ESTIA: A Real-Time Consumer Control Scheme for Space Conditioning Usage under Spot Electricity Pricing

P. Constantopoulos
University of Crete and Research Centre of Crete
Heraklion, Crete, Greece

F.C.Schweppe †
R.C.Larson
Massachusetts Institute of Technology
Cambridge, MA., USA

Scope and Purpose

With record daily electrical energy consumption recorded for both summer and winter months in many parts of the US and other countries over the past several years, the efficient use of electricity has become a global priority. On the one hand, the "greenhouse effect" presents a deterrent to building additional coal- and oil-fired plants; on the other, the fear of catastrophic accidents has all but stopped construction of nuclear power plants. The point of view of this paper is that communications and computer technologies are now relatively inexpensive, so that hourly (or more frequent) "spot prices" of electricity could be telecommunicated to electricity users. In turn, the users would utilize a computer-based operations research control algorithm to govern the discretionary use of electricity within the home (site, in general), the objective being a balancing of human needs and cost minimization. Consumer energy usage, if modified by spot pricing, could result in much more efficient utilization of scarce electrical generation capacity.

Abstract

Under spot electricity pricing the price of electricity represents actual demand and supply equilibrium conditions and consumers respond to price variations so as to achieve the best cost-service tradeoff over a given time period. Using a decision modelling approach developed for prescribing consumer response to a varying electricity price, the case of space conditioning usage is analyzed in detail and a real-time control scheme is proposed. Important classes of consumers, such as hospitals, commercial buildings and dwellings, are modelled and simulations are performed. It is found that significant

† Deceased, July 1988.

savings on a daily basis can be achieved through this scheme at the price of only small fluctuations of indoor temperature around its ideal value. Consumer preference consciousness results in static and dynamic savings. The latter are by far the largest and are only enabled by a varying price of electricity. A short-sight property limits the effective control horizon drastically, thereby rendering price forecasting practically unnecessary and reducing the data and computing requirements of the control scheme.

1. Introduction

Variable electricity rates, as opposed to direct load control, enable customer independence from the utility. Predetermined variable rate schemes, such as time-of-day energy and demand pricing, have invited the use of numerous computer-based energy management/demand control systems by large customers, often with impressive success in reducing customer bills with sometimes, but not always, a corresponding reduction in utility costs.

In this paper we consider the problem of determining electricity consumption schedules which reflect the reaction of individual consumers to a "spot" electricity price [9, 12, 13, 14]. Encouraged by the sustained dramatic drop in communication and computing costs, spot pricing establishes an energy marketplace where the price reflects actual demand and supply conditions and is set and communicated to the customers periodically, e.g., once a day or once an hour. Spot prices during a giver hour reflect actual utility costs during that hour much more precisely than predetermined variable rates (also known as "time rates"). Spot prices incorporate capital recovery factors directly in a time variable energy rate and thus the distortions caused by customer response to demand charges are avoided. A comprehensive treatment of spot electricity pricing can be found in [15]. A number of important issues concerning the economics and the implementation of spot pricing are discussed there. The particular implementation of a spot pricing scheme and the contract offered to the consumer by the utility are bound to influence the consumer's decision of subscribing to spot pricing. Such issues, though, are beyond the scope of this paper; rather, we are concerned with the *operational* problem that arises once a consumer subscribes to a spot tariff. The solution to this problem is to be taken into account when analyzing costs and benefits to decide whether to subscribe or not. It should be pointed out, however, that spot electricity pricing is not a hypothetical situation: Tabors et al. [16] give an account of the experience of several utilities with various versions of spot pricing, either at experimental or application scale.

The approach in the present work is based on a general model for individual load control under customer independence, suitable for automatic, computer-based, real-time control, which was presented in [5]. The purpose of this model is to determine power consumption schedules that best achieve the trade-off between total cost and levels of service obtained through various uses of electricity according to the preferences of the individual consumer. The criteria on which the choice of a power consumption schedule is based are total cost and end-use quantities describing quality of service (e.g., illumination for lighting; rate of supply and temperature for hot water; start of task execution for a machine-tool; indoor temperature and humidity for space conditioning). The consumer's choice process is represented as a maximization of joint utility with respect to these criteria. Under reasonable conditions the overall problem is decomposed into independent, single usage control problems, which greatly simplifies the entire control task.

Given this decomposition, we here concentrate on a detailed study of the control of space conditioning (heating and cooling). Space conditioning is an important use of electricity, common to all types of consumers - residential, commercial, industrial. It is known to offer significant load management potential and it is a major contributor to daily aggregate load peaks faced by electric utilities [10].

We begin by stating the space conditioning control problem. Quality of service is measured by the deviation of indoor termperature from a desired "ideal" point. We then present ESTIA, a computer-based control scheme. Subsequently we discuss properties and the performance of the scheme, the effects of price and weather, and those of certain model parameters. Finally, a simulated application study illustrates the performance of the model and its potential practical benefits under a realistic, marginal cost electricity tariff.

In summary, our main findings are the following:

1 With a constant price of electricity, a constant indoor temperature, generally not equal to the ideal one, must be maintained regardless of the weather. No special control algorithm is needed.

2 Under a spot price, different "optimal" indoor temperatures are prescribed at different times, depending on both price and weather. A variable price schedule enables considerable savings relative to an equivalent fixed price. Simulations involving relatively well insulated buildings (time constant equal to 25 hours) tariffs with maximum/minimum price ratios as high as 3 and quite adverse winter and summer weather conditions, show savings over the equivalent flat rate case not less than 10.9% and indoor termperature swings not larger than $6^o F$.

3 Changes in anticipated price or weather beyond a certain point in the future, do not affect current "optimal" power consumption and indoor temperature. This *short-sight property* of the control scheme is very important, especially because the *vision range* turns out to be much shorter than one day (our control horizon) - typically 6 and not more than 13 hours. Thus the forecasting horizon for future spot prices and the size of the control problem are considerably reduced. If spot prices are specified 6-13 hours in advance there is no relevant uncertainty in the price and consumers may reap the benefits of price variability at no risk.

2. Problem formulation

A space conditioning system has three main components: the building, the equipment and the users. Correspondingly, our model describes all three. Detailed treatments of the thermodynamics of buildings and equipment are available for design purposes in such sources as the ASHRAE handbooks [1]. The design procedures recommended there also account for the ways in which the intended use of the building modifies its thermodynamic behaviour. For our purposes we employ a very simple, linear, time-invariant dynamic model to represent the building and the equipment.

In reality, a building comprises a number of separate spaces, each with different temperature and humidity. Heat transfer and air mass transfer occur among spaces and between each space and the outside. So, there is an overall exchange of air mass and thermal energy between the building and the outside, as well as a circulation of air mass and energy inside the building. The heating or cooling loads come from external sources, including outside temperature, wind and insolation, and from internal sources, including people, lights, motors, etc. All sources are more or less random. In cooling applications we commonly distinguish sensible and latent heat loads, the first due to temperature and the second to the moisture of the air. Equipment is designed to meet weather conditions much worse than average and, in addition, it is typically oversized to guard against design or construction errors. Temperature control is achieved by means of thermostats, characterized by a temperature deadband around a set point, which results in finite on and off time for the equipment.

Here, we make the following assumptions:

1 We consider a single conditioned space. If this is one of many in a building, it interacts with the outside through its own external wall, if any, and it is supposed to be in thermal equilibrium with adjacent conditioned spaces. Or, it can be a whole building viewed as a single space where we neglect circulation effects and assume

uniform inside temperature and humidity, equal to the average values across the actual spaces.

2 The results of the heat load calculations are lumped together in a way that all heat loads are viewed as caused by a single, *effective outside temperature* T^o, from which the space is insulated by a shell with *overall thermal conductivity A*. The shell, the air mass and the other contents of the space have a *total thermal mass mc*.

3 We assume that no independent thermal storage is coupled to the main heating or cooling equipment, confining ourselves to the study of the basic system. Furthermore, a single piece of equipment is assumed.

4 We neglect the control of humidity and concentrate on temperature alone. At least for residential buildings, this can be approximately compensated by increasing the sensible heat load by 30% [1] - in our terms, by increasing the effective outside temperature. So, we can have the same dynamic model for summer and winter, but with different equipment parameters.

5 The cycling effect of the thermostat, due to the deadband, is neglected. This is primarily an analytical convenience, doing away with the annoying non-linearity introduced otherwise. The fractional power input prescribed corresponds to a specific duty cycle of the thermostat. When computer - controlled switching is available, the thermostat is unnecessary. The inside temperature is continuously monitored through a sensor and the applied control is revised periodically. With a small periodicity, say 5 minutes, and an on-off control, such a scheme can be expected to perform comparably with a thermostat. If continuous control is available, then the frequent, periodic scheme can achieve closer tracking of the desired temperature than a thermostat.

6 Internal heat sources are neglected.

Other, more general models appear in [6, 11].

The notation employed in our model is listed in Table 1. In a discrete time control procedure, the power input e_n, the spot price of electricity p_n, the desired indoor temperature T_n^d and the effective outside temperature T_n^o are assumed constant during each control period $n=0,\cdots,N-1$. The inside temperature T_n, is the instantaneous value of the inside temperature at the beginning of control period n (equivalently, end of period $n-1$).

As already mentioned, the performance criteria of the system are total cost and quality of service. Total cost over the control horizon is equal to

$$c=\tau(p_0e_0+\cdots+p_{N-1}e_{N-1}) \tag{1}$$

Table 1: Notation for the space conditioning model

Symbol	Definition	Unit
mc	total thermal mass	$kWh/{}^oC$
A	overall thermal conductivity	$kW/{}^oC$
τ	duration of control periods	h
E	power rating of equipment	kW
η	thermal conversion efficiency (heating)	
	coefficient of performance (cooling)	
$TC=mc/A$	time constant of the system	h
$\varepsilon=exp(-\tau/TC)$	factor of inertia	
T_n^d	desired inside temperature in period n	$''{}^oC$
T_n	inside temperature in period n	$''{}^oC$
T_n^o	effective outside temperature in period n	$''{}^oC$
e_n	electric power input in period n	kW
p_n	price of electricity in period n	$\$/kW$
$\hat{p}$	ceiling price of electricity	$\$/kW$
$l_n=T_{n+1}-T_{n+1}^d$	service loss in period n	$''{}^oC$
x'_n, x''_n	service loss limits in period n	$''{}^oC$
Δ'_n, Δ''_n	temperature rate limits in period n	$''{}^oC$
N	number of periods contained in control horizon	
E_T	total controlled connected load	kW
γ_n	service loss timing weight in control period n	
$b(\,)$	one-period service loss penalty function	
λ_L	utility weight on service loss	
$a=1/NE_T\tau\hat{p}$	utility scaling for total cost	
λ_C	utility weight on total cost	

A suitable measure for quality of service is the *service loss* trajectory $\{l_n \mid n=0,\cdots,N-1\}$ [4,5]. The one-period service loss is defined as the deviation of the actual from the desired inside temperature at the end of the control period:

$$l_n = T_{n+1} - T_{n+1}^d \quad n=0,\cdots,N-1 \tag{2}$$

It must be noted that the entire service loss trajectory is important to the user, while cost is of interest only as a total amount. Restrictions imposed by the user on the service loss trajectory can be expressed as tolerance ranges for service loss and rate of temperature change

$$-x'_n \le l_n \le x''_n \quad n=0,\cdots,N-1 \tag{3}$$

$$-\Delta'_n \le T_{n+1} - T_n \le \Delta''_n \quad n=0,\cdots,N-1 \tag{4}$$

In principle, humidity should also be monitored, but is here neglected for simplicity, as discussed earlier.

The consumer's decision problem can be mathematically stated as one of minimizing the expected value (in view of the uncertainly in weather and price) of a cost functional which is a joint merit index of the above performance criteria, subject to technical and use constraints.

The question of assessing an appropriate cost functional is discussed in [4,5]. The cost functional used here is:

$$f = \sum_{n=0}^{N-1} \lambda_C a \tau p_n e_n + \lambda_L \gamma_n b(l_n) \tag{5}$$

The consumer's attitude towards a single period service loss is represented by the penalty function b(.) which has the form of a bathtub function (see Figure 1). Usual, symmetrical bathtub functions are $b(x) = |x|, x^2, x^4$ if $-1 \le x \le 1$ and $b(x) = 1$, otherwise. The service loss timing weights, γ_n, measure the relative importance of the timing of the service losses and satisfy the relations:

$$\gamma_0 + \cdots + \gamma_{N-1} = 1$$

$$\gamma_n \ge 0, \quad n=0,\cdots,N-1 \tag{6}$$

Similarly, the utility weights on total cost and service loss, λ_C and λ_L respectively, measure the relative importance of the two performance criteria and satisfy the relations:

$$\lambda_C + \lambda_L \le 1$$

$$\lambda_C, \lambda_L \ge 0 \tag{7}$$

If space conditioning is the only controlled electricity usage, then the first of (7) holds with equality. Otherwise equality can be achieved by adding similar utility weights corresponding to the service loss trajectories of the other controlled usages. Procedures for assessing λ_C, λ_L, γ_n and $b(\,)$ are discussed in [4,7].

The technical constraints are expressed by a single dynamic energy balance equation along with control range inequalities which represent limits on the power inputs (eqs. 8a, 8d below).

In summary, our model for space conditioning using the notation of Table 1, is:

Energy balance: $$T_{n+1} = \varepsilon T_n + (1-\varepsilon)(T_n^o \pm \eta e_n / A) \tag{8a}$$
(+ : heating, - : cooling)

Initial condition: $$T_0 \tag{8b}$$

Service loss : $$l_n = T_{n+1} - T_{n+1}^d \tag{8c}$$

Control range : $$0 \le e_n \le E \tag{8d}$$

Service loss range : $$-x'_n \le l_n \le x''_n \tag{8e}$$

Rate range : $$-\Delta'_n \le T_{n+1} - T_n \le \Delta''_n \tag{8f}$$

for $n = 0, \cdots, N-1$

Cost functional: $$f = \sum_{n=0}^{N-1} \lambda_C a \tau p_n e_n + \lambda_L \gamma_n b(l_n) \tag{8g}$$

The above we call the *regular form* of the model, as distinguished from the normal form introduced later.

Note that the only thermodynamic parameters of the building needed to specify (8) are ε, related to the thermal time constant, and η / A, related to the equipment and the insulation. Computation of these values from first principles is non-trivial, but simple system identification procedures can be used to estimate them directly from observations of actual temperature behaviour.

If we assume a symmetrical, time-invariant tolerance range for service loss and a time-invriant ideal indoor temperature, i.e.

$$x'_n = x''_n = \hat{x} \text{ and } T_n^d = T^d \text{ for } n = 0, \ldots, N \tag{9}$$

then we can introduce the following *normal form* of the model using the normalized quantities defined in Table 2:

Energy balance : $$x_{n+1} = \varepsilon x_n + (1-\varepsilon)(w_n + \xi u_n) \quad n = 0, \ldots, N-1 \tag{10a}$$

Initial condition : x_0 (10b)

Service loss range : $-1 \le x_n \le 1 \quad n = 1,..., N$ (10c)

Control range : $0 \le u_n \le 1 \quad n = 0,..., N-1$ (10d)

Cost functional : $f = \sum_{n=0}^{N-1} \alpha \rho_n u_n + \beta_n b(x_{n+1})$ (10e)

The validity assumptions for the normal form are discussed in Section 7. Here, let us only note that the normal form has a number of advantages: (1) it is dimensionless, thus independent of measurement scaling; (2) all variables range over a common numerical scale; (3) all the characteristics of the usage are captured by the parameters α, β_n, ε, ξ and the function $b()$; and (4) it is suitable for developing a generic control scheme. Therefore, the normal form will be used in the rest of the paper. The rate range has been omitted from the normal form (10) only for simplicity, as it will also be neglected in the application below.

Table 2: Normalization of the space conditioning model

Normalized quantity	Symbol	Definition
service loss in period n	x_n	$(T_n - T^d) / \hat{x}$
outside temperature in period n	w_n	$(T_n^o - T^d) / \hat{x}$
power input (control) in period n	u_n	e_n / E
price in period n	ρ_n	$p_n / \hat{p}$
service loss coefficient in period n	β_n	$\lambda_L \gamma_n$
monetary cost coefficient	α	$\lambda_C a \tau E \hat{p}$
(steady state) temperature gain	ξ	$\pm \eta E / `A\hat{x}$ (+ : heat, - : cool)

Notice : $\hat{x}$, T^d are defined in (9)

3. ESTIA: A Computer-Based Controller for Space Conditioning

ESTIA, named after the ancient Greek godess of the home, is a real-time control scheme for space conditioning, based on the model of the previous section. Incidentally, the same word also means "burner" and "fireplace".

ESTIA aims at exploiting regularly available information on price, weather and inside temperature to achieve an overall balance between cost and comfort. Methodologically, ESTIA is a *certainty equivalent control* scheme, i.e., each time it makes a decision it uses the best available forecasts of weather and price as though they were certain. Randomness is then coped with by regularly revising the forecasts and the associated decisions. This choice of method has great computational advantages over the ideal dynamic programming solution and, although "suboptimal", yields good results as we shall later see. ESTIA comprises two levels of control:

1) A high level (HL) where it prescribes a space temperature to be reached by the end of the current hour, when it will prescribe a new temperature for the next hour and so on. At this level, hourly forecasts of price and weather are used for a horizon of one day.

2) A low level (LL) where the electric power input required required to achieve the target temperature is computed and applied every five minutes. So, weather fluctuations within the hour are compensated. Moreover, price fluctuations are taken into account and the nominal temperature trajectory over the hour prescribed by HL, is modified if necessary. In short, the low level acts as a regulator subject to high level directives.

The control period lengths of 1 hour at the HL and 5 minutes at the LL are only specific choices for the present implementation. In fact, ESTIA treats these time scales as parameters.

The overall structure of ESTIA is shown in Figure 2. The rectangular blocks represent the modules of ESTIA. MASTER is responsible for general coordination and for tuning the real-time controller. It furnishes the high level supervisor, HL, with all the system parameters, the control period lengths, the user's preferences and the probabilistic models for weather and price, the latter being determined by STAT. The high level functions include high level prediction and high level control, performed by HLP and HLC respectively. Using the weather and price models and the measurements taken, HLP predicts average hourly values for weather and price for the next N_H periods, where N_H is the length of the time horizon measured in high level control periods. At the beginning of each high level control period, HLC prescribes a power input (control) for the entire period and a target temperature to be reached by the end of the period. HLC always has a time horizon of fixed length N_H. LL coordinates the low level functions, including low level prediction (LLP) and low level control (LLC), and commands the equipment. LLP produces forecasts for price and weather for the remaining portion of the current high level period. LLC corrects the current high level control in each low level control period, with the objective of attaining the target temperature by the end of the current high level control period while responding to price and weather fluctuations. A distinctive difference from HL is that LL has a time horizon of variable length. The actual control commands to the equipment, issued periodically by LL, are precisely the controls determined by LLC.

The forecasts of future prices and weather could also be obtained from independent external agents.

4. Implementation of the High Level Controller

The method of certainty equivalent control is a far more important feature of ESTIA than the technique of two-level control. The latter merely expoits the computational advantage resulting from gradually refining the partition of the time horizon in discrete optimal control problems [2]. The degree of this advantage depends on the nature of the inputs and on the system parameters, chiefly the factor of inertia ε (related to the thermal time constant). For instance, with a small ε a one-level and a two-level scheme are equally satisfactory. In addition, the properties of the control scheme are derived from the certainty equivalent method and the nature of space conditioning [4]. Also the quality of the predictors HLP and LLP is determined by the choice of input models, a task assigned to STAT and totally independent of the control task. The most important part of ESTIA is the high level controller HLC, on which we concentrate in the rest of this paper.

Define the price and weather forecasts in period n to be equal to the expected values of future price and outside temperature conditional on their values in period n:

$$\bar{\rho}_i = E[\rho_i \mid \rho_n], \bar{w}_i = E[w_i \mid w_n], \quad i = n, \cdots, n+(N_H-1) \tag{11}$$

Assume that these forecasts are available at the beginning of each period n. The HLC scheme, in normal form, involves the following steps:

In each control period n:

1) Take the current forecast values of price and outside temperatures as certain, i.e., substitute $\bar{\rho}_i$ and $\bar{w}_i$ for ρ_i and w_i respectively.

2) Find the control sequence (power inputs)

$$\{ u_n, u_{n+1}, \cdots, u_{n+N_H-1} \}$$

which solves the deterministic problem:

$$J_n = \min \sum_{i=n}^{n+N_H-1} \alpha\bar{\rho}_i u_i + \beta_i b(x_i+1) \tag{12}$$

subject to:

$$x_{i+1} = \varepsilon x_i + (1-\varepsilon)(\bar{w}_i + \xi u_i) \quad i = n, \cdots, n+N_H-1 \tag{13a}$$

$$x_n : given \tag{13b}$$

$$-1 \le x_i \le 1 \quad i = n+1, \ldots, n+N_H \tag{13c}$$

$$0 \le u_i \le 1 \quad i = n, \ldots, n+N_H-1 \tag{13d}$$

3) Apply the control u_n. If there is a low-level controller, transmit u_n and x_{n+1} to it instead for further processing.

The problem (12), (13) is solved by combining two powerful methods with a proven, good performance record. The first is the *Multiplier method*, a combination of the penalty method with Lagrangian relaxation, which enables us to get rid of the inequality constraints while preserving the dynamic system equations. The second is a special version of *Newton's method* for unconstrained optimal control problems, which yields a fast solution of the problem resulting from the application of the multiplier method by replacing it with a sequence of linear, quadratic cost problems solved by means of the Riccati equation. It is beyond the scope of this paper to discuss these methods. The interested reader is referred to [2,3,8]. In particular, [3,8] explain the construction of the algorithms involved in great detail.

5. Behaviour of the High Level Controller

We now turn to the performance of the proposed space conditioning control scheme, in particular the high level controller, HLC, implemented as in the previous section. Evaluating HLC consists of examining the power input (control) and indoor temperature (service loss) trajectories determined by HLC ,as well as a number of comprehensive performance indices, as influenced by a set of relevant factors, and involves theroretical analysis, simulation and, eventually, field tests. Here we present features of HLC which are deemed significant from an application point of view. Results of simulation runs are provided in the next section. For proofs and discussion of modelling and design issues see [4].

The comprehensive performance indices used are:

- total (daily) cost, c($);

- percent total cost savings, Δ_o(%), over the "perfect comfort control" (λ_C=0);

- percent total cost savings, $\Delta_1(\%)$, over the "naive control" i.e. the case of a consumer with $\lambda_C > 0$ who naively sticks to the steady state indoor temperature that would be optimal if the price were equal to the daily mean value of the actual variable price; and

- indoor temperature swing $T_s(^oF)$, which is the difference between the daily maximum and minimum indoor temperatures prescibed by HLC.

The factors of interest here are mainly the price, the weather, the service loss timing weights and, to a lesser extent, the other system parameters. With regard to the inputs, we are concerned both with the effects of their entire daily trajectories and the sensitivity of the control scheme to input changes.

We first state an important **sensitivity result**: Given uniform service loss timing weights, the immediate target temperature in a given control period is insensitive to price fluctuations occuring after a number of $H+1$ periods, provided that the forecast price remains constant at the current level during the next H periods. In other words, on condition of a constant price forecast for the next H periods, equal to the current price, the future beyond these periods is irrelevant to the current decision.

The above statement describes a *Short-Sight Property* and the time span of H periods can be regarded as the *price vision range* of the control scheme. Computational evidence [4] suggests that the short-sight property enjoys a remarkably wider validity. Indeed, it is found to hold even if the price varies within the vision range. Moreover, computational experience [4] shows that the vision range length H is in practice almost entirely determined by the factor of inertia ε. Finally, it was found computationally that the immediate target temperature is also insensitive to weather changes occuring after a certain number of periods which we call the *weather vision range*. Figure 3 shows the relationship between the vision ranges and the thermal time constant observed computationally [4].

Both vision ranges are considerably shorter than one day, the length of the control horizon, for all practical values of the time constant. This is very important for a practical implementation of ESTIA: the *effective control horizon* is equal to the longest of the vision ranges, i.e. the price vision range, rather than the original, one-day horizon. So, the high level prediction and control tasks, HLP and HLC, are significantly reduced. Furthermore, if spot prices are specified 6-13 hours in advance, there is no relevant uncertainty in the price and consumers may take advantage of price variability at no risk.

A proposition equivalent to the short-sight property is the *Long Horizon Property* which concerns the settling of a steady state: When a constant price is forecast for a period of time greater than or equal to the vision range H and the service loss timing weights are uniform, then it is best to achieve and maintain a certain constant indoor temperature regardless of weather. This steady state temperature corresponds to a service loss x such that:

$$db(x)/dx = -\alpha\rho/\beta\xi \qquad (14)$$

It must be stressed that only a constant price forecast is needed to establish the long horizon property, with no restriction on weather. Therefore the long horizon property applies to flat rates as well as to time-of-day rates with long blocks of constant price levels, and provides a very simple control algorithm.

When the price varies it is impossible to completely dissociate price from weather effects. The general condition of the weather (i.e., cold or warm) together with the price trajectory determine the sequence of zero and non-zero controls and, to a large extent, prescribe the service loss trajectory. In particular, zero controls are assigned to neighbourhoods of the local maxima of price. These neighbourhoods are the largest allowed by the overall weather condition. The variation of the weather influences the service loss trajectory and determines the actual values of the non-zero controls. The effects of price and weather on the control and service loss trajectories are evidenced in the simulations of the next section, where it is clearly seen that price variability makes these trajectories non-trivial and calls for the use of computer control.

The thermal time constant of the building has a very strong bearing on the control and service loss trajectories, the performance indices and the effective control horizon. However, being a design, rather than an operational parameter it will not be discussed here any further. See [4].

The adjustable elements which can be used as "tuning knobs" of ESTIA, are the preference parameters λ_C, λ_L, γ_n and the penalty function $b(\,)$. In what regards the penalty function, we remark that the usual symmetrical forms $b(x) = |x|, x^2, x^4$ present an increasing deadband effect as their order increases, which results, naturally, in higher savings and larger indoor temperature swings. On the utility weights, λ_C and λ_L, we have to note that even with moderate values of λ_C, considerable savings are achieved with only small to moderate indoor temperature swings. This is demonstrated in the next section.

Finally, we mention three important cases of consumer attitude toward the timing of service loss and their respective models using the weights γ_n:

a. Uniform attitude:

$$\gamma_n = 1/N_H, \text{ for } all\ n \qquad (15)$$

This is the case of a building with invariable comfort requirements, e.g., a hospital.

b. Two-level attidute:

$$\gamma_n = \gamma \text{ for } n \text{ in } N_O, 0 \text{ for } n \text{ in } N_V \quad (16)$$

This is the case of a commercial building with periods of occurancy, N_O, and vacancy, N_V.

c. Three-level attitude:

$$\gamma_n = \gamma_W \text{ for } n \text{ in } N_{W,} \gamma_S \text{ for } n \text{ in } N_S, 0 \text{ for } n \text{ in } N_V \quad (17)$$

This is the case of a dwelling, a typical day in which comprises periods of wakefulness, N_W, sleep, N_S, and vacancy, N_V.

All three cases are demonstrated in the next section.

6. Simulation Studies

The operation of the proposed space conditioning control scheme, in particular the high level controller, has been simulated under realistic conditions.

Two consumers are considered, one in Madison, Wisconsin on a day of December and the other in Dallas, Texas on a day of July. The consumers are assumed to have identical preference parameters, λ_C, λ_L, γ_n, $\hat{x}$, and penalty function, $b(\,)$, as well as identical building thermal time constant, TC, and overall thermal conductivity, A. It is also assumed that space conditioning is their only controlled electricity usage, so $\lambda_C+\lambda_L=1$. The parameter values of the two systems, chosen to represent reasonable designs and user preferences, are listed in Table 3.

The price is taken to be the actual hourly marginal operating cost of a certain utility in the midwestern U.S. in representative days of December and July. The marginal operating cost is a good approximation to the notion of a spot electricity price. The weather consists of hypothetical temperature trajectories based on average values and average swings quoted by the U.S. Weather Bureau for the specific locations and months. It is assumed that temperature extremes occured at 4 a.m. and 1 p.m.. The price and weather inputs are shown in figures 4 and 5.

Each consumer assumes in turn the roles of a hospital, a commercial building and a dwelling. This is done to demonstrate the three main attitudes toward service loss timing. Strictly speaking, the technical characteristics of the consumers correspond to private houses. Therefore, the simulation results refer to alternative usage patterns of given buildings, rather that to actual hospitals or commercial buildings. For each consumer role three cases are solved: a) the perfect comfort case with x_n = 0 during periods of occupancy (corresponding to $\lambda_C = 0$); b) the naive control case with $x_n = x$, the steady state value corresponding to the specific $\lambda_C > 0$; and c) the ESTIA case with variable x_n.

The corresponding values of the performance indices are listed in Tables 4 and 5. The normalized control (fractional power input) and indoor temperature trajectories, resulting from the simulations in the case of cooling (consumer in Texas) are shown in Figures 6, 7, 8, and 9.

Finally, a word about the choice of the monetary cost utility weights λ_C. The values of λ_C are chosen such that the consummer would be prepared to tolerate the maximum service loss, i.e., $x_n = \pm 1$ for all n, if the price were equal to some fixed, perhaps unrealistically high value, p^*. This price, here taken to be 10 times the ceiling price of electricity, is called "deadlock price" and offers one way of assessing the utility weights [4].

7. Discussion

The features of HLC mentioned in section 6 are clear from the simulation results. It is, though, worth stressing that with a variable price the optimal control and indoor temperature trajectories are indeed non-trivial.

Tables 4 and 5 indicate that ESTIA enables important savings with acceptable temperature swings. At worst, a swing of $6^o F$ and savings of 10.9% were observed. These results were obtained with a mildly cost-minded (λ_C=0.67) or even neutral (λ_C=0.5) consumer attitude, quite adverse weather conditions and a rather usual variation of utility marginal operating cost (daily maximum/minimum ratio < 3) which we took as the electricity price. Therefore, under a spot electricity tariff, ESTIA achieves significant savings from electric space conditioning with rather small sacrifice in comfort.

The simulation results suggest that better overall performance (higher savings and lower swings) is possible in the summer than in the winter. This is reasonable since the weather relative to the ideal indoor temperature is milder in Texas in the summer than it is in Wisconsin in the winter. An actual comparison of the two cases, however, would require marginal operating costa data of the respective local utilities.

The percent savings, Δ_0 and Δ_1, may be interpreted as follows. By definition, Δ_0 is the total savings over the perfect comfort case (λ_C=0), while Δ_1 is the savings over the "naive control" case. The "naive control", by prescribing a constant service loss, achieves "steady state" savings over the perfect comfort case, because it uses a λ_C>0. Thus, Δ_1 represents "dynamic" savings which must be attributed to the variation of price. The

Table 3: System parameter values for simulation studies

Parameter	Wisconsin,heating	Texas,cooling
$A(kW/^oF)$	0.14	0.14
$TC(h)$	25	25
$E(kW)$	21	4.2
η	1(efficiency)	2.5(COP)
ξ	10	-5
$\tau(h)$	1	1
$T^d(^oF)$	70	75
$\hat{x}(^oF)$	15	15
$\hat{p}(\$/kWh)$	0.20	0.20
$p^*(\$/kWh)$	2.00	2.00
λ_C	0.67	0.50
$b(x)$	x^2	x^2

Table 4: Wisconsin, heating: performance summary

Case	c($)	$\Delta_o(\%)$	$\Delta_1(\%)$	$T_s(^oF)$
uniform γ	4.88	14.3	13.8	6
2-level γ	3.81	24.0	23.5	2.5
3-level γ	4.66	11.6	10.9	4.5

Table 5: Texas, cooling: performance summary

Case	c($)	$\Delta_o(\%)$	$\Delta_1(\%)$	$T_s(^oF)$
uniform γ	0.44	18.8	16.8	2.2
2-level γ	0.34	30.3	28.5	2.4
3-level γ	0.40	17.0	15.0	3.4

"steady state" savings are then $\Delta_o - \Delta_1$, much smaller than the "dynamic" savings. Therefore the benefits of a decision model based control, such as ESTIA, are obtained mainly under a variable electricity price.

The rather arbitrary choice of symmetrical service loss tolerance range and penalty function is ex post justified by noticing that overheating or overcooling were always pretty small.

The significance of the short-sight property should be underlined. In actual implementations of spot pricing, the price review period is usually longer than half a day, which makes price forecasting practically unnecessary in those cases, and the implementation of the control algorithm requires less data and computing resources. In view of the large effective elimination of uncertainty, the suboptimality of the certainty equivalent control scheme is less of an issue. In fact, the control scheme does take, to some extent, uncertainty into account, by using a rolling horizon and applying only the immediate control computed each time.

Even though ESTIA is a certainty equivalent controller, the attitude of consumers towards the risk arising from uncertainty in external signals (price and weather) could have been described within our modelling framework. Yet a direct account for risk would have resulted in computational overheads of questionable payoff.

In what regards service levels, it was found that significant monetary savings are obtained with only small deviations from the ideal indoor temperature. At the modelling level, the service loss range can be set narrow enough, so as to represent no significantly dangerous outcomes, without discounting the accuracy of the model.

Money-related risk, on the other hand, is ignored by the cost functional used, which is linear, therefore risk-neutral, with respect to money. The working hypothesis underlying this choice of cost functional is that the electric bill is not a dominant part of the consumer’s total spending. When this assumption does not hold and the attitude towards risky monetary outcomes must be explicitly taken into account, the various uses of electricity can no longer be controlled independently. This is due to the fact that a non-linear cost functional with respect to money has to be employed, and the monetary cost argument actually is the total cost for all electricity uses. The computational implications are very significant: a coordination problem arises with subproblems corresponding to the various electricity uses, the solution of which requires a multiple of the total solution time for the subproblems. Moreover, this does not include the computing time increase due to considering alternative outcomes. Apart from these computational consequences, accounting for risk would not have affected the basic structure and properties of the subproblem considered in this paper, namely space conditioning control.

The simulation results reported here are only indicative of the performance of ESTIA under realistic conditions. For cost-benefit analyses, extensive simulation should be carried out with realistic variable price and weather data at least over a period of one year.

References

[1] American Society of Heating, Refrigerating and Air Conditioning Engineers, *Handbook of Fundamentals*, 1977.

[2] Bertsekas, D.P., "Newton's Method for Linear Optimal Control Problems," Proc. of IFAC Symposium on Large Scale Systems, Udine, Italy, June 1976, pp. 353-359.

[3] Bertsekas, D.P., *Notes on Nonlinear Programming and Discrete-Time Optimal Control*, LIDS-R-919, Massachusetts Institute of Technology, January 1982.

[4] Constantopoulos, P., *Computer-Assisted Control of Electricity Usage by Consumers*, Doctoral dissertation, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, May 1983.

[5] Constantopoulos, P., et al., "Decision Models for Electric Load Management by Consumers Facing a Variable Price of Electricity", in *Energy Models and Studies*, (Lev, ed.), North-Holland Publishing Co., 1983, pp. 273-291.

[6] Ihara, S., and F.C.Schweppe, "Physically Based Modeling of Cold Load Pickup," *IEEE Trans. of Power Apparatus and Systems*, Vol. *PAS-100*, No. 9, Sept. 1981, pp.4142-4150.

[7] Keeney, R., and H. Raiffa, *Decision with Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, 1976.

[8] Kort, B.W. and D.P. Bertsekas, "Combined Primal-Dual and Penalty Methods for Convex Programming," *SIAM J. Control and Optimization*, Vol. 14, No. 2, Feb. 1976, pp. 268-294.

[9] Massachusetts Institute of Technology Energy Laboratory, New Electric Utility Management and Control Systems: Proceedings of Conference Massachusetts Institute of Technology, Center for Energy Policy Research and Electric Power Systems Engineering Laboratory, Report MIT-EL-79-024, June 1979.

[10] Morgan, M.G., and S.N. Talukdar, "Electric Power Load Management: Some Technical, Economic Regulatory and Social Issues," *Processing of the IEEE*, Vol. 67, No. 2, Feb. 1979, pp.241-313.

[11] O’Rourke, P., and F.C.Schweppe, "Space Conditioning Load Under Spot or Time of Day Pricing," IEEE PES Summer Meeting, 1982, 82SM433-1.

[12] Schweppe, F.C., "Power System 2000," *IEEE Spectrum*, July 1978, pp. 42-47.

[13] Schweppe, F.C., et al., "Homeostatic Utility Control," *IEEE Trans. on Power Apparatus and Systems*, Vol. *PAS-99*, No. 3, May/June 1980, pp. 1151-1163.

[14] Schweppe, F.C., et al., "Homeostatic Control for Electric Power Usage, *IEEE Spectrum* July 1982, pp.44-48.

[15] Schweppe, F.C., et al., Spot Pricing of Electricity, Kluwer, 1988.

[16] Tabors, R.D., Schweppe, F.C., and Caramanis, M.C., "Utility Experience with Real Time Rates", IEEE Trans. on Power Systems, Vol. 4, No. 2, May 1989, 463-471.