM) uses a star wireless network of multiple sensors to create an information-rich environment. Wireless actuators complete the control loop in replacing thermostat relays to HVAC equipment. The sophisticated controller at the heart of DREAM uses this information to learn information about the house and its equipment to optimize the temperature setpoint for both cost and comfort. The user interface provides a friendly and informative face to the system, allowing people full control over the system but also teaching and advising them how to save energy and money.

The DREAM system was tested via simulations, lab tests and field implementations. We successfully demonstrated the promise of a smart, adapting, demand responsive disaggregated thermostat that uses wireless technology. The optimization algorithm in conjunction with the user-specified cost/comfort index worked well in the field tests to provide appropriate comfort for the price. Meanwhile, the tests provide insight on potential issues for DR policy and technology (described in Peffer et al, 2008). We hope this work provides a springboard towards further research and development in learning systems and human behavior in developing a demand response future.

## Acknowledgements

The authors gratefully acknowledge support for this work from Cliff Federspiel, Charlie Huizenga (Center for the Built Environment, Architecture Dept). Student researchers at UC Berkeley providing exemplary contribution include Carrie Brown, Kyle Konis, Anna LaRue (graduate student researcher) and Po-kai Chen, Reman Child, Sun Chen, Ashik Manandar, Jonathan Ellithorpe, William Watts, Yi Yuan (undergraduate student researcher). This project was sponsored by a PIER contract from the California Energy Commission.

## References

- Arens, E., et al. 2008. Demand Response Enabling Technology Development, Phase II Report from the Thermostat/Controls Group, December 2005 – December 2007.
- Arens, E., et al. 2006. Demand Response Enabling Technology Development, Phase I Report: June 2003-November 2005, Overview and Reports from the Four Project Groups.
- Archacki, Raymond. 2003. Carrier Thermostat Mode Summary.
- ASHRAE. 2004. Standard 55-2004: Thermal Environmental Conditions for Human Occupancy. Atlanta, GA: American Society of Heating, Refrigeration, & Air-Conditioning Engineers.
- Auslander, D.M., J.R. Ridgely, J.D. Ringgenberg. 2002. Control Software for Mechanical Systems, Upper Saddle River: Prentice-Hall.
- CEC, 2004. Residential Appliance Saturation Survey. Sacramento: California Energy Commission.
- EnergyStar 2006. Program Requirements for Programmable Thermostats: DRAFT 1 - version 2 [www.energystar.gov/ia/partners/prod\\_development/revisions/downloads/thermostats/ProgramThermDraft1.pdf](http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/thermostats/ProgramThermDraft1.pdf)
- EPRI. 2002. The Western States Power Crisis: Imperatives and Opportunities. [www.epri.com/WesternStatesPowerCrisisSynthesis.pdf](http://www.epri.com/WesternStatesPowerCrisisSynthesis.pdf)
- Peffer, Therese, Ed Arens, Xue Chen, Jaehwi Jang, and David Auslander. 2008. A tale of two houses: the human dimension of demand response technology from a case study of an adaptive wireless thermostat. In *Proceedings of the ACEEE 2008 Summer Study on Energy Efficiency in Buildings*.
- Peffer, T., Do, A., Whitney, C., & Langford, K. 2005. SIMS Spring 2005 DREAM interface design and development, from [www.sims.berkeley.edu:8000/academics/courses/is213/s05/projects/thermostat/index.php](http://www.sims.berkeley.edu:8000/academics/courses/is213/s05/projects/thermostat/index.php)
- PG&E. 2004. [www.pge.com/notes/rates/tariffs/advice/tariffsheets/19882-19901.pdf](http://www.pge.com/notes/rates/tariffs/advice/tariffsheets/19882-19901.pdf)
- USCensus. 2000. Census Data. <http://factfinder.census.gov/servlet/SAFFHousing?sse=on>