4235

Architecture for Intelligent Thermostats That Learn from Occupants' Behavior

Alex Boisvert **Ruben Gonzalez Rubio, Ph.D.**

ABSTRACT

This paper proposes a new approach to thermostat design. For many years, thermostats have been "dumb" devices, meaning that they react to their environment either by direct user control or by previous user programming. This new approach details an intelligent thermostat that learns about the behavior of the occupants and their environment and controls ambient temperature to maintain comfort according to human specifications. In that way, the thermostat reduces the number of interactions with the user and eliminates the need for them to learn how to program the device. Additionally, the thermostat reduces energy consumption by setback when occupants are absent. While the proposed architecture fundamentally changes the functionality of today's conventional thermostats, it retains their simple user interface.

This article presents the modular software architecture of this new intelligent thermostat design. The functionality of the thermostat in different states is described and how each module specializes in learning a certain pattern is explained. At the end, the results obtained using neural networks as a technique for learning are presented.

INTRODUCTION

Today, a growing number of economic and environmental considerations are leading us to look at new ways to reduce energy consumption. In northern regions, the activity that uses the most energy is the heating of buildings. To some extent, the same logic applies to regions where air conditioning is widely used. In this context, it is important to consider what determines the quantity of energy spent.

Figure 1 illustrates the thermal system under consideration. First, the external environment is uncontrollable and is the most important factor that influences energy consumption. Second, the physical building characteristics are important, but in most cases they are adequate and fixed. Third, a thermostat is used to monitor ambient temperature and control a heating or cooling system. Finally, one or more occupants operate the thermostat to obtain a given ambient temperature that reflects their thermal comfort.

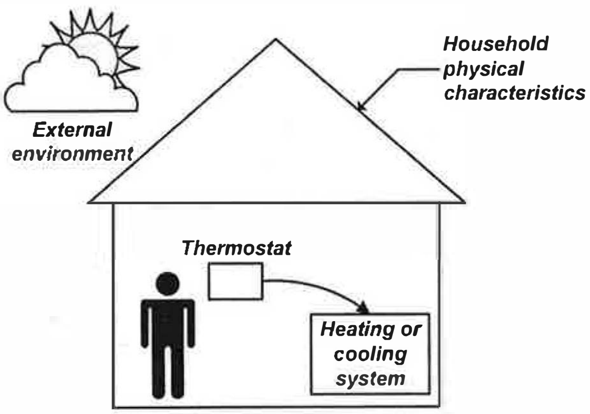

Figure 1 *Thermal environment.*

In this view, a person interacts with a thermostat, which, in turn, interacts with a heating/cooling system. In our mind, it is this interaction that mostly influences energy consumption. Occupants are often lazy about frequently adjusting their thermostats. Also, Harmon (1981) argues that "many people do not fully understand the proper operation of the most common residential thermostat." For these reasons, the authors designed an intelligent thermostat that automates comfort control and energy conservation.

Alex Boisvert is a graduate student and **Ruben Gonzalez Rubio** is a professor in the département de génie électrique et de génie informatique, Université de Sherbrooke, Québec, Canada.

BACK TO PAGE ONE

Types of Thermostats

Two types of thermostat inspired this design—conventional and programmable thermostats. The conventional thermostat is, by far, the most popular. This type of thermostat is characterized by a simple interface and very simple functionality. Usually, a rotary dial is used to directly specify a set-point temperature. The thermostat measures ambient temperature and, if it does not correspond to the set point, it activates the heating/cooling system to restore ambient temperature to the set point. Newer versions of this thermostat use digital controls (buttons) instead of a rotary analog dial. The problem with this type of thermostat is that it must be manually adjusted every time an occupant wants to change the ambient temperature. In order to vary the temperature according to their lifestyle, the occupants must repeatedly interact with the thermostat. Moreover, the occupants are directly responsible for the energy consumption of the heating/cooling system. If they leave for a period of time, they must (again) adjust the temperature in order to save energy. Generally, however, people have little concern for energy conservation. In fact, studies have shown that in cold climates, people have a tendency to elevate temperature without restoring it afterward (Benton 1992). Clearly, it is difficult for an average user to put into practice an efficient energy-saving plan.

The second type of thermostat design is the "programmable thermostat," which is an extension of the conventional thermostat. This design automates the task of adjusting the thermostat to the lifestyle and the personal preferences of the occupants. The thermostat is programmed in advance to adjust the set point with respect to a given schedule. This is an advantage because it reduces the number of interactions between the user and the thermostat, but it doesn't solve the problem completely because the occupants still have the responsibility of correctly programming the thermostat to reduce energy consumption. In fact, long-term energy consumption will decrease significantly only if (1) the thermostat is correctly programmed, (2) the occupants follow the programmed schedule, and (3) they update the schedule as their lifestyle changes. Also, if it is compared with a conventional thermostat, the number of interactions with the occupants decreases but each interaction now takes more time and is more difficult due to programming and the added complexity of the user interface.

Design Goals

In light of the previous discussion, the design goals of an ideal home thermostat are to

- maximize comfort,
- minimize energy consumption (without sacrificing comfort),
- keep user interaction to a minimum,
- present a simple user interface, and
- keep to a minimum the cost of installation (i.e., use existing heating/cooling system) and operation.

The Approach

The design approach considers each of the above goals; the order in which the goals are stated reflects their respective importance in this design.

The first step in this approach is to borrow concepts from both conventional and programmable thermostats. With respect to the former, we wish to retain the simple user interface, and from the latter, we would like to retain the concept of following the lifestyle of the occupants but without explicitly programming the device.

However, both of these thermostats lack an important concept: the presence of occupants. This additional information can lead to an interesting compromise between comfort and energy savings because the notion of comfort is not applicable if nobody is present. In fact, this simple rule is the cornerstone of this approach to minimizing energy consumption.

The new design avoids numerous interactions between the user and the thermostat because the thermostat learns the occupants' behaviors and automatically programs itself. With a conventional thermostat, every interaction with the user reflects a need to adjust the set-point schedule of the thermostat. According to this, the information needed to program the thermostat is obtained implicitly, one interaction at a time, instead of having a complete schedule programmed explicitly. Consequently, the thermostat can learn the schedule as time passes, recording the interactions and generalizing them to depict a pattern and generating a schedule from this pattern.

Learning, then, is what differentiates this approach from others. As it turns out, learning can be used in many ways in a thermostat. For example, provided that we have information on the history of the occupants' presence, we can predict their absence and save energy by not maintaining the ambient temperature within the comfort interval.

We call the resulting design an intelligent thermostat. We use the term "intelligent" because (1) the thermostat learns, (2) it imitates the behavior of the occupants (with respect to their interaction with a conventional thermostat), and (3) it is more intelligent than a "dumb" conventional thermostat.

This paper focuses on the high-level design and architecture of the proposed intelligent thermostat. The implementation of such a device benefits from, but does not require, the use of advanced artificial intelligence concepts (i.e., artificial neural networks, fuzzy logic). A description of such concepts or algorithms is beyond the scope of this article.

ASSUMPTIONS

The design of the intelligent thermostat depends on a number of available devices. Each of the following devices is required:

- A thermometer to measure ambient temperature. Since precision of measurement can be an issue, we consider a precision of at least 1°C.
- A simple user interface, similar to actual conventional thermostats. A thermostat with simple digital interface

BACK TO PAGE ONE

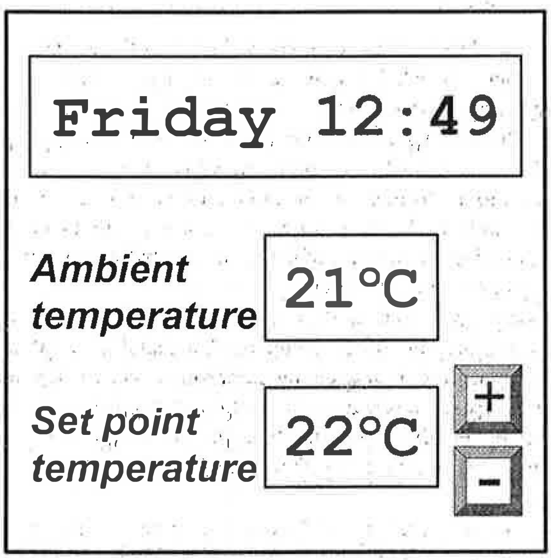

Figure 2 *Simple thermostat user interface.*

is illustrated in Figure 2.

- A microcontroller with random access memory and a real-time clock. We consider that the microcontroller is capable of doing floating-point arithmetic calculations.
- A nonintrusive device to determine if the occupants are present or absent (i.e., a motion detector).

FUNCTIONAL DESCRIPTION

Let us first describe the functionality of the thermostat. The intelligent thermostat has three states, which are illustrated as a finite-state automaton in Figure 3. Following is a brief description of each state:

Comfort mode—In this state, the temperature set point is determined automatically by the thermostat with respect to the set-point history learned from the occupants (called the set-point schedule).

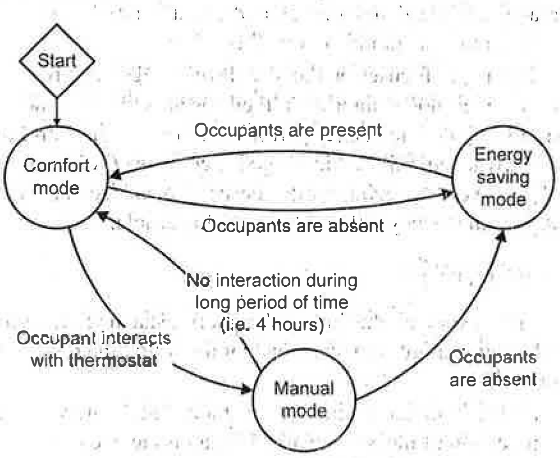

Figure 3 *States of the intelligent thermostat.*

Energy-saving mode—When the occupants are absent, the thermostat minimizes energy consumption. Consumption is reduced because the set-point temperature is modified to allow a certain deviation from the normal set-point schedule. This concept has been called "setback" in the literature (Schade 1978). The extent of the deviation depends on the present history learned from the occupants (present schedule).

Manual mode—When an occupant interacts with the thermostat to dictate a set-point temperature, the thermostat enters this state for an arbitrary period of time (i.e., eight hours). During this time, the set-point temperature is the one given by the user. If during this time the occupants leave, the thermostat switches to the energy-saving mode.

SOFTWARE ARCHITECTURE

In electronic devices, intelligence and function are often provided by embedded software. Using a microcontroller with memory allows us to use a software design approach to solve our problem.

The software architecture of the intelligent thermostat is composed of three main modules, as shown in Figure 4. Each module defines a logical boundary separating independent units of processing that perform specific functions.

Modules are combined to provide higher-level functionality. In the intelligent thermostat, the interactions between each module depend on the state of the thermostat. The logic that controls how modules interact together can be thought of as an additional higher-order module. For simplicity, it is not illustrated as such.

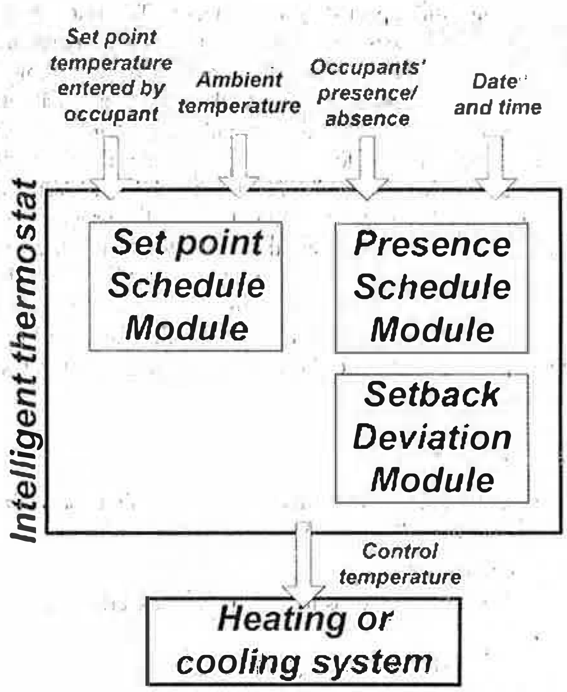

Figure 4 *Intelligent thermostat design overview.*

BACK TO PAGE ONE

Modules are described below.

Inputs and Output

The intelligent thermostat has four inputs:

1. An ambient temperature readout, obtained from a thermometer
2. A time indicator
3. A set-point temperature entered by an occupant using the thermostat console
4. An indication of the presence (or absence) of occupants

Compared with other common thermostats, only the last input is a novelty.

Since we did not want to change existing heating/cooling system equipment, the intelligent thermostat has only one output: the control temperature. This control temperature can be compared to ambient temperature to generate an on/off indicator to activate the heating/cooling system. In this way, the intelligent thermostat can be used as a front end to an existing thermostat.

PER-STATE PROCESSING

Comfort Mode Processing

When the thermostat is in the comfort mode, the occupants are present but they do not interact with the thermostat. Three operations are executed in parallel, as shown in Figure 5:

1. The set-point temperature is chosen according to the set-point history learned from the occupants (set-point schedule).
2. The thermostat records the ambient temperature and reinforces the set-point schedule. The motivation for doing so is because the occupants are present and they are implicitly agreeing with the ambient temperature because they do not interact with the thermostat to modify the set-point temperature.
3. The thermostat records the presence of the occupants. This information is stored and used to sketch an estimated presence schedule.

Energy-Saving Mode Processing

When the occupants are absent, the thermostat switches to energy-saving mode. The following operations are performed simultaneously, as illustrated in Figure 6:

1. The thermostat fixes a setback temperature to reduce energy consumption. The setback temperature represents a deviation from the normal comfort schedule set-point temperature. To calculate the allowable deviation, the thermostat considers the probability that the occupants return home at this moment. Generally speaking, if a return is highly probable, then the deviation will be zero; if it is improbable, then the deviation will be high (i.e., up to 5°C). Therefore, when the occupants return home according to their usual schedule, the ambient temperature is restored to their comfort level when they arrive.

   Since the power of heating/cooling systems varies from one installation to the other, it is possible that the rate of change of the setback temperature is higher than the rate of change of the ambient temperature that the heating/cooling system can provide. For this reason, the thermostat controls the heating/cooling system so that the ambient temperature always remains within the setback deviation. To achieve

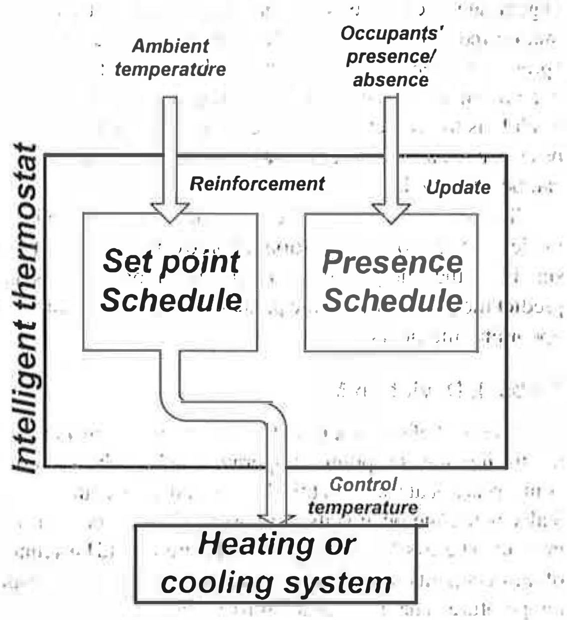

Figure 5 *Comfort mode logic.*

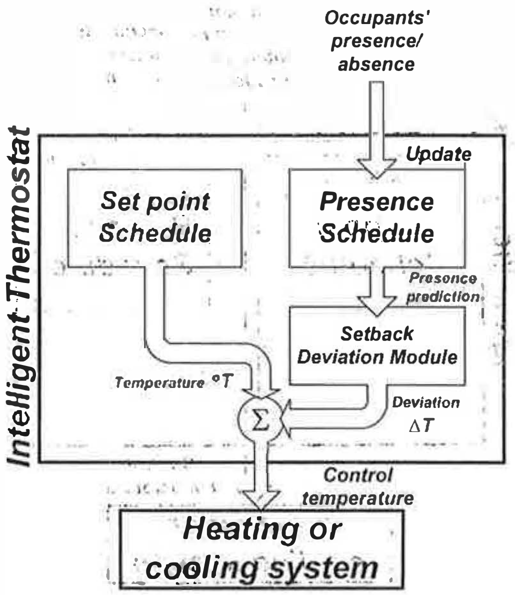

Figure 6 *Energy-saving mode logic.*

BACK TO PAGE ONE

this, the thermostat predicts upcoming ambient temperature setback deviations and activates the heating/cooling system in advance, if needed. (This mechanism is beyond the scope of this article, so it is not covered in detail.)

2. The thermostat records the presence of occupants and updates the presence schedule.

Manual Mode Processing

Manual mode is the simplest of all three modes. It is assumed that the user interacted with the thermostat and specified a set-point temperature that corresponds to a desired comfort level. The following tasks are performed in parallel, as shown in Figure 7:

1. The thermostat activates the heating/cooling system with regard to the user-specified set-point temperature.
2. The thermostat records the user-specified set-point temperature and updates the set-point comfort schedule.
3. The thermostat records the presence of occupants and updates the presence schedule.

PER-MODULE DESCRIPTION

Comfort Schedule Module

This module encapsulates a schedule of ambient temperature set points that the occupants determine. With a programmable thermostat, this module is a fixed schedule that users explicitly program. Within the intelligent thermostat, the schedule is not explicitly programmed at once but, rather, it is implicitly and incrementally generated by the occupants who modify the set-point temperature as time goes by.

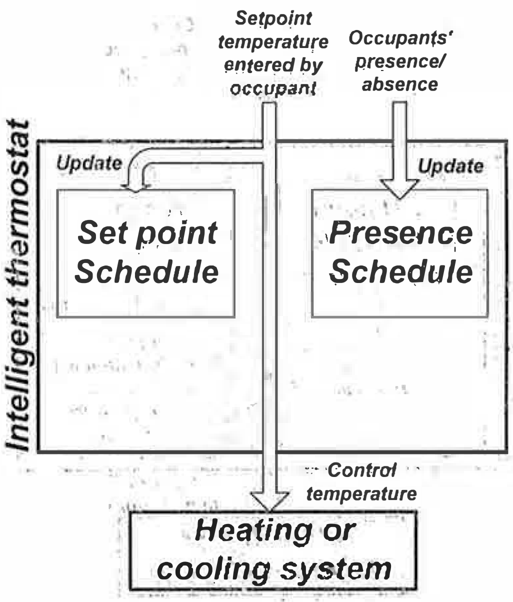

Figure 7 *Manual mode logic.*

It is assumed that the temperature set points specified by the occupants reflect an easy comfort level. Then, in essence, this module tries to establish a temporal pattern of desired set-point temperature.

Each set point given by an occupant is recorded along with the time of the day and the day of the week when the interaction took place. The sum of all these interactions forms a data set used to generate a schedule that corresponds to the instructions given to the thermostat.

When the thermostat is in the comfort state, this module is used to determine the ambient temperature with which the occupants are most likely to be comfortable. While in this state, the thermostat collects information on the actual ambient temperature and feeds it to this module to reinforce the current schedule. Again, the motivation for doing this is because the occupants are present and are implicitly agreeing to the ambient temperature since they do not interact with the thermostat.

In the manual state, the module is fed the new interactions. In this case, the actual set-point temperature is determined by the user, not by this module.

Finally, in the energy-saving mode, the module determines the temperature against which a deviation is permitted. The temperature given is the same as in the comfort mode, but a deviation from this temperature is allowed since no occupants are present and reducing energy consumption becomes a priority.

Presence Schedule Module

The presence schedule module encapsulates a schedule of probability of presence of the occupants. In that respect, it is very similar to the comfort schedule module because both try to generate a schedule based on temporal data sets. The information fed to the presence schedule module is collected by a "presence detector" and consists of a boolean value (present or not present) coupled with the current time. The role of the module is to predict the presence behaviors of the occupants based on their presence history. This assumes that presence can be predicted.

The module records presence information in all three modes of thermostat operation. Moreover, when the thermostat is in the energy-saving mode, the module is asked to predict the probability of the presence of the occupants for the upcoming moments.

Setback Deviation Module

The goal of setback is to decrease energy consumption by modifying the set-point temperature when the occupants' comfort is not at stake. In this design, setback occurs automatically when the occupants are absent. The novel idea that is presented here is that setback temperature should be a function of the occupants' probability of presence and desired set-point temperature. The mathematical formula is:

$$T_{setback} = T_{desired} \pm T_{deviation}(P_{presence}) \quad (1)$$

BACK TO PAGE ONE

As shown, the deviation allowed from the desired temperature should be a function of the occupants' probability of presence. The relation should be zero when $P_{presence}$ is 100% and should be a significant number of degrees (i.e., 5°C) when $P_{presence}$ is 0%. We have used the sigmoid (S-shaped) function and have had good results:

$$T_{deviation} = 5°C \times \frac{P_{presence}}{1 + e^{-P_{presence}}} \tag{2}$$

The setback deviation module is, therefore, just a function taking the probability of the occupants' presence as input and calculating an allowable deviation as output. This module is coupled with the presence schedule module and is used only when the thermostat operates in the energy-saving mode.

IMPLEMENTATION

We have developed a prototype of this intelligent thermostat in a computer simulation environment. To do this, we used a commercial mathematical programming environment and a transient system simulator.

Learning

Both learning modules (comfort and presence schedules) were based on artificial neural networks. This technique was chosen because neural networks are known for their capability to generalize relations and because they show a strong level of robustness when faced with noisy inputs, which, in our case, translates to exceptions within the schedules. In an effort to reduce redundant information and to allow the learning of new patterns rapidly, we automatically filter the collected data sets as time passes based on redundancy and coherence with known patterns.

RESULTS

To evaluate the relative performance of the thermostat design against existing ones, we set up the simulation environment shown in Figure 8.

We used the simulation program to model the building and heating/cooling system characteristics. The exterior environmental conditions are modeled by collected data. Using the math programming environment, we simulated three types of thermostats: conventional, programmable, and intelligent. Combined with simulated human behavior, this system reproduces the system of Figure 1.

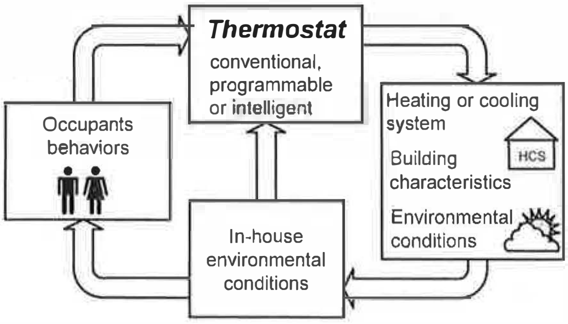

Figure 8 *Simulation environment.*

To measure each thermostat's performance, we used three metrics: (1) the number of interactions between the occupants and the thermostat, (2) the period of comfort for users, and (3) the setback period. Both comfort and setback are expressed as a percentage of comfort/setback time over the total time. We avoid any comparison with the programmable thermostat because this would lead to arbitrary assumptions on the correctness and frequency of this thermostat's programming with respect to the occupants' behaviors.

The first case studied is a theoretical situation used to demonstrate how many interactions are needed by each thermostat to obtain equivalent energy savings. The results are shown in Table 1. We assumed that occupants are absent during weekdays for work and that setback is used during the day and night. During the weekend, occupants are mostly present, and, therefore, setback only occurs during the night. Also, we took for granted that occupants follow this exact pattern for 28 consecutive days. Over the course of this period, the thermostat is in energy-saving mode about 57% of the time. It can be noted that with a conventional thermostat, occupants cannot be comfortable 100% of the time since they have to be physically present (and awake) to manually adjust the thermostat for the setback period to end, whereas with the other thermostats this is done automatically before the occupants wake up or come back. We chose a period of 28 days because it is unlikely that a behavior pattern lasts for more than this period of time, and it is unlikely that occupants will "reprogram" their thermostat more often.

TABLE 1
Theoretical Case with Equal Energy Savings for Each Approach (28-Day Period)

Metric	Conventional Thermostat	Intelligent Thermostat	Programmable Thermostat
Number of Interactions	96	4	1 (complex)
Comfort Period	96%	100%	100%
Setback Period*	57%	57%	57%

* When occupants are present

The second case studied (Table 2) is more realistic. We assumed that occupants do not exactly follow the same schedule every day. During the weekdays, occupants come back home for lunch at a random time between 11:00 a.m. and 2:00 p.m. Also, the moment they leave and when they come back varies up to an hour. During the weekend, they leave for about three hours at a random moment. We also consider a more typical use of the conventional thermostat where occupants only adjust it once a month.

BACK TO PAGE ONE

TABLE 2
Typical Use of Each Thermostat Over a 28-Day Period

Metric	Conventional Thermostat	Intelligent Thermostat	Programmable Thermostat
Number of Interactions	1	7	1 (complex)
Comfort Period*	100%	98%	96%
Setback Period	0%	56%	57%

* When occupants are present

The results of the second case studied illustrate that people who have fairly foreseeable behavioral patterns can significantly reduce their energy consumption (almost as much as a correctly programmed thermostat) using an intelligent thermostat. In our test, the occupants adjusted the thermostat five times during the first two days and only two times for the remaining twenty-six days. This shows that once a pattern is learned (after two days), the intelligent thermostat can accurately predict the occupants' behavior and provide both comfort and energy savings.

If behavior were less predictable, comfort would decrease slightly as the thermostat could not accurately predict the presence of the occupants and determine the correct time to set back or recover. In this case, the thermostat would become more conservative with setback to avoid long recovery time where the occupants are uncomfortable. This, in turn, would reduce potential energy savings.

Compared to a conventional thermostat, we have measured energy savings varying from 9% to 16% with the intelligent thermostat, which is equivalent to the results obtained with the programmable thermostat.

CONCLUSION

We have presented a new architecture for intelligent thermostats that provides better comfort for occupants and also reduces energy consumption. Both of these advantages are obtained because the thermostat learns the occupant's behavior and decides when setback should begin and when it should recover, as well as what set-point temperature the occupants prefer as a function of time.

Technically speaking, we have covered the high-level modular software architecture of the thermostat by describing each module's role and explaining how processing occurs within each of the thermostat's states.

The proposed thermostat is thought to be more intuitive than a programmable thermostat, which requires tedious programming and has a complex user interface. Instead of being programmed "all at once," it is programmed gradually, as the behavior occurs. Moreover, it is believed that the learned set-point schedule follows the occupants' preferences more closely because it is adjusted every time a user interacts with the thermostat. And, finally, we have shown that this thermostat requires fewer interactions by occupants compared to conventional thermostats.

ACKNOWLEDGMENTS

We would like to thank David Bestman-Smith and Hugo Jacques, both former students at the University of Sherbrooke, who contributed to the design of this architecture.

This research was partially funded by FCAR Fond Canadien de l'Avancement de la Recherche.

REFERENCES

Benton, R. 1992. The impact of thermostat performances on energy consumption and occupant comfort in residential electric heating systems. *ASHRAE Transactions* 88(1): 1029-1041.

Harmon, K.S. 1981. Advanced control strategies for energy conservation in building operation. *ASHRAE Journal*, July, pp. 55-57.

Schade, G.R. 1978. Saving energy by night setback of a residential heat pump system. *ASHRAE Transactions* 84(1): 786-798.