SCHEDULING OF ELECTRICITY CONSUMPTION UNDER SPOT PRICES

by

BAHMAN DARYANIAN

M.S., Massachusetts Institute of Technology
Technology & Policy Program
(June 1986)

M.S., Massachusetts Institute of Technology
Mechanical Engineering
(February 1980)

B.S., Massachusetts Institute of Technology
Mechanical Engineering
(June 1977)

SUBMITTED TO THE DEPARTMENT OF
MECHANICAL ENGINEERING
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 1989

Signature of Author ______________________
Department of Mechanical Engineering
February 1, 1989

Certified by ______________________
Dr. Richard D. Tabors, Thesis Supervisor

Accepted by ______________________
Professor Ain A. Sonin, Chairman
Departmental Graduate Committee
Department of Mechanical Engineering

SCHEDULING OF ELECTRICITY CONSUMPTION UNDER SPOT PRICES

by

BAHMAN DARYANIAN

Submitted to the Department of Mechanical Engineering
on February 1, 1989 in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy

ABSTRACT

This thesis is concerned with the development of scheduling algorithms for residential electricity consumers under spot price based electricity rates. Spot pricing is an *optimal* time varying rate reflecting the true cost of electricity supply. It accomplishes various objectives of load and demand side management through the market signal of electricity prices. In practice, each customer will be given 24 hour schedules of prices a few hours prior to the hour of application. Price levels will be constant for each hour.

Electricity usage is classified according to the storage capabilities of end use devices and processes. For each class, various static control logics are presented. These are conceptually simple and easy to implement.

For dynamic control, storage type devices and processes are modeled as discrete time linear dynamic systems. The state and output equations comprise multiple storage elements, controllable (electricity) and deterministic exogenous (weather) inputs, and product outputs (heat or hot water). The control problem, scheduling of electricity usage at minimum cost subject to input and output constraints, is formulated as a linear programming problem. The need for a fast and near real time control algorithm has motivated the search for alternative optimization techniques.

It is found that thermal systems are asymptotically stable positive dynamic systems. Based on this result, an algorithm is introduced for the control of discrete time linear dynamic thermal systems subject to linear costs and bounds on any variable. This algorithm is to be used in a Feed Forward Control scheme. The algorithm finds a particular initial feasible solution, and at each iteration moves linearly in the feasible space to a new solution (constraint) set with a lower total cost. Each movement is based on a linear combination of impulse response vectors. The algorithm terminates when no further cost improvements are possible. The algorithm takes simple forms for first-order systems. The computer codes written for the case study problems result in optimal solutions. Depending on the size of the problem, they are up to 10 times faster than a commercially available linear programming code. Further improvements in efficiency are possible.

The first case study focuses on the scheduling of electricity for heating and cooling of buildings. Using published estimated data for building parameters, a 2 resistance, single storage (massive internal structure) building is studied in some detail. The output variable (inside air temperature), is a function of the state variable (storage temperature). Four different, yet general, price patterns were studied. Impact of different thermal parameters and problem constraints on the total savings, compared to equivalent flat rates, were examined. Depending on the price patterns and the price levels, savings associated with optimal scheduling varies from zero to 30%. Results also demonstrate that associated with each price pattern, there exists a unique optimal design values for the storage size.

Also considered is the scheduling of electricity use for dual element hot water heaters. In this case, the output is the storage variable itself, which results in a very simple version of the algorithm. Depending on the price patterns and price levels, results demonstrate possibility of up to 50% savings under spot pricing relative to comparable flat rates.

To study the effects of uncertainty, an stochastic dynamic programming algorithm is developed, which is computationally slow. The most important result for a single storage system without losses is: for highly uncertain exogenous inputs, the certainty equivalence assumptions, i.e. using expected values of exogenous inputs, result in suboptimal control.

Thesis Committee: Dr. Richard D. Tabors (Supervisor)
Title: Assistant Director and Principal Research Associate
Laboratory for Electromagnetic and Electronic Systems, MIT

Dr. Thomas B. Sheridan (Chairman)
Title: Professor of Engineering and Applied Psychology
Department of Mechanical Engineering, MIT

Dr. Shahryar Motakef
Title: Assistant Professor of Mechanical Engineering, MIT

Dr. Ignacio J. Pérez-Arriaga
Title: Professor of Electrical Engineering and Director of the
Instituto de Investigación Tecnológica at
Universidad Pontificia Comillas , Madrid, Spain

Acknowledgments

In the name of the common spirit which bonds all souls,
and its boundless benevolence which makes each one unique.

This project was started under the joint supervision of Professor Fred Schweppe and Dr. Richard Tabors. Professor Schweppe died in July of 1988. He was brilliant in mind and soul. His joyous personality and his deep intelligence created a truly unique being. It is sad that he is no longer with us. Yet, it is good to know that if I happen to be in Main, I may see his spirit sitting on a hill top, looking at the sunset, smiling at all the wonders of the world, thinking and singing at the same time.

Speaking of good spirits, Richard Tabors has my deepest gratitude. For the past few years, there could have been no substitute for his guidance, support, and patience. Without his direction and supervision, this thesis would not have come into being. Roger Bohn of Harvard Business School has been a continuous source of good advice. My earlier work with him provided the foundation for this thesis. Tom Sheridan facilitated this long journey through his genuine concern and guidance. Throughout the years, I have learned many lessons from Shahryar Motakef, whether or not he meant to teach them. I thank him for his continuous support, critical comments, and valuable suggestions. In the short time I have known him, Ignacio Pérez-Arriaga contributed a lot to this thesis. His comments were extremely helpful in clearing some of the theoretical problems, and in shedding light on the interpretation of the results.

Professor Richard de Neufville, chairman of the Technology and Policy Program, and Professor Ain A. Sonin of the mechanical engineering department at MIT made it possible for me to continue my studies after a few years of being absent from the school. Dan Flag expertly disciplined the unruly computers whenever the need aroused. The whole operation could not have been run without the help of the Ladies of the LEES (Laboratory for Electromagnetic and Electronic Systems) headquarters.

I also wish to thank the support of Laurence Carmichael, the project manager of EPRI RP 2830-5 under which this research was carried out.

The environment at MIT has been very hospitable, all because of my good office-mates and friends at LEES and beyond. My friends outside MIT have always been a source of happiness for me and my family.

And finally, I owe it all to Hassan, Parvin, Jilla, Eskandar, and all of my extended family in that little known Shangrila in Azerbaijan, Iran.

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

1.1 Thesis Overview

This thesis considers the problem of optimal scheduling of electricity usage under spot price based rates. The overall results of this study are applicable to residential, commercial, and industrial customers; however, the case studies consider residential customers only.

Spot pricing is an optimal time varying rate scheme for pricing of electricity. Also called flexible, dynamic, real time pricing, or time varying tariffs, it is based on the concept that prices should change with time to reflect the variable cost of electricity supply. Spot pricing provides a market based incentive for modification of supplier and consumer behavior, and as such acts as a unifying tool in the realization of various demand side management objectives.

This thesis limits itself to a particular scheme of pricing where prices are different for each hour (or any fixed period), and are communicated to the customers at least 24 hours in advance of the application. A short overview of spot pricing is presented in the following section of this chapter.

The third section of this chapter considers various types of end uses of electricity consumption, and classifies them according to their possible response behavior.

The last section of this chapter discusses previous work in this area, and their relations to this thesis.

Residential electricity consumption can be classified as belonging to one of two broad categories. One type possesses some form of storage capability which enables the postponement of service from the actual time of electricity consumption. Electric water heaters, space cooling, and space heating fall into this category. Price responding algorithms for this category are non-trivial and are the main concern of this thesis. Appliances such as dish washers, laundry machines, and dryers comprise the other type in which electricity is consumed at the same time that the service is provided. The control schemes for this category of end uses are simple to conceptualize and easy to implement. The simpler group of end uses are considered first and the classifications and the control schemes are described in some detail in Chapter 2. Also discussed is the static control of storage type end uses.

For dynamic control, the storage type electricity usage is modeled as a discrete time linear dynamic system. Service considerations such as comfortable temperature range in a building and electricity usage capacity impose bounds on the input and output variables of the system. The only cost considered is the cost of electricity usage based on spot price rates. Thus, from a theoretical perspective, this thesis considers a particular class of linear programming problems. However, the desire for a fast and efficient algorithm for an on- line *feed forward* control of electricity usage has motivated the search for an alternative optimization method.

Chapter 3 discusses the properties of thermal systems. It shows that linear thermal systems are asymptotically stable positive dynamic systems. Hence, the impulse response vector for these systems has a particularly simple form. Based on these results, Chapter 4 introduces a general algorithm for the control of asymptotically stable positive dynamic systems. For a finite time horizon, the system equations, with bounds on inputs, outputs, and storage variables, define a convex feasible solution space. The algorithm first finds a particular initial feasible solution set for the inputs. Then, at each iteration it moves as much as possible in a feasible linear direction to a new solution point. It is shown that any feasible direction is a combination of some elementary feasible directions. The algorithm considers only those movements which reduce the total cost. It terminates when no elementary feasible direction with net reduction of total cost is found. The general framework of the algorithm is provided for single-input single-output, deterministic, discrete-time, linear time variant, thermal systems subject to linear costs and bounds on any variables.

Chapter 5 is devoted to the study of heating and cooling in buildings under spot prices. It considers various building types and determines their particular dynamic models in relation to the framework of the application of the general algorithm. It then focuses on a particular thermal model applicable to simple buildings, namely a single thermal storage building. A case study is performed using published estimated data for the thermal properties of these type of buildings. The general algorithm is simplified when applied to this model. Of particular interests are the impact of price patterns and the thermal properties of the building on the total savings relative to comparable flat rates of electricity.

Chapter 6 considers the case of dual element electric water heaters. It provides a model for the water heater, and defines the optimal scheduling of hot water production in terms of the control of the volume of the hot water produced. Again, the general algorithm is simplified further for this particular class of applications.

Chapter 7 drops the assumption that the exogenous inputs are deterministic. It presents a preliminary study on the question of stochastic inputs, and it provides a basis for further work in that area. A stochastic dynamic programming framework is developed for a simple storage type problem and various assumptions are tested. The most interesting result is that, when the exogenous inputs are highly uncertain, the assumption of the *certainty equivalence* results in sub-optimal solutions. Under the certainty equivalence assumption, the optimization under stochastic conditions is held to be equivalent to a deterministic optimization with the uncertain inputs fixed at their expected values.

Finally, Chapter 8 provides a brief summary of the thesis, lists the main conclusions, and makes recommendations for future research.

1.2 Spot Pricing of Electricity

1.2.1 Why spot prices ?

Electricity production and transmission is a regulated economic activity in the United States and most other countries. Electricity prices, generally fixed and determined by regulatory commissions, do not reflect the real-time economic costs of providing electricity service. The result is that many generating and transmitting facilities are either over-worked or idle. This economic reality has always been a source of concern and activity by the electric utilities in the U.S. and abroad [G2]. Presently, it is an important element forcing a reassessment of ways in which electric utilities operate and do business [N1,N2,N3].

Many electric utilities have been considering various methods of influencing and changing customer behavior in ways that would result in the utilities' desired load shapes [M4]. The resulting activities and programs have evolved into more sophisticated demand-side management initiatives such as spot pricing [G2].

A utility's load shaping objective depends on the physical and financial characteristics particular to that utility. Gellings [G1] classifies these objectives as peak clipping, valley filling, load shifting, strategic conservation, strategic load growth, and flexible load shaping. It is possible for two different utilities to pursue two differing objectives. However, each of these objectives can, in economic terms, be viewed as the desire for a reduction in operating and capital costs.

There are a variety of tools and methods at the disposal of utilities to meet these objectives. Some are individualized alternatives, which, as described by Delgado

[D4], take into account the specific type of the electricity customer and particular end uses of electricity. Demand charges, interruptable rates, and time of use (TOU) rates are three tariffs which influence customers' electricity use behavior. However, none of the rate systems in use reflect the real-time variable service cost of electricity. Hence, as incentives they provide only limited information to the customers, and as a result, produce only limited response.

Spot pricing is a time varying pricing system for electricity which reflects the actual time and space-dependent costs of electricity generation, transmission, and distribution [B6,B5,B8,C2]. The notion of spot prices also includes purchase of power from independent power producers as well as sale of power to customers.

The basic idea of spot pricing is that if at any time, cost of generation increases, either because of increased demand or unscheduled outages, prices should increase as well. Raising up the prices would not only reflect the higher costs of additional generating units being brought on line, it would also cause a decrease in the electricity demand. Some customers would experience a loss of benefits as a result of increases in the real cost of electricity and would attempt to reduce their consumption accordingly. In this respect, spot pricing can induce spontaneous supply and demand equilibrium by acting as a control feedback mechanism within electric power system.

The concept of spot pricing in its present form was introduced by Schweppe, et all [S2] under the name of homeostatic utility control. An earlier but less exact concept of responsive pricing had been proposed for general utility commodities by Vickrey [V1]. The most up to date and comprehensive reference on spot pricing of electricity is the recent book be Schweppe, Caramanis, Tabors, and Bohn [S4].

In one possible implementation, a typical customer would see prices changing every hour. However, spot prices would be communicated 24 hours in advance through the mass media or private communication channels such as a phone line. An energy conscious residential or industrial customer may depend on a sophisticated microprocessor based energy management system to shift most of its load to the times of lower prices. A less sophisticated residential customer may respond only when prices become unbearably high by turning off the non-essential appliances. Although the degrees of benefits may vary for different class of customers, it is expected that in time, experience and economic considerations will extend the opportunity of more flexible responses to all classes of customers.

The following list includes some of the expected benefits of spot pricing of electricity:

- Spot prices, reflecting the true cost of electricity, result in a more efficient allocation of resources in the society.

- Spot pricing directly benefits the electricity suppliers, because utility and non-utility generating units would operate more efficiently, and additional units would be introduced into the network as they become economically profitable.

- Customers benefit directly by knowing the real cost of electricity at different periods, and by planning their consumption behavior in a manner that takes advantage of price variations.

- Electricity customers with a consumption pattern that tracked peak production costs, would not be subsidized by other customers with a different load pattern.

- The inherent control characteristics of spot pricing would achieve more comprehensively the objectives of other traditional demand-side management tools, central dispatching, and rate structure techniques.

- Rotating blackouts and involuntary rationing (interruptable rates) would not be necessary because of the self-adjustment of electricity supply and demand. This would result in an added freedom of choice for customers who can voluntarily curtail their consumption if economic rationality does dictate so.

- Independent producers of electricity from wind power, solar power, cogeneration, and other unconventional sources can be more easily integrated with the electricity network system. Electricity from an independent source will be introduced to the network when prices are economical. Their economic contribution would be subject to market forces rather than regulation and legislation.

- Spot pricing provides a framework both for the short-term operational decisions and the long-term investment decisions of suppliers and customers [C1].

- Spot pricing provides a vehicle for partial deregulation of the industry if so desired.

- Spot pricing could reduce the need for new capacity.

In summary, If implemented properly, spot pricing of electricity can accomplish, single handedly, all the goals of various and disparate demand-side management programs.

1.2.2 Spot prices in theory

From a mathematical point of view, economic efficiency is equivalent to the maximization of some social welfare function. A theoretical derivation of optimal spot prices is provided by Bohn [B7], and by Caramanis, Bohn, and Schweppe [C2]; and expanded further by Bohn, Caramanis, and Schweppe [B8] to reflect the location-specific constraints within the network. A complete and detailed theoretical treatment is given in Schweppe, Caramanis, Tabors, and Bohn [S4]. The social welfare function considered is comprised of terms which reflect net economic benefits to the suppliers and consumers and is subject to some network-associated constraints.

Theoretically derived optimal spot prices vary continuously over time and space; thus reflecting the stochastic nature of supply and demand, electrical energy balance, and physical network constraints. More specifically, spot prices are the sum of the following components:

- Variable operating costs component, which includes the marginal fuel, operation, and maintenance cost of electricity generation. This component is referred to as the "system lambda" in most of the literature on utility systems.

- Energy balance quality of supply component, or quality of supply for generation, which reflects the cost of meeting the physical constraint of maintaining an equality or balance between electricity supply and demand. It is equal to the prorated cost of additional capacity needed to meet unserved demand; or equivalently from the customers' point of view, it is equal to the prorated value of unserved energy. In case there is any unused generating capacity in the system, this component is zero. Otherwise, it is added to the price in order to induce a decrease in demand.

- Transmission loss component, which reflects the cost of dissipated energy, or resistive loss over the network. This component is location dependent. It is frequently represented as proportional to the square of the load, and as a result its value becomes more pronounced as the transmission load increases.

- Transmission quality of supply component, or quality of supply for network, which reflects the voltage magnitude and power flow constraints. This component is also location dependent. It is a function of stability limits and also the probability of network burn-out due to over-loading of the network. Since the probability of network outage increases with load, this component of spot prices also becomes more significant as the transmission load increases.

All of the above components vary with time and location, and as a result so does their sum. At off-peak hours the first component of the spot price is dominant. At peak hours other components become significant.

1.2.3 Spot prices in practice

At the present time, spot pricing is in an experimental phase, though some industrial development rates closely approach implemented spot price rates. Comprehensive studies of spot pricing from utilities' perspectives have been performed for power companies in Wisconsin by Caramanis, Tabors, and Stevenson [C3]; and for power companies in California by Schweppe, Caramanis, Tabors, and Flory [S1].

Spot pricing does not completely do away with regulation. There would still be a need for a central authority which would regulate the electrical energy markets, maintain and develop electrical networks, periodically evaluate the spot price based rates, and communicate these rates to the suppliers and customers. Bohn, Golub, Tabors, and Schweppe [B10] have proposed creation of spot markets for bulk power as a partial substitute for the deregulation of electricity generation.

Theoretical spot prices fluctuate rapidly in time, and instantaneous changes in its components leave no room for a continuous evaluation and communication of prices. In addition, on-line adjustments to prices by consumers on a short notice is rather impractical. Therefore, theoretical prices can only serve as a framework for a more practical rate structure. For this purpose Caramanis, Bohn, and Schweppe [C2] provide a methodology for the evaluation of spot prices which can be utilized in practical situations.

An important issue is what kind of rate structures should consumers see and how often should price changes be communicated to the customers. Bohn [B7] gives a theoretical presentation that demonstrates the tradeoffs between transaction and communication costs and benefits of more frequent price updates. Given the rapid pace of technological innovations, these costs are expected to decrease substantially within the coming years. In this connection, Flory and Parker [F1] provide an overview of communication and metering technology for electricity spot pricing.

For the purpose of this work, and as suggested in the literature, it is reasonable to assume some form of arrangement in which a customer or a homogeneous class of customers is provided with some form of electricity price forecasts in advance. A price forecast cycle can be a day, a week, a month, or any agreed-upon time horizon. Furthermore, each price forecast cycle or time horizon would be divided into smaller periods such as minutes or hours. Price levels would be constant within each period, only changing in a step-wise manner between the periods within a price forecast cycle. Note that in this thesis, these price forecasts are taken to be deterministic.

1.3 Customer Response to Spot Prices

1.3.1 Customers and spot prices

Spot pricing of electricity provides the opportunity for the electricity customers to schedule their usage in a manner that results in reduced costs of consumption relative to a comparable flat rate. The savings are possible because of the price differentials among different periods. A customer may schedule to use non-essential appliances when prices are low, and use high priced electricity only if necessary. In addition, having a storage capability, enables the customers to enjoy the service even when prices are high by delaying the time of service use from its time of production.

From the utilities' point of view, scheduling of electricity by a significant portion of the consumers reduces demand during peak hours and increases consumption during hours of low demand, and consequently, helps the utilities to reduce costs and to increase efficiency [S4,B9].

A typical customer would see prices varying every hour which would be communicated 24 hours (or a week) in advance through mass media or private communication channels such as a phone line.

The main point is that customers have the flexibility to schedule electricity usage in order to capture the benefits of price differentials. The feedback aspects of spot pricing have no effect if customers are not responsive. What is envisioned is that the residential customer, once in a while, would specify the overall strategic rules, i.e. criteria and constraints that define the limits of acceptable service; and then, the local controlling system would automatically operate the various end-use devices subject to time varying spot prices, ever changing weather conditions, and service constraints set by the customer.

It is expected that industrial customers will find the benefits of spot pricing more attractive than smaller residential customers. Bohn's initial study [B4] has indicated the potential for significant savings by industrial customers. References [D1] and [D2] also provide a theoretical evidence for responsiveness and savings associated with spot pricing for a real industrial customer. Although the material of this thesis are applicable to industrial customers, the case studies are developed for residential or commercial thermal processes.

1.3.2 Price responding algorithms and end-use devices

Various end use devices have different behavior characteristics, and thus, different implications for response under spot pricing. The following list presents a preliminary classification of end use devices according to the expected response behavior under spot pricing. The choices for response are:

I- To schedule the operation at the minimum cost periods for appliances which only use electricity and which have no storage capabilities.

II- To switch to a substitute fuel when electricity price become higher than the price of the substitute fuel. This is an on-off operation, except that the operation at the off periods is continued using a substitute fuel.

III- To schedule electricity consumption load at the lowest price hours subject to production and storage capacity constraints and the deterministic or stochastic schedule of end product demand.

I- End use devices with no storage or fuel substitution: The first situation applies to the type of appliances which have no storage capability, and which can not use other energy substitutes for electricity. In this category we can include dishwashers, washing machines, dryers, etc. Electricity usage by these appliances can be controlled by simple decision logics. Absence of storage simplifies the control action, because the service of the appliance use is obtained at the same time that electricity is consumed. The decision logic can be as simple as turning the device on only when the prices are lower than some pre-specified threshold. This and relatively more sophisticated control actions are described in Chapter 2.

II- End use devices with fuel substitution capabilities: This situation is the least likely to be encountered in a residential setting. Possible examples are dual fuel stoves, burners, and water heaters. The decision logic would continuously compare the electricity spot prices to the cost of the substitute fuel (taking into account the efficiency), and then switch the device to the electricity or the fuel depending on their relative costs.

Examples are the dual fuel water heaters used in the residential and the commercial sectors, and the electric/gas ovens used in the glass industry. The same concept also applies to industries with self generation or cogeneration, where the electricity is switched from the grid to in-house supply.

III- End use devices or processes with storage: This category includes thermal devices such as space heating and cooling, supplementary thermal (hot or cool) storage, electric water heaters, refrigerator/freezer compressor and defrost control, swimming pool heaters, and water bed heaters.

What unites these devices or processes is their ability to store the end product (heat and hot water). The time at which the service is provided and the time of electricity use do not have to be the same. The required algorithms are more involved and are based on *dynamic logics*, because the decision at a certain hour will be based on the spot prices and the states of the system at previous and future hours. Dynamic logics may allow an increase in energy usage while reducing costs and improving benefits. Many industrial processes can also be modeled as storage processes as described in references [B7], [D1], and [D2].

The essential idea here is to use electricity during times of low price to heat the house and the water heater and build up stored heat and then to allow them to *float* through the high price times.

Chapter 2 describes some simple static logics for storage devices. A more sophisticated dynamic logic is discussed in Chapter 4.

1.3.3 Other examples of storage-type customers or producers

In addition to the type of end-use devices described above, many other industrial customers or even some small electricity producers have storage capabilities which can be utilized in order to respond to spot priced based electricity buyer and seller rates. For producers, the output or the end product of the process is electricity, and from a mathematical point of view, their problem is one of maximization of their sales. Some examples as mentioned in previous references are:

- **Finished goods storage by an air liquefaction plant:** These plants use electricity intensive processes to liquefy air into its constituent parts. These plants already have large storage tanks in place in order to smooth out operations in times of power interruptions and demand fluctuations. The electricity usage rating and storage size determine the bounds on the operations.

- **Municipal water plants:** Water is pumped into a water tank from a reservoir. The timing for replenishing of the tank is flexible. The input is the electricity usage as determined by water flow.

- **Low head hydro generation:** Independent power generators can also be subject to spot prices when selling to the main power grid. The reservoir level can be adjusted slightly by the producer to provide additional storage, and turbines are driven mostly at times of high prices. The size of the reservoir determines the maximum storage size.

- **Independent waste burning generators:** The input is the waste and the output is the electricity and waste storage size is limited.

- **chilled water storage using ice in an office building:** The chiller size determines the maximum production capacity, and the ice tank size limits the storage.

Many other similar examples should exist in electricity intensive industrial applications.

1.4 Previous Work in Response Scheduling Under Variable Prices

This section presents an overview of some related works in customer behavior modeling and formalization of response to time varying prices.

Most of the models considered in literature define storage type processes as linear dynamic systems. Scheduling of electricity usage under spot pricing is formulated as a minimum cost optimization, or control under linear costs. The earliest work is due to Tsitsiklis [T2] which provides the same general discrete time dynamic model as the one used in this thesis. Tsitsiklis also introduces an efficient algorithm for scheduling under spot prices for a scalar (one-storage) system with losses, where the controlled output is the storage level. The methodology in the present thesis, when applied to a first order system with losses and the storage level as the controlled variable (electric water heater), reduces to Tsitsiklis' algorithm. In fact the code used in the water heater study later in this thesis was written independently, and without the knowledge of Tsitsiklis' work. The algorithm is described in Appendix C in the context of the case study of Chapter 6. The algorithm is faster than the simplex method applied to the same problem.

Previous work by the present author [D1] and [D2] provides an algorithm for a single storage system without losses, which in fact, is a special case of the single storage system with losses. Hence it is again a special case of both the general algorithm presented in Chapter 4 of this thesis, and that of Tsitsiklis [T2]. The proof of optimality provided in Reference [D1] is different from the one provided by Tsitsiklis for the storage system with losses. This algorithm is also applicable to scheduling of electricity usage of water heaters when losses are lumped with the demand for the hot water as explained in Chapter 6.

Tsitsiklis [T2] also discusses the case of multi-storage (higher-order) systems, and introduces decomposition techniques which break down the master problem into linked smaller size storage optimization subproblems. Decomposition is possible when special requirements for matrix sparseness and invertibility are met. Each subproblem can be solved using more efficient techniques, however the master problem is solved by the simplex or equivalent methods. In contrast, the present thesis develops a different methodology for multi-storage systems based on the block representation and impulse response of the linear systems. All the aforementioned problems and also the general problem of this thesis are linear programming problems.

O'Rourke and Schweppe [O1], [O2] consider the modeling of space condition load under spot pricing. Their thermal model is again a single-storage (first-order) linear system. However, they consider only periodic weather and price data which enables them to expand the variables in a Fourier series for a fixed interval, and algebraically solve for the optimal power input. The results give some insights into the behavior of space conditioning load under spot pricing. However, the periodicity assumptions are too restrictive to be realistically acceptable. Their study can be used as a planning tool for the evaluation of inputs of price responding HVAC control.

A more general analysis of electricity usage by customers is carried out by Constantopoulos [C6]. He also considers thermal linear dynamic systems. However, he introduces the customers' utility functions and their explicit valuation of service and comfort. Therefore, the cost or the objective function for the optimization problem becomes nonlinear, and thus, more complex nonlinear optimization techniques are employed. Furthermore, only a first order model is studied, and *certainty equivalence* for stochastic situations is assumed. Chapter 7 of the present thesis employs a stochastic dynamic programming method to demonstrate that the certainty equivalence may not hold for similar systems with highly variable exogenous inputs.

A different approach by Chang [C4] introduces an algorithm based on the method of the Lagrange multipliers to solve an optimization problem with quadratic cost function which combines load shifting and self-generation.

Finally, in his recent work, David [D3] formulates a multi-storage optimization problem very similar to the present work. However, only on-off or integer values are assumed for the power inputs, and established linear and integer programming techniques are employed. Also discussed are models of other types of customer loads which have not been explicitly discussed in the present work.

Chapter 2

CONTROL LOGICS AND END USE DEVICES

2.1 Introduction

Static control applies both to storage type and non-storage type end use devices. In this chapter, steady state conditions are assumed for storage type devices, and the savings inherent to dynamic scheduling of power inputs are ignored. Dynamic control for storage type devices is discussed in more detail in the following chapters. Non-storage end use devices require simple static control logics for operation under spot prices. In contrast with storage-type processes, the service can not be stored for later use, and a decision is made for either termination of service at a particular period, or switching to a substitute fuel. Only the former case is considered here. Examples are limited to those of residential customers. Except for a few minor exceptions, the materials presented in this chapter are completely based on a recent paper by Schweppe, Daryanian, and Tabors [S5]. Again, it is assumed that the forecasts of future hourly prices and weather variables are provided to the customers in advance.

This chapter first describes the main features of a complete energy management system which provides the actual control of end-use devices. Then, it suggests a classification of end use devices according to the service they provide, and then presents a list of possible static control logics for each category and discusses the related issues. Dynamic controls for storage type devices are listed, but no algorithms are provided. The last section of this chapter discusses the utility-customer interaction.

2.2 Energy Management Systems

The residential price responding algorithms discussed here are to be implemented in the residential *Load Control Emulation System*, or LCES, which is a hardware/software system supported by the Electric Power Research Institute (EPRI). LCES is designed to assist utilities in their experimentation with various demand side management techniques [E3]. Spot pricing is only one of the many load management components of the LCES. There exist other similar price responding energy management systems which have already reached implementation stage [P2], [T1]. Description of the envisioned energy management system is general enough to be applicable to the storage-type end use devices.

There are four main functions performed by a complete energy management system:

- **Tactical Control:** Provides the real time control of the end use devices.
- **Behavior Modeling:** Supplies the behavior models to the Tactical Control which uses them to find the best operational modes.
- **Usage Diagnostics:** Provides the users with the data and statistics of the costs and modes of operations for each end use device.
- **Strategic Planning:** Supplies the Tactical Control with the service criteria and limits which are in turn furnished by the customers.

Tactical Control is divided into two parts. The first part consists of the Control

Logic. The decisions for the types of actual signals sent to the end use devices are determined by the Control Logic. In some cases these signals are directly sent to the customer who turns the devices on or off. The second part consists of the State Estimator which determines the current on- off status of various devices. The information is used by the Tactical Control as the initial data or the state values for the current decision interval. Tactical Control must operate continuously.

Behavior Modeling provides the necessary information to the Tactical Control concerning the models of the systems to be controlled. These include both engineering parameters such as coefficients for the state equations (heat transfer and thermal capacity coefficients for heating and cooling), and also the load data such as hot water usage patterns and weather information. For simple non-storage end use devices, the relevant data include power rating and kWh/Usage values. Some of the models require statistical averaging techniques to collapse the observations made during longer periods. A time interval of one day would be sufficient for the update of the data by the Behavior Modeling.

Usage Diagnostic provides bookkeeping information by assembling data on energy usage and behavior patterns. These are furnished to the user to assist him/her in choosing the criteria and limits, and also to provide evaluations of customer's previous strategies. Information from Usage Diagnostic can be invoked on a monthly basis.

Strategic Planning is where the customer inputs the limits of the service requirements into the energy management system. For example, the customer sets the acceptable range of the inside temperature, and the price thresholds above which appliances must be turned off. Information is to be supplied once at the start up, and then at any time at the demand of the customer.

These functions can be developed into more sophisticated expert type systems as further experience is gained from the field operations. At earlier stages of implementation some of the functions would be performed manually by the customer or the utility personnel. It is expected that at the early stages, only the Tactical Control is going to be automatically operational. Other functions would be automated at later stages of implementation.

2.3 Classification of End Use Devices

Statically controlled storage type and non-storage type end use devices can be classified according to the type of service they provide. A partial listing may include the following:

- **Thermal Storage, Temperature Controlled:** These include HVAC, swimming pool heater, hot tub, water bed, etc. Usage can be rescheduled for these devices, however, the service can be used at a different time.
- **Water Heating :** This category is similar to the previous one, except that the controlled variable is the volume of the hot water not its temperature.
- **Periodic Use:** This group includes defrosting of refrigerators and swimming pool pumps. They are activated only once in a while.
- **Reschedulable Appliances:** These include dishwashers, washing machines, dryers, etc. For this class of end-use devices, both service and usage can be rescheduled.
- **Discretionary Activity Devices:** These include stoves, vacuum cleaning, and devices used in hobbies and chores. Here also, both service and usage can be rescheduled.
- **Non-Reschedulable Appliances:** These include the lights, TV, and HiFi. For these devices, service and usage can only be reduced rather than be rescheduled.

The main difference of the storage-type devices (the first two categories) from the others is that for the former the service and the usage need not occur simultaneously. The dynamic models and the associated control algorithms for the storage-type devices are described in more detail in the following chapters.

2.4 Tactical Control Logics

This section provides a list of possible static control logics for each category of end use devices. A particular energy management system may use one or more of these logics in its Tactical Control, or may even add other logics not mentioned here. The following notations are used in this chapter:

k	Time index, one hour step
$p[k]$	Price of electricity during hour k, \$/kWh

Thermal Storage: Temperature Controlled

Static models of thermal storage ignore the cost savings associated with consideration of price differentials among different periods. The static control logic considers only the electricity price of the current period. Define:

$T_a[k]$: Thermostat Temperature setting at hour k
T_{min} : Minimum acceptable temperature
T_{max} : Maximum acceptable temperature
T_e[k] : External temperature at hour k
$U[k]$: Electrical energy used during hour k

The acceptable bounds on the temperature are provided by the customer to the Strategic Planning, and they can be time varying. The cost per hour, \$/Hour, is $p[k]\ U[k]$. The external temperature refers to the outside temperature for HVAC and to the room temperature for the water bed. Three possible control logics are:

- TSS1: The device is turned off when \$/Hour exceeds some prespecified threshold. The control is overridden when the temperature falls outside the acceptable range.

- TSS2: The customer is warned when $/Hour exceeds the threshold. The customer can decide whether to turn off the device or not.
- TSS3: This logic is similar to TSS2 except that future values of $/Hour are provided to the customer to help in planning of a decision strategy for the day.

Thermal storage for the building can be either due to the main structure and the furniture in the house, or due to the special storage devices such as hot rocks or ice. Other factors for HVAC, such as inside humidity and the rate of temperature change are ignored here.

Water Heating

Under normal operations, electric water heaters maintain a constant temperature volume of hot water in the storage. The heating elements are activated when the water is cooled either because of heat loss, or because of cold water inflow when hot water is drawn from the storage. Only dual element water heaters are considered here. Three possible static control logics are:

- WH1: The lower element is allowed to operate only within the time intervals of lowest cost. Thus, if it takes a water heater 2 hours to fill up the storage, then it does it during the continuous 2 hours of the day when the total cost is minimum. This will be done only once a day.
- WH2: The customer specifies various usage times, and the lower element is allowed to operate during the minimum cost intervals between the specified usage times. If there are N specified usage times, then the tank is filled with hot water N times during the day.

The upper element is not controlled. This provides a measure of safety and a limited provision of hot water when the need arises.

Periodic Use Requirement

This category includes defrosting of refrigerators and swimming pool pumps. They are activated only once in a while, but the exact time is not important. Two possible control logics are:

- PUR1: Given the required time interval for the operation, the device is activated during the hours of minimum cost.
- PUR2: This involves a more sophisticated logic whereby a limit is also put on the total cost of operation.

The second control logic can be made more sophisticated by considering tradeoffs between the cost and the run times.

Reschedulable Appliances

This category includes dishwashers, washing machines, and dryers. To minimize damage to the device and the service, no appliance is cut off during its cycle of operation. Let's define the following:

- kWh/Usage : Energy used during a single *average* run cycle
- $/Usage : Cost of a single *average* run cycle.
- $/Deferred : Money saved if a single run-cycle is deferred from the present hour k to some future hour during the next 24 hours when the price is the minimum.

Values of the thresholds are provided by the customer to the Strategic Planning. Six possible control logics are provided here:

- RA1: Don't start operation if $/Usage exceeds the specified threshold.
- RA2: The customer is warned that $/Usage exceeds the limit. The customer then decides whether to turn on the appliance or not.

- RA3: The customer is provided with the values of \$/Usage for the next 24 hours, who then decides for the best time of operation.

- RA4: Similar to RA1 except that the start up is based on \$/Deferred instead of \$/Usage.

- RA5: Similar to RA2 except that the customer is warned if \$/Deferred exceeds the limit.

- RA6: Similar to RA3 except that the customer is provided with the future values of \$/Deferred.

- RA7: The energy management system automatically reschedules usage to the minimum price time between the present time k and a prespecified future time.

For performance studies and possibly for control purposes, it may be necessary to learn the kWh/Usage of each appliance. This value is not necessarily a constant and may change in time. Therefore, if direct metering or off-line calibration is used, then the values must be averaged. Albeit, the simplest thing to do is to read the name plate of the appliance, but this may not prove helpful.

Discretionary Activity Devices

This category is rather broad and includes end use devices associated with cooking and hobbies. Control logics similar to the ones suggested for the category of reschedulable appliances may be used. However, determination of a single *average* run time is difficult due to the variability of usage time. The simplest control actions are:

- DAD1: Don't turn on the device if the price is higher than a threshold.

- DAD2: Inform the customer of the schedule of the prices and let him/her to decide when to cook and do hobbies and chores.

The customer provides the threshold information to the Strategic Planning.

Non-Reschedulable Appliances

This category includes lights, TV, etc. The only control action is whether to do without the service if prices are high enough. Again two simple control actions are possible:

- NRA1: Don't turn on the device if the price is higher than a limit.
- NRA2: Inform the customer of the schedule to the prices and let him/her to decide when to use lights and watch TV.

2.5 Dynamic Control Logics for Storage Type Devices

What unites storage type devices or processes is their capability to store the end product (heat and hot water) for a later use. The time at which the service is provided and the time of electricity use do not have to be the same. The required algorithms will be more involved and will be based on *Dynamic Logic*, since the decision at a certain hour will be based on the spot prices and the states of the system at previous and future hours. The *Dynamic Logic* may allow an increase in energy usage while reducing costs and improving benefits.

Dynamic Logic for Space Heating and Cooling

Three possible dynamic control logics are:

- Minimize bill over next N hours subject to Constraints on inside temperature bounds.
- Minimize bill plus cost of deviation over next N hours.
- Minimize bill to achieve a specific temperature at a specific future hour.

Some of the information such as spot prices and weather data required for the implementation of dynamic logic will come from a central information processing computer. However, as noted later, the scheduling algorithm is based on the state equation of the system involved. The system has to be identified and its parameters have to be estimated using a method such as the ARMX structure. These parameters will be unique for each particular system, for the simple fact that each physical system such as a house, has its own particular characteristics.

Dynamic Control Logics for Electric Water Heaters

Again only the dual-element water heaters are considered. A simple static decision logic is to cut off the electricity when \$/Gallon exceeds a certain limit. This limit must be higher for the upper element or there may even be no limits for the upper element. The cut-off limit may also be based on \$/Gallon- Deferred. The information required for static control are spot prices, and possibly, power ratings of the heating elements.

Dynamic decision logic may have the following objectives :

- Minimize bill for lower element subject to meeting the predicted hot water demand.
- Minimize bill for lower element subject to having a full tank of hot water by a specific future hour.

The dynamic decision logic requires the information on spot prices, inside and outside temperatures, thermostat settings, heating element ratings, and also the system parameters for the state equation of the water heater. In addition, future demand for hot water need to be predicted from time-based thermostat and heating element data.

2.6 Utility-Customer Interaction

The final service provided by the utility, in order to be comprehensive, will include different types of control options. Each particular utility will offer one or a subset of options to customers, and the customers will make the final choices. It is assumed that the end use devices will undergo no *major* modifications, and no fancy sensors are contemplated. The overall control actions are to be kept simple and practical and within the general means of the residential customers.

One decisive factor contributing to the success in modification of customer behavior will be the utility-customer interface. Since the various end use devices in a house do not share common characteristics, such an interface should have the capability of addressing the needs of each end use requirement without overwhelming the customer's patience and without inhibiting his/her interactions.

At initial stages of implementation, all the communication between the utility and the customer can be done by phone or mail. The utility can call the customer and ask the limits and criteria selected by the customer, and suggest actions at the same time. It can also mail, on a monthly basis, the processed information on customer's savings and performance. With the advent and penetration of personal computers and modern communication tools, utility-customer interaction will evolve into something more sophisticated. Figure 2.1 illustrates the lines of communication and control in the customer-utility interactions.

It is expected that a future utility-customer interface will have the following *positive* characteristics:

- Will require no hardware development.
- Will be easy to install.
- Will allow many levels of sophistication.
- Will look like an ordinary personal computer which many people already

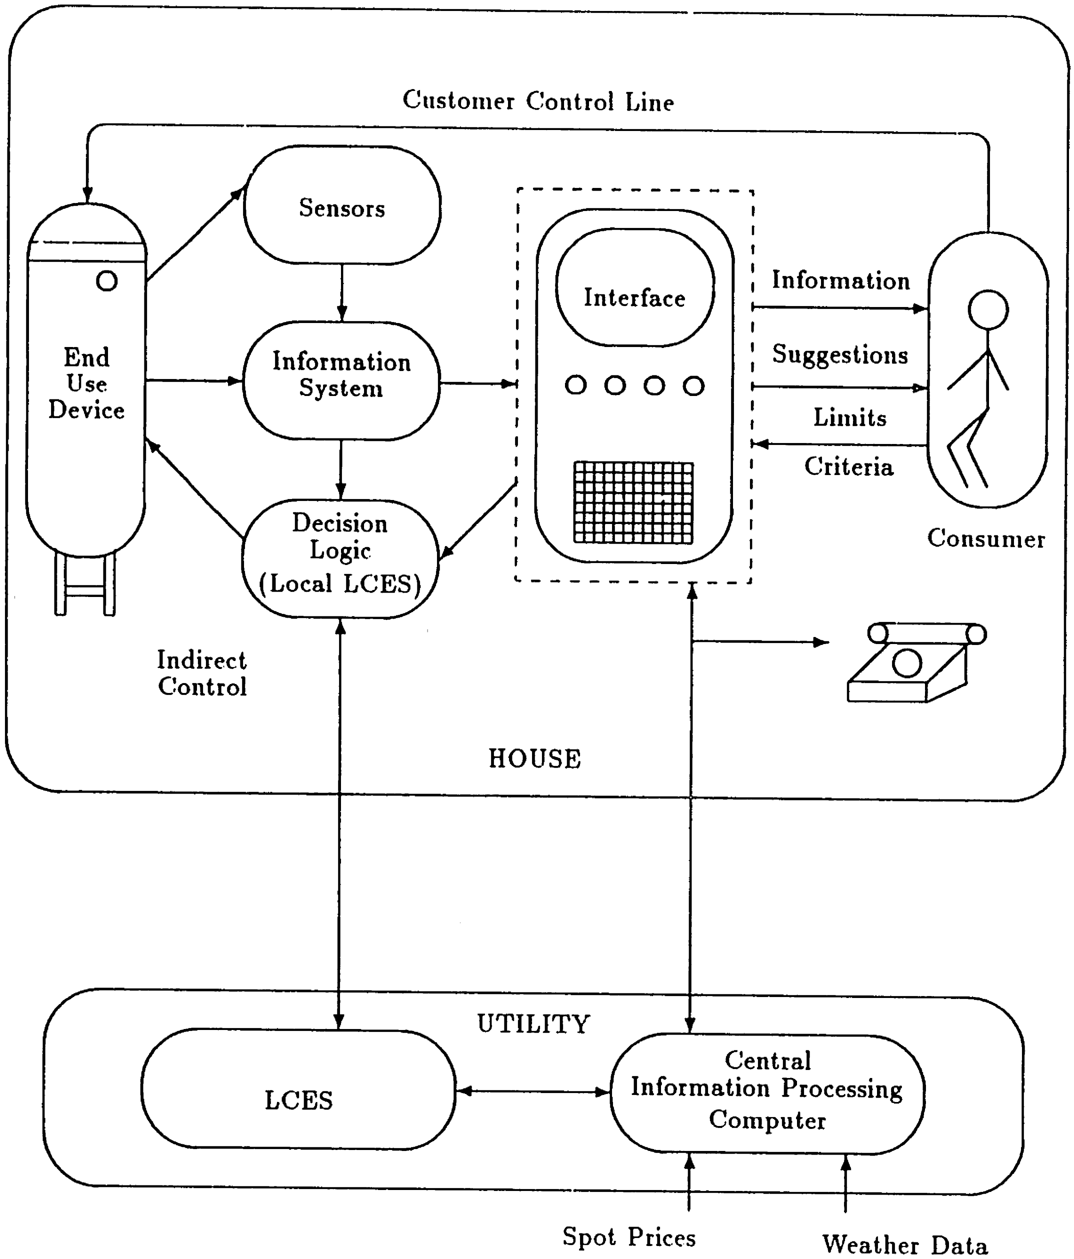

Figure 2.1: Utility-Customer Interface

have.

On the negative side, it is expected that:

- Cost per point will be unnecessarily high.
- Some features will be considered an overkill for many of the applications.

As shown in Figure 2.1, the interface will provide the customer with information from two sources. One set of information will come from the utility, which would include spot prices, weather information, and utility bills. Each utility can also provide some analysis of customers performance such as a monthly or yearly diagnosis such as $/End-Use, kWh/End-Use, dollars saved, and suggestions for future action. Another set of information will be provided from the customer's own end use devices in terms of amount, duration, and times of electricity usage. Each customer will also provide inputs to the interface, such as acceptable performance limits and service criteria which are to be utilized by the utility and/or the local LCES for control and performance computation purposes. An important feature would be the ability of the central LCES to down-load programs to the local LCES.

The customer, if he/she so desires, can reserve the right to interfere at any time and to assume direct control of each end use device. The interface must also allow for frequent and easy changing of customer's limits and criteria. An added feature would be to display various options and strategies for indirect control.

It is expected that a residential customer will apply direct control over the end use devices of the first and second category which have no storage capabilities. However, the decision of when to operate will be based on the information put out by an *expert system* type interface. Such information may include only spot prices, or the result of some computation suggesting the best times to use the appliances during each day.

The indirect control is more logical in case of storage type processes such as space heating and cooling and electric water heating. For example, the customer obtains information from the interface on the trade off between service and costs, i.e. how many more dollars saved for an additional discomfort in having wider deviation in the house temperature. Then, knowing the expected trade off, the customer sets the minimum and maximum acceptable temperature and the control system takes over from there to take advantage of spot price differentials and utilize the thermal storage in order to minimize total costs without violating the limits set by the customer. The interface software should include the necessary built-in programs to explain the purpose of the system, to provide an introduction and tutorial on the use of the system, and to act as a help facility when the customer needs it.

Chapter 3

DYNAMIC THERMAL SYSTEMS

3.1 Introduction

Space heating and cooling, together with water heating, account for a significant portion of the residential and commercial electricity loads, and to some extent that of the industrial loads. They constitute the main components of the demand side management programs as described in References [E1], and [M4]. Most of the previous studies have considered consumer response to time of use and interruptable rates.

Some recent studies consider space conditioning under time of day pricing. O'Rourke and Schweppe [O2] assume a cyclical price and weather patterns which limits the applicability of their findings. The study by Constantopoulos [C6] uses a first-order thermal model only, but the savings vs. service trade offs are explicitly taken into account. The result is a non-linear optimization methodology.

A work related to this thesis is that of Tsitsiklis [T2] which introduces an algorithms for a minimum cost scheduling of a state-controlled first order system. His algorithm is a special case of the general algorithm introduced in Chapter 4 of this thesis. In the case study of Chapter 6, the state-controlled first-order model is used to describe the dynamics of a dual element water heater. The results are based on an algorithm similar to that of Tsitsiklis' but written without the prior knowledge of its existence. Tsitsiklis also describes an algorithm for higher-order systems. The optimization methodology involves diagonalizing of matrices and matrix inversions which are computationally more complicated than the routines used by the algorithms of this thesis. It is possible that the two methodologies are some how related.

An important characteristic of space heating and cooling and water heating is their storage potential, which provides some flexibility in the ways customers can respond to spot prices. In this study the emphasis will be on the inherent storage capability of buildings within their internal and external thermal masses. However, the general framework provided here can be easily extended to cover the auxiliary and external heat and cool storage facilities.

Control of a thermal system requires a model which describes its dynamic behavior. It is impossible to simplify behavior of all thermal systems–such as the dynamics of buildings of many types and characteristics–into a single and all-encompassing model. However, it is possible to provide general simple mathematical models which can be applied to any situation with some elementary modifications.

A major distinction between space heating/cooling and water heating is that for the former, the variable to be controlled is the temperature. Space heating and cooling may also involve a multitude of storage mediums. In contrast, for water heating, the controlled variable is the volume of the stored hot water. In addition, water heating uses a single storage only, and therefore, it can be modeled as a first order dynamic system. Apart from this physical distinction, the dynamic behavior of both types of thermal systems can be described by discrete time state space equations, which is the form utilized by the algorithms introduced in this thesis.

The thermal dynamic model of each building should provide a reasonable match for the complexities of its structure. This necessitates utilization of control algorithms which are tailored to the complexity of the model under consideration. Furthermore, even if two buildings look similar, they may have different thermal characteristics. Hence, each building will have its own unique model parameters.

Although linear programming techniques can be applied in the case of linear models, however, the need for a fast and efficient scheduling algorithm that can be used both in on-line control and also for large scale simulation, has motivated the search for an alternative optimization methodology. The result is a general algorithm presented in Chapter 4 which is applicable to any asymptotically stable positive dynamic systems. Thermal dynamic systems belong to the same category.

Some thermal systems are simple enough to afford a faster and less complex derivative versions of the general algorithm. The case studies in Chapters 5 and 6 are based on such models.

This chapter introduces the necessary background needed for the understanding of the algorithm of the next chapter. It starts with a discussion of the models of thermal systems. An example is provided wherein the appropriate state equations are derived for a representative building thermal system.

Next, based on the generalization of the example, and the mathematics of the positive dynamic systems, the special properties of thermal systems are discussed. These properties, and also the general properties of linear systems are utilized in the development of the general algorithm.

The chapter ends with a brief discussion of the suggested control scheme under uncertainty and the problem of parameter estimation for the thermal models involved. The last issue in not a main concern of this thesis, but it is included because it is a very important component of any practical scheduling scheme.

For the simplicity of discussions, most of the developments in this thesis refer explicitly to space heating. However, all the arguments and results of the thesis are also applicable to space cooling.

3.2 Descriptive Overview

3.2.1 Space heating and cooling

An ordinary thermostat-controlled heating system maintains the temperature inside a house within a narrow band around a desired temperature which is set by the residents of the house. The width of this band is determined by the hysteresis characteristics of the thermostat itself. The times of electricity use are independent of prices, which are normally fixed at a constant level anyway; instead, they depend on the set temperature, external temperature, and the heat transfer and thermal storage characteristics of the house.

The external and internal structures in any building act as thermal masses which store thermal energy. The stored thermal energy is returned to the surrounding areas at a rate determined by the internal and external temperatures, and also by the thermal characteristics of the building. Thus, although the rate and the timing of heat input into the storage can be completely specified in an arbitrary manner, the stored heat is transferred back to the air or the outside at a rate determined by the dynamic equations and the external temperature and the state of the system. This is another example of a situation where it is easier to give than to take.

The thermal dynamics of the house indicate that the heat input to the house is transferred both to the outside and also to the thermal masses in the building. If at any moment the heat input is stopped then the heat stored in the internal masses is slowly returned back to the air.

Under spot pricing, and if the future prices and the external temperatures are known and the behavior of the system can be predicted, it is possible in principle to use more electricity (than required in normal operation) in times of lower prices in order to substitute for the electricity usage at later times when prices are high. For instance, when the electricity prices are low, the building can be heated up to the maximum temperature allowed. This would increase the stored heat by heating up the thermal mass in the building. Later on, when prices are high, the heating system can be turned off to let the temperature to coast down to the minimum allowable temperature. In this manner, the stored heat replaces the actual heating at times of high prices.

The thermal masses represents the potential size of the thermal storage elements available in the house. These storage elements are leaky, and they loose heat. The severity of the heat loss, or the capability for retention of the thermal energy, is determined by the heat transfer resistances among the inside air, the thermal masses, and external ambient temperature.

The problem, from the point of view of economics, is to decide when to turn on or off the heat in order to minimize total cost of electricity used for heating or cooling, while maintaining a comfortable range of air temperatures in the house. For an economical scheduling of electricity under spot prices it is suspected that a wider comfortable temperature zone is required compared to the narrow hysteresis range of the thermostat. The bounds of the comfortable range, denoted by T_{min} and T_{max}, indirectly limit the amount by which the thermal storage can be charged up. For instance, if the inside air temperature is required to stay absolutely constant, then the on-off cycle of the heating system becomes extremely short. However, if there are no limits on the maximum attainable temperature inside the house, then the heating system can be turned on at full blast at times of lower prices to drive the temperature of the internal mass as high as possible in order to have some reserve of thermal energy in the future when prices become higher.

It should be emphasized that temperature constraints apply only to the air temperature T_a and not the other internal thermal mass temperatures. Therefore, there are no explicit constraints on the temperatures of the storage elements.

Another physical constraint is the maximum rate of heating. For example, in a house with a powerful heating system, it would be possible to heat up the house to a higher temperature in a shorter period of time compared to a similar house with a less powerful heating system.

3.2.2 Dual element water heaters

Only the dual element electric water heaters are considered in this study. The reason is that the two thermostats are used as marks on the extent of storage to be controlled. As noted before, the volume of the hot water is the variable to be controlled. Under normal operation, a full tank of hot water at a fixed temperature is maintained. As hot water is drawn from the top, cold water enters from the bottom and is heated up to the set temperature.

Under spot prices, and predictable demand, it is possible to control the lower heating element such that the tank is filled mostly at times of lower prices. Chapter 6 provides a more detailed description of the workings of a dual element water heater.

3.3 Dynamic Models for Space Heating and Cooling

Thermal model of a building can be described mathematically in terms of a system of differential or difference equations. These equations describe the dynamic system which is to be controlled. The control action is provided by the heat input into the building. The performance criteria is the total electricity cost which is to be reduced as much as possible as long as the controlled variables of the system, inside air temperature in our case, are kept within the given acceptable bounds.

The basic assumption is that the thermal characteristics of any building can be represented by a few simple *lumped* system elements, thus requiring only a manageable number of simple thermal coefficients. More complex buildings require more of these simple elements in their models. The dynamics of lumped systems can be described by simple differential or difference equations. The complexity of the system is manifested by the order of these equations. For example, a building with two distinct thermal mass areas can be modeled as a two storage system, which in turn can be represented by a second order differential or difference equation.

The modeling process involves the derivation of the state equations for the lumped thermal system. Possible methodologies include:

Bond Graphs This is a modeling technique based on the energy and information flow within the physical system. First developed by Professor Henry Paynter of MIT, it is a powerful tool equally applicable to mixed mechanical, electrical, fluid, and thermal systems of any level of complexity [K1]. The modeling results in the derivation of equations in state-space form.

Linear Graphs This methodology is based on the similarity and equivalence of basic elements in various physical systems, and it is also well suited to the modeling of thermal systems [S6].

Equivalent Circuit Models This technique is based on the representation of the thermal system in terms of equivalent electrical elements. Circuit theory is then used to derive the equations [S7].

Energy Balance Methods This method is based on simply writing the heat transfer equations between various elements, and combining the equations. This method is suitable for very simple thermal systems.

Each method should result in equivalent state equations for the system. The number of independent thermal storage elements determine the order of the system. For our purposes, it is essential that the system be modeled as a discrete time system. This is possible if the dynamics of the system are considered for sufficiently short time periods, during which, the state of the system is assumed to remain constant. Each of the above methods can be used to derive differential equations representing the dynamics of the system studied. Those equations are then transformed into discrete time equations, by assuming that *input level* remains constant for small time steps. Coefficients of the equivalent difference equations will include *matrix exponential* terms, as shown below.

However, the last method above can be used to derive the difference equations directly. In this case, it is assumed that *all variables* remain constant during the time step. Thus, the coefficients in the discrete time equations have simple forms. It is also the easiest to understand in physical terms. Therefore, for simplicity and clarity of presentation, only the last method is used in this chapter. The difference in the two representation is described in the following section.

Lumped thermal systems are more restrictive than mechanical, fluid, electrical, and other dynamic systems. Two general observations worth mentioning are:

- Thermal systems can be completely represented by source, resistive, and capacitive elements similar to their electrical counterparts. However, no similar *inductor* element exists in thermal systems.
- Thermal systems can be classified as special cases of *asymptotically stable positive dynamic systems*. Mathematically, this means that the state transition matrix for the thermal dynamic systems is always positive, and has a stable equilibrium point [L3].

The significance of these and other properties of thermal systems is described in the following sections by means of an example.

3.4 Discrete Time and Continuous Time Equations

If a time-invariant system is modeled as a discrete time dynamic system, the state-space equations can be written directly in the form of difference equations. The basic assumption is that *all variables* remain constant during small time steps. For a single-input single-output system, the result in state-space form is:

$$\mathbf{X}[k+1] = \mathbf{A}\ \mathbf{X}[k] + \mathbf{B_u}\ u[k] + \mathbf{B_w}\ \mathbf{W}[k] \tag{3.1}$$

$$y[k] = \mathbf{C}\ \mathbf{X}[k] + d_u\ u[k] + \mathbf{D_w}\ \mathbf{W}[k] \tag{3.2}$$

where, $\mathbf{X}$ is the vector of state variables, y is the output, u is the control input, $\mathbf{W}$ is the vector of exogenous variables, and k is the time index. Note that the input$u[k]$ is given in units of *energy*. For time invariant systems, the coefficient matrices have simple constant forms.

If the system is modeled as a continuous time system, equations for a single-input single-output system in state-space have the following form:

$$\frac{d\mathbf{X}(t)}{dt} = \mathbf{A}^C\ \mathbf{X}(t) + \mathbf{B_u}^C\ u(t) + \mathbf{B_w}^C\ \mathbf{W}(t) \tag{3.3}$$

$$y(t) = \mathbf{C}^C\ \mathbf{X}(t) + d_u^C\ u(t) + \mathbf{D_w}^C\ \mathbf{W}(t) \tag{3.4}$$

The subscript C indicates that the coefficients are for the continuous time representation. Note that the input $u(t)$ must be in units of *power*. If it can be assumed that *input level* changes only at equally spaced time steps T, then the continuous time state equations can be discretized. In the following, the notations kT and $(k+1)T$ are used instead of k and $k+1$. The discrete time representation takes the form:

$$\mathbf{X}((k+1)T) = \mathbf{A}^D(T)\ \mathbf{X}(kT) + \mathbf{B_u}^D(T)\ u(kT) + \mathbf{B_w}^D(T)\ \mathbf{W}(kT) \tag{3.5}$$

$$y((k+1)T) = \mathbf{C}^D(T)\ \mathbf{X}(kT) + d_u^D(T)\ u(kT) + \mathbf{D_w}^D(T)\ \mathbf{W}(kT) \tag{3.6}$$

The subscript D indicates that the coefficients are for the discrete time representation. Note that the coefficient matrices depend on the time step T. These coefficients can be expressed in terms of the continuous time coefficients. For the continuous time system, give an initial state $\mathbf{X}(0)$ and inputs $u(t)$ and $\mathbf{W}(t)$, the solution for $\mathbf{X}(t)$ is:

$$\mathbf{X}(t) = e^{\mathbf{A}^C t}\, \mathbf{X}(0) + e^{\mathbf{A}^C t} \int_0^t e^{-\mathbf{A}^C \tau} \mathbf{B_u}^C\, u(\tau)\, d\tau + e^{\mathbf{A}^C t} \int_0^t e^{-\mathbf{A}^C \tau} \mathbf{B_w}^C\, \mathbf{W}(\tau)\, d\tau \tag{3.7}$$

where the *matrix exponential* $e^{\mathbf{A}t}$ is equivalent to the *state transition matrix* and has the following form:

$$e^{\mathbf{A}t} = \mathbf{I} + \mathbf{A}t + \frac{1}{2!}\mathbf{A}^2 t^2 + \cdots + \frac{1}{k!}\mathbf{A}^k t^k + \cdots \tag{3.8}$$

or in compact form:

$$e^{\mathbf{A}t} = \sum_{k=0}^{\infty} \frac{\mathbf{A}^k t^k}{k!} \tag{3.9}$$

Writing the solution expression for periods $(k+1)T$ and kT and then multiplying the second result by $e^{\mathbf{A}^C}$ and subtracting from the first result, we obtain:

$$\mathbf{X}((k+1)T) = e^{\mathbf{A}^C T}\, \mathbf{X}(kT) + \int_0^T e^{\mathbf{A}^C \lambda} \mathbf{B_u}^C\, u(kT)\, d\lambda + \int_0^T e^{\mathbf{A}^C \lambda} \mathbf{B_w}^C\, \mathbf{W}(kT)\, d\lambda \tag{3.10}$$

where $\lambda = T - t$. By comparison, we find:

$$\mathbf{A}^D(T) = e^{\mathbf{A}^C T} \tag{3.11}$$

$$\mathbf{B_u}^D(T) = \left(\textstyle\int_0^T e^{\mathbf{A}^C t} dt\right)\, \mathbf{B_u}^C \tag{3.12}$$

$$\mathbf{B_w}^D(T) = \left(\textstyle\int_0^T e^{\mathbf{A}^C t} dt\right)\, \mathbf{B_w}^C \tag{3.13}$$

Thus, the discrete time coefficients are given in terms of the time step T and continuous time coefficients. For very small time steps, the higher order terms in matrix exponential can be ignored, resulting in the following:

$$\mathbf{A}^D(T) = \mathbf{I} + T\, \mathbf{A}^C \tag{3.14}$$

$$\mathbf{B_u}^D(T) = T\, \mathbf{I}\mathbf{B_u}^C \tag{3.15}$$

$$\mathbf{B_w}^D(T) = T\, \mathbf{I}\mathbf{B_w}^C \tag{3.16}$$

One problem with this formulation is that the control input in many applications is either on or off (bang-bang control), and usually has a constant power rating.

Suppose the power rating of the control input is ν power units. Then for each time step, the control input is going to be on for a portion of the time T. Let us denote the proportionality factor as $\eta[k]$ for period k. The discrete time equations become:

$$\mathbf{X}((k+1)T) = e^{\mathbf{A}^C T}\,\mathbf{X}(kT) + \int_0^{\eta[k]T} e^{\mathbf{A}^C \lambda}\mathbf{B_u}^C\,\nu\,d\lambda + \int_0^{T} e^{\mathbf{A}^C \lambda}\mathbf{B_w}^C\,\mathbf{W}(kT)\,d\lambda \tag{3.17}$$

where $\lambda = T - t$. Note the change in the expression for the first integral. Again assuming small time step, and ignoring higher order terms:

$$\mathbf{A}^D(T) = \mathbf{I} + T\,\mathbf{A}^C \tag{3.18}$$

$$\mathbf{B_u}^D(T) = \eta[k]T\,\mathbf{I}\mathbf{B_u}^C \tag{3.19}$$

$$\mathbf{B_w}^D(T) = T\,\mathbf{I}\mathbf{B_w}^C \tag{3.20}$$

In this case, at each time step, the input power level does not change, except in a step manner. The only variable changing at each period is $\eta[k]$ which is a number between 0 and 1. In the optimization problem, $\eta[k]$ is the variable to be determined. The maximum energy input per period is νT. The discretization of the continuous system results in the same formulation based on the assumption of all variables being constant for each period, only if the time step T is sufficiently small.

Another method for discretization is based on transforming the differential equations directly into difference equations. For sufficiently small Δt, the continuous time equations become:

$$\frac{\mathbf{X}(t_{k+1}) - \mathbf{X}(t_k)}{\Delta t} = \mathbf{A}^C\,\mathbf{X}(t) + \mathbf{B_u}^C\,u(t) + \mathbf{B_w}^C\,\mathbf{W}(t) \tag{3.21}$$

$$y(t) = \mathbf{C}^C\,\mathbf{X}(t) + d_u^C\,u(t) + \mathbf{D_w}^C\,\mathbf{W}(t) \tag{3.22}$$

which can be put in the following form:

$$\mathbf{X}(t_{k+1}) = (\mathbf{I} + \Delta t\mathbf{A}^C)\,\mathbf{X}(t) + \Delta t\mathbf{B_u}^C\,u(t) + \Delta t\mathbf{B_w}^C\,\mathbf{W}(t) \tag{3.23}$$

$$y(t) = \mathbf{C}^C\,\mathbf{X}(t) + d_u^C\,u(t) + \mathbf{D_w}^C\,\mathbf{W}(t) \tag{3.24}$$

Again, it can be seen that for sufficiently small time step, the expressions for the discrete time coefficients are similar to the earlier results. For simplicity, the algorithm of Chapter 4 assumes that at each period the level of energy input can be changed continuously.

3.5 Example: A 3R2C Building

A general model of a building includes all the distinctive thermal masses as the storage or the capacitor elements. It also includes all the significant internal heat transfer interactions among the storage elements, and also between the internal and external ambience. These are represented as energy transfers through resistive elements. Heat and temperature sources are represented, respectively, by equivalent current and voltage sources. In short, heat flow acts like current, and temperature difference acts like a voltage difference.

The variables of interest are the following:

Δt	length of each period (time step)
$T_a[k]$	inside air temperature at period k (C^o)
$T_i[k]$	temperature of storage i at period k
$T_e[k]$	external temperature (or weather variable at period k
$T_{min}[k]$	minimum acceptable inside air temperature at period k
$T_{max}[k]$	maximum acceptable inside air temperature at period k
$U[k]$	electric heat input at period k (kWh/period)
$U_{min}[k]$	minimum possible heat input at period k
$U_{max}[k]$	maximum possible heat input at period k
$Q[k]$	other miscellaneous heat inputs at period k

A building with massive external and internal walls can be modeled as a multi-storage system. The other assumptions and restrictions considered here are:

- During each discrete time period, the state, output, and input variables, and also the spot prices, all remain constant. Therefore, there are no exponential terms in the coefficient matrices.

- Although not considered here, for more accuracy, the inside air can also be modeled as a thermal mass, but its value would be negligible compared to the thermal mass of the walls.

- The deriving force is the ambient temperature or any other weather variable that can be represented by a temperature equivalent. This exogenous input is assumed to be predictable and given for the duration of the time horizon.

- Certain important quality factors such as humidity and the rate of temperature change have been ignored. A more exact representation of the thermal dynamics in a building may require the explicit consideration of humidity in the dynamic equations.

- The miscellaneous heat sources includes solar heating, non-electrical heating, heating due to appliances, impact of humidity changes, and the heating due to the presence of people. If they are not negligible, they are assumed to be given and predictable.

- Electric heat inputs are proportional to the amount of electricity used. Under this assumption, the cost function is linear.

- The heat input is described in terms of the thermal energy output of the source, not the actual amount of electricity consumed. Therefore, the bounds on the heat input must also be described in terms of maximum and minimum amount of thermal energy output of the source. The lower bound is usually zero, and the upper bound is determined by the maximum power rating of the heating source, taking into consideration the efficiency of the source and the duration of the time step.

- The lower and upper bounds on the temperature describe the limits of the comfort range for the residents.

- Costs and delays associated with start ups and shut downs are considered to be negligible.

The thermal parameters unique to each building are:

C_i	thermal capacity of storage i (kWh/C^o)
h_{ae}	coefficient for heat transfer between inside air and outside (kW/C^o)
h_{ei}	coefficient for heat transfer between outside and storage i
h_{ai}	coefficient for heat transfer between inside air and storage i
h_{ij}	coefficient for heat transfer between storage i and storage j
h_f	coefficient for heat transfer through air infiltration
R	heat transfer resistance, equivalent to $1/h$

If the heat transfer between the inside air and outside occurs only through the air infiltration, then h_f and h_{ae} are the same. All heat transfer coefficients are described in the same units. In the related literature it is common to use heat transfer *resistance* instead of heat transfer coefficient. A heat transfer resistance is simply the reciprocal of the respective heat transfer coefficient. It is more convenient when the models are described in equivalent circuit form. Heat transfer coefficients do not depend on the direction of the heat flow. Therefore, h_{ae} is equal to h_{ea}. A typical model of a two storage system is shown in Figure 3.1.

In this model, the two storage units are the internal walls (subscript i) and the external shell (subscript s) of the building. The outside temperature acts as a voltage source, and the solar heat inputs are parallel with the thermal capacitances. For internal thermal mass, solar gain is through the windows. In this example, heat transfer modeling is used to derive the state equations in the state-space form, which is the formulation used in the general algorithm of the Chapter 4 for cost optimization. Another motivation is to demonstrate the special positiveness of thermal dynamic systems.

An issue which will not be discussed here, is the possible dependence of some

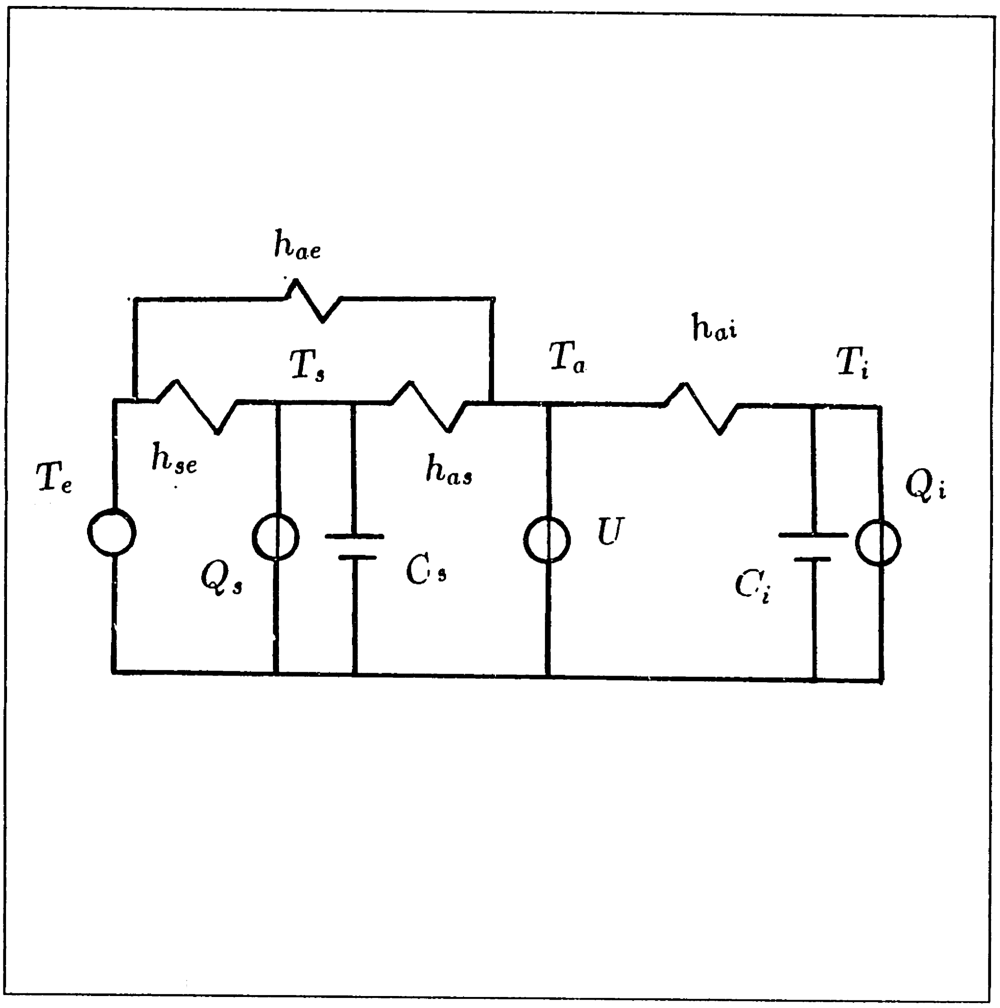

Figure 3.1: 3R2C Heat Transfer Model of a House

storage elements. Bond Graph modeling properly eliminates the dependent storage variables from the system equations, which is sometimes difficult to do by simply relying on intuition.

The energy balance equations are:

$$U[k] = h_{as}\Delta t(T_a[k] - T_s[k]) + h_{ai}\Delta t(T_a[k] - T_i[k]) + h_{ae}\Delta t(T_a[k] - T_e[k]) \quad (3.25)$$

$$Q_s[k] = C_s(T_s[k+1] - T_s[k]) + h_{es}\Delta t(T_s[k] - T_e[k]) + h_{as}\Delta t(T_s[k] - T_a[k]) \quad (3.26)$$

$$Q_i[k] = C_i(T_i[k+1] - T_i[k]) + h_{ai}\Delta t(T_i[k] - T_a[k]) \quad (3.27)$$

The state variables are T_s and T_i, the controlled or the output variable is T_a, the control input is U. All the other variables, i.e. outside temperature (or weather variables) T_e and solar heat gains Q_s and Q_i, are assumed to be deterministic exogenous variables. In addition, all the thermal parameters are assumed to be known.

To put these equations in the state-space form, first the output variable T_a is solved in terms of other variables from the first equation:

$$T_a[k] = \frac{h_{as}}{H}T_s[k] + \frac{h_{ai}}{H}T_i[k] + \frac{h_{ae}}{H}T_e[k] + \frac{1}{H\Delta t}U[k] \quad (3.28)$$

where

$$H = h_{as} + h_{ai} + h_{ae} \quad (3.29)$$

As can be seen, the output (inside air temperature) at any period is a weighted average of the other temperature variables plus a contribution from the electric heat input. This value of T_a is substituted in the next two equations in order to solve for $T_s[k+1]$ and $T_i[k+1]$. The final state-space form is:

$$\begin{bmatrix} T_s[k+1] \\ T_i[k+1] \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} T_s[k] \\ T_i[k] \end{bmatrix} + \begin{bmatrix} b_{u_1} \\ b_{u_2} \end{bmatrix} U[k] + \begin{bmatrix} b_{w_{11}} & b_{w_{12}} & b_{w_{13}} \\ b_{w_{21}} & b_{w_{22}} & b_{w_{23}} \end{bmatrix} \begin{bmatrix} T_e[k] \\ Q_s[k] \\ Q_i[k] \end{bmatrix} \quad (3.30)$$

and the output equation, which we have derived already, is:

$$T_a[k] = [c_1 c_2] \begin{bmatrix} T_s[k] \\ T_i[k] \end{bmatrix} + d_u U_k + [d_{w_1} d_{w_2} d_{w_3}] \begin{bmatrix} T_e[k] \\ Q_s[k] \\ Q_i[k] \end{bmatrix} \quad (3.31)$$

These have the format of the equations used in the general algorithm of Chapter 4. The Coefficients for the matrix **A** above are:

$$a_{11} = 1 - \frac{h_{as}\Delta t}{C_s} - \frac{h_{es}\Delta t}{C_s} + \frac{h_{as}^2\Delta t}{C_s H} \tag{3.32}$$

$$a_{12} = \frac{h_{as}h_{ai}\Delta t}{C_s H} \tag{3.33}$$

$$a_{21} = \frac{h_{ai}h_{as}\Delta t}{C_i H} \tag{3.34}$$

$$a_{22} = 1 - \frac{h_{ai}\Delta t}{C_i} + \frac{h_{ai}^2\Delta t}{C_i H} \tag{3.35}$$

The $\mathbf{B_u}$ vector is:

$$\mathbf{B_u} = \begin{bmatrix} \frac{h_{as}}{C_s H} \\ \frac{h_{ai}}{C_i H} \end{bmatrix} \tag{3.36}$$

The $\mathbf{B_w}$ matrix is:

$$\mathbf{B_w} = \begin{bmatrix} \frac{h_{as} hae \Delta t}{C_s H} & \frac{1}{C_s} & 0 \\ \frac{h_{ai} hae \Delta t}{C_i H} & 0 & \frac{1}{C_i} \end{bmatrix} \tag{3.37}$$

And as shown before, the **C** vector is:

$$\mathbf{C} = \begin{bmatrix} \frac{h_{as}}{H} & \frac{h_{ai}}{H} \end{bmatrix} \tag{3.38}$$

The other coefficients of the output equation are:

$$d_u = \frac{1}{H\Delta t} \tag{3.39}$$

and

$$\mathbf{D_w} = \begin{bmatrix} \frac{h_{ae}}{H} & 0 & 0 \end{bmatrix} \tag{3.40}$$

The patterns observed in these results can be used as equation charts from which equations of more complex thermal systems can be constructed without going through the algebra.

3.6 Size of the Time Step

There are a few observations to be made on the coefficients of the **A** matrix. One is that for the equations to be physically meaningful, the diagonal elements must be positive and less than one. This is due to the fact that the temperature of any thermal storage, starting from a positive value, should decrease to zero in the absence of any thermal source. In other words, if all sources are zero, and the temperature of storage i is positive at period k, then at period $k+1$ the temperature should be a positive fraction of the temperature of the previous period. Of course, with perfect insulation, the temperature should stay constant.

This physical fact can be used to properly select the time step Δt. To transform a continuous time dynamic system into a discrete time system, it is clear that the time step must be at least smaller than the smallest time constant of the system. To meet these conditions the following equalities are to be satisfied:

$$0 < 1 - \frac{h_{as}\Delta t}{C_s} - \frac{h_{es}\Delta t}{C_s} + \frac{h_{as}^2\Delta t}{C_s H} < 1 \tag{3.41}$$

$$0 < 1 - \frac{h_{ai}\Delta t}{C_i} + \frac{h_{ai}^2\Delta t}{C_i H} < 1 \tag{3.42}$$

Recall that $H = h_{as} + h_{ai} + h_{ae}$. Then, solving for Δt in both inequalities, the result is:

$$0 < \Delta t < \frac{1}{\frac{h_{as}}{C_s} + \frac{h_{es}}{C_s} - \frac{h_{as}^2}{C_s H}} \tag{3.43}$$

$$0 < \Delta t < \frac{1}{\frac{h_{ai}}{C_i} - \frac{h_{ai}^2}{C_i H}} \tag{3.44}$$

The right hand terms are always positive since H is larger than its constituting individual heat transfer coefficient. The above inequalities put an upper bound on the value of time step or the length of the time period to be selected. Only within this bound would the discrete equations be a physically meaningful representation of the thermal system under study.

In addition, the time step must be taken to be smaller than any other system time constant. The other system time constants are: C_i/h_{ai} , C_s/h_{as} , and C_s/h_{es}.

3.7 Properties of the A Matrix

In this section, the discrete time 3R2C model is used to demonstrate the special properties of thermal systems. Results can be generalized to more complex thermal systems. The coefficients of the 3R2C state-space equations satisfy the following relationships:

$$a_{11} + a_{12} + b_{w_{11}} = 1 \tag{3.45}$$

$$a_{21} + a_{22} + b_{w_{21}} = 1 \tag{3.46}$$

$$c_{11} + c_{12} + d_{w_{11}} = 1 \tag{3.47}$$

These results can be checked by explicitly writing out the terms and doing the summations. From the above, the following generalizations can be made:

- Elements of the **A** matrix are non-negative.
- Diagonal elements of the **A** matrix are non-zero.
- Sum of each row of the **A** matrix is less than 1.

A similar case can be made for a multi-storage discrete time thermal dynamic system. Multiplication of state variables by appropriate coefficients would result in state variables which have units of energy. Then the first argument is based on the fact that each storage level, starting from a positive value, diminishes to zero in the absence of any replenishing source. The second argument is that in the interactions between different storage units, one storage unit may not get more than what exists in the other storage units. In terms of temperatures, in the absence of other energy sources, because of thermal interactions, the temperature of each storage at any period may not be more than the weighted average of the temperatures of all of the storage units at the preceding period.

In the terminology of dynamic systems, the matrix **A** is called a positive matrix, in the sense that all of its elements are non-negative, with at least one element being strictly positive. Following the exposition of Luenberger [L3], let us define the set of real numbers λ such that

$$\mathbf{AX} \geq \lambda \mathbf{X} \tag{3.48}$$

for some $\mathbf{X} \geq 0$ where $\mathbf{X}$ is some strictly nonnegative vector. It can be seen that $\lambda = 0$ satisfies the inequality. Also, λ may not be arbitrarily large. This can be seen from the fact that the size of the elements of $\mathbf{A}$ put a limit on how large the elements of the resulting vector at the left-hand side can be. Define λ_0 as

$$\lambda_0 = \max \{ \lambda : \mathbf{AX} \geq \lambda \mathbf{X}, \quad \text{some } \mathbf{X} \geq 0 \}$$

The following theorem shows that in the case where $\mathbf{A}$ is a nonnegative matrix, λ_0 defined above is the dominant eigenvalue for $\mathbf{A}$.

Extension of the Frobenius-Perron Theorem [L3]. If $\mathbf{A} \geq 0$, then there exists a positive eigenvalue $\lambda_0 \geq 0$ and a positive eigenvector $\mathbf{X}_0 \geq 0$ such that:

a) $\mathbf{AX}_0 = \lambda_0 \mathbf{X}_0$;

b) If $\lambda \neq \lambda_0$ is any other eigenvalue of $\mathbf{A}$ then $|\lambda| < \lambda_0$.

The basic result of the Ferobenius-Perron theorem is that the matrix $\mathbf{A}$ has a dominant positive eigenvalue λ_0, referred to as the *Frobenius-Perron eigenvalue* of $\mathbf{A}$. Hence, the behavior of the system is dominated by this eigenvalue, and its associated eigenvector.

Furthermore, a set of bounds can be derived for the value of λ_0. Let $\mathbf{X}_0$ be the corresponding *normalized* positive eigenvector with components $x_1, x_2, \cdots, x_n$. We have:

$$\mathbf{AX}_0 = \lambda_0 \mathbf{X}_0 \tag{3.49}$$

Writing this matrix equation in detail as n equations, and summing them results in:

$$\lambda_0 = \Delta_1 x_1 + \Delta_2 x_2 + \cdots + \Delta_n x_n \tag{3.50}$$

where Δ_i is the sum of the elements in the ith column of $\mathbf{A}$. Since vector $\mathbf{X}_0$ is normalized, then λ_0 is a weighted average of the column sums of $\mathbf{A}$ and therefore, it must lie between the two extreme values of the column sums. Since the transpose of $\mathbf{A}$ has the same dominant eigenvalue, it can be shown that λ_0 is also a weighted average of the row sums of matrix $\mathbf{A}$. The results can be summarized as:

- λ_0 is bounded by the minimum and the maximum column sums of the matrix $\mathbf{A}$.
- λ_0 is bounded by the minimum and the maximum row sums of the matrix $\mathbf{A}$.

We have already shown that the row sums of the matrix $\mathbf{A}$ are more than zero and less than one. This yields the following bounds:

$$0 < \lambda_0 < 1 \tag{3.51}$$

From the above results we can make the following related statements:

- Thermal systems are *asymptotically stable*, since the matrix $\mathbf{A}$ has all of its eigenvalues strictly within the unit circle of the complex plane.
- The homogeneous thermal systems have their equilibrium point at zero.

The second result is based on the fact that the dominant eigenvalue is less than one. To show this, suppose $\bar{\mathbf{X}}$ is the equilibrium point of the homogeneous system satisfying

$$\bar{\mathbf{X}} = \mathbf{A}\,\bar{\mathbf{X}} \tag{3.52}$$

Since all eigenvalues are less than one, then the matrix $\mathbf{I} - \mathbf{A}$ is nonsingular, and the equilibrium point $\bar{\mathbf{X}}$ is a zero vector.

The above results indicate that a higher-order (multi-storage) thermal system will have a dynamic behavior very similar to a simple first order system. Simplicity of the system response can best be explained in the form of its associated *impulse response* vector. The simplicity of the scheduling algorithm of Chapter 4 is direct result of the simple form of the impulse response vector. Indeed, the algorithm itself is mostly based on the utilization of linearity of the systems and the their associated impulse response vectors.

3.8 Properties of the Impulse Response Vector

This section discusses the special properties of the impulse response vectors for thermal systems, which are based on the results of the preceding section. As discussed more fully in the next chapter, the general single-input single-output time invariant discrete time linear system has the following form:

$$\mathbf{X}[k+1] = \mathbf{A}\ \mathbf{X}[k] + \mathbf{B_u}\ u[k] + \mathbf{B_w}\ \mathbf{W}[k] \tag{3.53}$$

$$y[k] = \mathbf{C}\ \mathbf{X}[k] + d_u\ u[k] + \mathbf{D_w}\ \mathbf{W}[k] \tag{3.54}$$

where u is the controllable input, and $\mathbf{W}$ is the vector of exogenous inputs. The impulse response vector is:

$$\mathbf{I} = \begin{bmatrix} d_u \\ \mathbf{CB_u} \\ \mathbf{CAB_u} \\ \vdots \\ \mathbf{CA} \cdots \mathbf{AB_u} \end{bmatrix} \tag{3.55}$$

Subsequent elements of the impulse response vector for $k \geq 3$ are of the form $i_k = \mathbf{CA}^{k-2}\mathbf{B_u}$. Results of the previous section ensure that for $k > 2$:

- As k increases, i_k decreases in a monotone manner.
- Ratio of i_{k+1} to i_k increases in a monotone manner to the value of the dominant eigenvalue.

To understand these results, recall that the sum of the elements of $\mathbf{C}$ is less than one. Thus $\mathbf{CB_u}$ is less than the sum of the elements $\mathbf{B_u}$. Inclusion of the matrix

$\mathbf{A}$ in subsequent elements of the impulse vector results in a smaller number, since the sums of each rows of $\mathbf{A}$ are also less than one.

In addition, the ratio of i_2 to i_1 is always less than 1. This can be observed by explicitly writing out the ratio for the 3R2C model:

$$\frac{i_2}{i_1} = \frac{\mathbf{CB_u}}{d_u} = \Delta t \left(\frac{h_{as}^2}{C_s H} + \frac{h_{ai}^2}{C_i H} \right) \tag{3.56}$$

and also recalling that the following relations always hold:

$$\Delta t < \frac{C_s}{h_{as}} \tag{3.57}$$

$$\Delta t < \frac{C_i}{h_{ai}} \tag{3.58}$$

$$H = h_{as} + h_{ai} + h_{ae} \tag{3.59}$$

Now, let's take the largest of the time constant terms, say h_{as}/C_s, then

$$\frac{i_2}{i_1} < \Delta t \left(\frac{h_{as}}{C_s} \frac{h_{as} + h_{ai}}{H} \right) < \Delta t \frac{h_{as}}{C_s} < 1 \tag{3.60}$$

Furthermore, by taking Δt small enough, it can be ensured that:

$$\frac{i_2}{i_1} < \frac{i_3}{i_2} < \cdots < \frac{i_{k+1}}{i_k} \tag{3.61}$$

In fact, if Δt is decreased, the diagonal elements move closer to unity, and simultaneously, the off-diagonal elements move closer to zero. Beyond a certain value Δt, all the submatrices of $\mathbf{A}$ will have a determinant greater than zero, which ensures that the $\mathbf{A}$ matrix is a *positive definite* matrix, and as a result all its eigenvalues become positive. In fact as Δt is made nearly equal to zero, the $\mathbf{A}$ matrix becomes nearly diagonal with elements approaching one, and the diagonal elements correspond to the eigenvalues of the matrix.

Thus, for a homogeneous system, the result of a impulse input is a sudden jump in the output of the first period, and asymptotical decay afterwards to an equilibrium value of zero.

To most important result, the one which is going to be the basis of the algorithm introduced in Chapter 4, is:

- If starting from a given schedule of inputs (initial solution), the input $u[i]$ at some period i is increased, then the output of all the future periods will increase.

- Then, the input at some later period j, where $j > i$, can be decreased so that the output at period j is returned to its initial level without having any of outputs at later periods becoming less than their initial values.

The above statements do not hold for systems with impulse response vectors which exhibit increasing or oscillatory behavior. They are, also, not true if the ratios of the subsequent elements of the impulse response vector do not *monotonically* and *asymptotically* increase to a number less than one. For the systems satisfying the conditions under which the algorithm of Chapter 4 is applicable, this limiting number is the Ferobenius-Perron eigenvalue, which it is less than 1 for the thermal systems.

The following chapter uses the above result and introduces a general algorithm for the minimum cost scheduling inputs for deterministic discrete-time time-invariant linear dynamic thermal systems. The following section of this chapter describes the control scheme under which the algorithm will be applied.

3.9 How to Handle Uncertainty

In practice, the scheduling will be implemented under uncertainties which exist in the behavior models and the forecasts of the exogenous inputs. Four of the possible methods are listed below:

CERTAINTY EQUIVALENT CONTROL: Use the expected values of the uncertain exogenous variables (weather, hot water demand, etc.), and solve the deterministic optimization problem for the next N periods, and apply the resulting control inputs accordingly for the whole time horizon.

FEED FORWARD CONTROL: Use the best possible prediction of exogenous variables, and then solve a deterministic optimization Problem for the next N time periods. But in contrast with the Certainty Equivalent Control, implement the resulting control only for the very next time period. At the next period, go back to the first step again. A major advantage of the first two methods is that it is not necessary to model the uncertainties explicitly. However, the Feed Forward Control takes into account the new information about future values of the exogenous variables as they become available.

STOCHASTIC CONTROL: Model uncertainties explicitly and apply stochastic dynamic programing. Depending on the probabilistic nature of the exogenous inputs and the structure of the controlled system, this method may be computationally very time consuming.

MIN-MAX CONTROL: Use of *unknown but bounded* models for the uncertainty wherein only upper and lower limits, i.e bounds, on the uncertainty are assumed [G3]. An explicit probabilistic/stochastic structure is not required. One possible "min-max" criteria is to find the control that minimizes the maximum cost subject to meeting the various constraints. Take the extreme values of uncertain variables and solve the deterministic optimization problem.

Based on the preceding discussions, this thesis recommends the implementation of the **Feed Forward Control.** The deterministic optimization problem is to be solved by the algorithm introduced in Chapter 4.

In Chapter 7, stochastic control is applied to a simple first order system in order to evaluate its practicality, and lay the foundation for future work.

3.10 Parameter Estimation

As discussed in Chapter 2, an important function of a complete energy management system is to provide the Tactical Control with the behavior models of the system to be controlled. For space heating and cooling, the requirements are twofold:

- To provide a representative model of the topology of the system configuration in terms of its constitutive elements.
- To provide the values of the parameters of the system

The first requirement falls under the subject of system modeling and *system identification* [L2]. It must be actually carried out by a person familiar with dynamic system modeling. The task is to figure out the shape of the equivalent circuit diagram for the system. This is equivalent to determining the order of the system, or the number of significant storage elements, and their location in the thermal network. The final product is similar to the network shown in Figure 3.1.

The second requirement is to provide the values of the thermal properties for the elements which constitute the network. There are two very different approaches:

1. **Calculations from Heat Transfer Models:** This method is based on calculating the values of thermal mass and heat transfer coefficients from the physical properties of the materials used in the construction and the geometric configuration of the building.
2. **Parameter Estimation:** This method assumes a relatively simple lumped parameter thermal dynamic structure for the building, and includes only the most basic heat storage components. The thermal coefficients are then estimated using statistical and regression techniques by fitting the observed performance data on temperatures and heat inputs to the equations.

The first method requires many man hours of work by a person with a professional engineering background. This is a rather involved procedure, and requires a knowledge of the layout of the building, and the materials used in its construction, and information on thermal properties of individual walls, windows, insulation, etc. Usually, the result is a model of a very high order. Heat transfer parameters for different materials can be found from ASHRAE (American Society of Heating, Refrigeration, and Air Conditioning Engineers) handbooks. There are also software packages available that can be used in evaluation of heat transfer properties of buildings of different shapes and materials.

The second methodology is a more recent development. System identification and parameter estimation techniques have been applied to the determination of equivalent thermal parameters of buildings in References [S8], [W3], and [R1]. The case study of this chapter relies on published data from these sources. This method is well suited for implementation on energy management systems.

Chapter 4

OPTIMAL SCHEDULING ALGORITHM FOR THERMAL SYSTEMS

4.1 Introduction

This chapter introduces an algorithm which finds the optimal scheduling of electricity usage for thermal systems. The controlled system is a *single-input single-output multi-storage deterministic discrete-time linear time-invariant thermal system* subject to *linear costs and bounds on variables.* The linear cost function is the total cost for the time horizon, which is the sum of the cost of the electricity usage at each period for all the periods in the time horizon. Recall that the products of the electricity usage (heat or hot water) are assumed to be proportional to the electricity used. The bounds on the input are the physical limits on the size of electricity usage at each period. The bounds on the output are either due to the physical constraints (water heater tank size) or due to the comfort limits (inside air temperature) set by the electricity consumers.

This chapter begins with a mathematical statement of the optimization problem. Next, a descriptive presentation of how the algorithm works is provided, followed by a summarized list of the steps of the algorithm. The main section of this chapter presents the mathematical exposition of the steps of the algorithm using the *block representation* of the discrete time linear systems. The final sections address the issues of convergence, optimality, and the relation of the algorithm to the simplex method. A more formal and step by step description of the algorithm, and its programming code is presented in the Appendix A.

The simple structure of the problems considered in the Chapters 5 and 6 result in special and simpler cases of the algorithm. The simplifications and actual codes for the case studies are presented in the Appendices B and C.

4.2 Statement of the Optimization Problem

This section presents a mathematical statement of the optimal scheduling problem. The variables of interest are defined as:

N	total number of periods in the time horizon
$x_i[k]$	ith state variable at period k, element of column vector $\mathbf{X}[k]$
$y[k]$	output variable at period k
$u[k]$	control input at period k
$w_i[k]$	ith exogenous input at period k, element of column vector $\mathbf{W}[k]$
$p[k]$	unit cost of input at period k
$y[k]_{min}$	minimum acceptable value of $y[k]$
$y[k]_{max}$	maximum acceptable value of $y[k]$
$u[k]_{min}$	minimum allowable value of $u[k]$
$u[k]_{max}$	maximum allowable value of $u[k]$

In this chapter,

- Bracket notation $[k]$ is reserved to denote time.

- Upper case letters denote vectors and matrices, whose elements are represented by lower case letters.

- Subscripts indicate the position within the vector or a matrix.

The general single-input single-output time invariant discrete- time linear system has the following form:

$$\mathbf{X}[k+1] = \mathbf{A}\ \mathbf{X}[k] + \mathbf{B_u}\ u[k] + \mathbf{B_w}\ \mathbf{W}[k] \tag{4.1}$$

$$y[k] = \mathbf{C}\ \mathbf{X}[k] + d_u\ u[k] + \mathbf{D_w}\ \mathbf{W}[k] \tag{4.2}$$

where, following the notations given above, the state vector is:

$$\mathbf{X}[k] = \begin{bmatrix} x_1[k] \\ x_2[k] \\ \cdot \\ \cdot \\ \cdot \\ x_n[k] \end{bmatrix}$$

Similarly, $\mathbf{W}$ is a $q \times 1$ column vector. The time invariant matrix $\mathbf{A}$ is :

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & \cdots & \cdots & \cdots \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix}$$

$\mathbf{B_w}$ has a similar form with dimensions $n \times q$.

Other vector coefficients are :

$$\mathbf{B_u} = \begin{bmatrix} b_{u_1} \\ b_{u_2} \\ \vdots \\ b_{u_n} \end{bmatrix}$$

and

$$\mathbf{C} = [c_1\ c_2\ \cdots\ c_n]$$

For multi-input multi-output systems, $\mathbf{C}$ and $\mathbf{B_u}$ are matrices instead of vectors. As can be seen, input has been parted into two groups. The control inputs $u[k]$ are specified by the controller. The exogenous inputs $\mathbf{W}[k]$, on the other hand, are specified by the environment. In most cases the state variables are the storage terms, the exogenous variables are the demand terms, and the control variables are the input terms. For heating, the demand terms are negative, where they act to deplete the storage, and the input variables are specified so that the depletion of the storage is avoided. The reverse is true for space cooling.

The output variable y can be one of the state variables. This becomes possible when a state variable is also one of the measurable variables to be controlled.

Let Z denote the objective function. The mathematical problem considered here can be stated as finding solutions $u^*[k]$ for each k which minimize the following objective function for a given time horizon of N periods:

$$\min_{u[k]} Z = \min_{u[k]} \sum_{k=1}^{N} p[k] \cdot u[k] \tag{4.3}$$

subject to the following constraints for $k = 1$ to N :

$$\mathbf{X}[k+1] = \mathbf{A}\ \mathbf{X}[k] + \mathbf{B_u}\ u[k] + \mathbf{B_w}\ \mathbf{W}[k] \tag{4.4}$$

$$y[k] = \mathbf{C}\ \mathbf{X}[k] + d_u\ u[k] + \mathbf{D_w}\ \mathbf{W}[k] \tag{4.5}$$

$$y[k]_{min} \leq y[k] \leq y[k]_{max} \tag{4.6}$$

$$u[k]_{min} \leq u[k] \leq u[k]_{max} \tag{4.7}$$

Values of all the coefficients, and also all the initial values $\mathbf{X}[1]$, and *all* the elements of $\mathbf{W}[k]$ and $p[k]$ are assumed given for all k. As the notation implies, the bounds on the inputs and the outputs may also vary with time, which will be discussed later.

The above formulation is indeed in the form of a linear programming problem, and therefore, any of existing methodologies, such as simplex, can be utilized to find the optimal solutions. However, the intention here is to take advantage of the special properties of the asymptotically stable positive systems in order to develop a more efficient algorithm.

Proper usage of variables and units can result in the state equations representing changes in the storage of *energy* terms in time. However other convenient physical entities can be used as the state and control variables, where the resulting coefficients in the state equations would reflect physically meaningful properties.

4.3 Descriptive Presentation

The basic feature of the problem considered here is that input at some initial time can be stored for use at later times subject to some storage losses and interactions with other storage units. In mathematical terms, this means that an increase in the input at any period will increase the state and the output variables of future periods. The magnitude of the increases depend on the *impulse response* of the system for time-invariant systems.

In the specific setting of electricity spot pricing, the control variable $u[k]$, i.e. consumption of electricity, has a different unit cost $p[k]$ for each period of usage k. Given an initial feasible pattern of electricity usage, it is possible to re-arrange electricity consumption in a manner which lowers the total cost for the time horizon without violating any of the output (service) and the input (control) constraints. One way of doing this is first to find some appropriate initial feasible solution, and then to search for a period with a low electricity price and to increase usage at that period. Doing so results in a net increase in the storage, some of which is lost with the passing of time. Whatever remains can be used later, particularly when the price is higher, to decrease the electricity usage of that period from its initial level.

One particular initial feasible solution, which can be easily found without consideration of the given price pattern, is the one based on the least amount of the total electricity consumed during the time horizon. In fact, the control inputs specified by a flat-price operation are the initial feasible solutions, provided that the controlled variables are kept at their lowest permissible level.

After finding the initial feasible solution, the algorithm checks to see if there are any possible savings achieved by increasing the input at any of the periods and substituting for some of the future inputs. The period with the highest potential savings is designated as the *charge period*, i.e. the period at which the input will be increased in order to decrease some future inputs.

The algorithm considers every period as a potential charge period, and it computes the maximum amount by which *every one* of future control inputs can be reduced in *chronological* order if the control input of the charge period is increased by one unit.

The effect of an increase in the input of any period is a subsequent increase in the level of future state and output variables. Therefore, although control inputs at future periods will be at their initial values, future state and output variables of each period are going to be above their initial values. The term *initial* refers to the values of the previous solutions for the same periods. In this thesis, the term *reference* is used interchangeably with *initial*.

As shown in Chapter 3, the relative increase in the output of each period will decay to zero as time goes by, i.e. as k increases. Therefore, if the input of any selected period, i.e the *charge period*, is increased, then all the future outputs $y[k]$ will be raised above their initial level, and it becomes possible to reduce the control input of future periods from their initial level as long as the output and the input variable constraints are not violated. Calculation of these reductions and also of the resulting changes in the state and output variables are made easier by the linearity of the system.

These reductions are accomplished by moving in time from the charge period, to the last period, and one be one reducing the inputs to their lowest possible values. Once a feasible direction vector is found for a charge period, the procedure of moving to a new solution point require simple scaler-vector multiplication or vector additions.

The elements of a feasible direction vector associated with a charge period are the amount of the decrease in the input of every future period when the input of the charge period is increased by unity. This vector can be thought of as a feasible direction vector in the space defined by the decision variable $u[k]$, along which the *block* vector of inputs can be changed at each iteration.

Once at a solution point, the algorithm does not consider all possible directions. It restricts itself to the consideration of the particular class of directions computed by the above procedure. In fact, there are other feasible directions along which inputs can be changed simultaneously. However, orderly movement in the directions determined by the algorithm ensures convergence to the optimal solution.

Using these possible feasible directions, the algorithm checks to see how economical is the consideration of each period as a potential charge period. In a sense, the algorithm computes the marginal savings associated with each period as if it were a charge candidate. At each iteration it searches for the period with the highest marginal saving, and then, it increases the input of that period and simultaneously reduces the control input of the future periods (moving along the feasible direction) as long as it is economically and physically feasible.

Economic feasibility means that the cost of the input increase at the charge period must be less than the savings accrued from the reduction of future inputs. The physical feasibility requires that inputs and outputs be kept within acceptable levels for all periods.

After carrying out the rescheduling and finding the new solution,the algorithm searches for the next set of possible feasible direction vectors. The new set is different from the previous one, because each iteration is completed when a new constraint becomes active. As a result, the algorithm must also determine the new set of marginal savings.

After updating the set of feasible direction vectors and the marginal savings vector, the iteration loop is started again by selecting a new period with the highest new marginal saving as the new charge period, and the process is repeated as before.

Iterations are continued until no positive marginal savings can be found. At each iteration, the algorithm preserves feasibility, and proceeds with rescheduling of control inputs as long as it is economically viable. With each intermediate solution, the total cost of inputs within the time horizon is decreased.

Most of the calculations involve vectors because the linearity of the system enables us to take advantage of the *impulse response* representation of the system output. For discrete time systems these can be defined as the elements of row or column vectors whose rank would not exceed the length of the time horizon. Impulse response representation, which is one manifestation of the superposition property of the linear systems, results in simple vectorized computational routines for most of the algorithm.

4.4 Summary of the Algorithm

What follows is a descriptive summary of the steps of the algorithm:

- Find the reference (initial) feasible solution : Assume a flat price schedule and find initial control inputs $u^{init}[k]$ for all k which result in the minimum total control inputs. This can be done by solving the state equations one period at a time starting at period $k = 1$ while keeping all output variables are kept at their minimum.

- Find the feasible direction vector and marginal savings s_k associated with a unit increase in the input at period k in the following way: For each period increase the input by one unit and reduce all the future inputs one be one by as much as possible without violating the constraints, and compute the resulting savings and costs. A feasible direction vector for each period is a vector whose elements are the amount of the feasible reductions in the inputs of the following periods for a unit increase in the input of the current period.
- Order the periods by their marginal savings.
- Select the period with the highest marginal saving as the charge candidate. If no period with positive marginal savings can be found, terminate the process.
- Otherwise, increase the input at the charge candidate. This will increase the state and output variables of the future periods above their previous levels. Then start with the first period after the charge period, and one by one reduce inputs at future periods until all the output variables are back to their lower bounds. The combined movement in the feasible space is represented by the feasible direction vector. Do this for every period except for those which have been selected before as charge candidates unless the output variable at that period is at its maximum possible value. Continue increasing the control input at the charge candidate and reducing future inputs in the above manner until one of the following happens:
  - The control input at the charge period reaches the maximum level.
  - The output variable at the charge period reaches the maximum level.
  - A control input at any future period reaches the minimum level.
  - An output variable at some of the future periods reaches the maximum level.

The last case may happen for those future periods at which the control input is already zero and no further reduction in the control input is possible.

- Find the new set of solutions.
- Recalculate marginal savings for all the periods. Note that the control inputs of previous charging periods should not be reduced unless the output variables at those periods reach the maximum level possible. This results in a new set of marginal savings values. At each iteration, values of the total cost and the maximum marginal savings are reduced.
- Select a new charging period based on the highest new marginal saving among all of the periods. A previous charge period can be selected again.
- Continue with the iterative scheduling until no period with positive marginal saving can be found. At this point, the optimal schedule of the control input is found.

The number of steps, calculations, procedures, and iterations would be much larger if instead of using impulse response vectors, all the computations were based on solving the state equations recursively.

4.5 Mathematical Presentation

4.5.1 Introduction

The heuristic algorithm presented here is simply a formalization of the preceding summary. Rescheduling of inputs, from a mathematical point of view, is equivalent to moving in an N dimensional space from an initial solution point to a new solution point in the direction of a *feasible direction vector*. We can think of two distinct multi-dimensional space. One belonging to the input variables and the other to the output variables.

The algorithm starts by generating a particular feasible initial solution, and then by using an efficient methodology, it iteratively reschedules electricity usage so as to reduce total cost at each iteration.

The algorithm moves towards the optimal solution by a series of moves in the feasible space. Let $\mathcal{U}$ be a vector whose elements are $u[1]$, $u[2], \cdots, u[N]$. In its most general form our constrained optimization problem is of the form

$$\text{minimize } f(\mathcal{U})$$
$$\text{subject to } \mathcal{U} \in \Omega$$

The set of constraints together describe the feasible set Ω which is a subset of E^N. Starting from a set of initial feasible solutions $\mathcal{U}^{init}$ on a constraint surface, the algorithm moves in the feasible space to a new solution set in a feasible direction until it hits a new constraint surface. In general, there is an infinite number of feasible direction vectors at each point in the feasible set. The algorithm finds a suitable feasible direction vector $\mathbf{F_u}$ for each solution point. Therefore, given a solution point $\mathcal{U}^{old}$ and a feasible direction vector $\mathbf{F_u}$ the new solution point is found by

$$\mathcal{U}^{new} = \mathcal{U}^{old} + \alpha \mathbf{F_u} \tag{4.8}$$

where α is the proportion of the length of the direction vector $\mathbf{F_u}$ which must be travelled before a new constraint surface is met.

At each iteration, a new feasible direction is chosen and the solution is moved closer to the optimal. During each iteration, feasibility is maintained. In addition, the total cost is monotonically decreased at each iteration, i.e.

$$f(\mathcal{U}^{new}) < f(\mathcal{U}^{old}) \tag{4.9}$$

Reduction of the total cost at each iteration ensures convergence to the optimal solution.

4.5.2 Effect of input increase on the output: state transition matrix

One way of studying the effects of an input increase on the future outputs is to look at the transitions of the states. It is of interest to know if an input can be increased at some period, in order to replace the input of some future period, while the output levels are within the accepted bounds. The expressions for the state and the output transition depict the effect of an input increase at some period on the states and the outputs of the future periods. By recursively substituting for the state variables, states and outputs of all periods can be expressed in terms of the initial states and the streams of all the past inputs. The result is the expression for the state transition, where the time index starts from 1:

$$\begin{aligned} \mathbf{X}[k+1] = \mathbf{\Phi}[k+1,1]\mathbf{X}[1] &+ \sum_{l=1}^{k} \mathbf{\Phi}[k+1,l+1]\mathbf{B_u}u[l] \\ &+ \sum_{l=1}^{k} \mathbf{\Phi}[k+1,l+1]\mathbf{B_w}\mathbf{W}[l] \end{aligned} \tag{4.10}$$

and the output transition:

$$\begin{aligned} y[k] = \mathbf{C\Phi}[k+1,1]\mathbf{X}[1] &+ \sum_{l=1}^{k-1} \mathbf{C\Phi}[k,l+1]\mathbf{B_u}u[l] + d_u u[k] \\ &+ \sum_{l=1}^{k-1} \mathbf{C\Phi}[k,l+1]\mathbf{B_w}\mathbf{W}[l] + \mathbf{D_w}\mathbf{W}[k] \end{aligned} \tag{4.11}$$

where the *state transition matrix* is defined as:

$$\mathbf{\Phi}[k,l] = \mathbf{A}[k-1]\mathbf{A}[k-2]\cdots\mathbf{A}[l], \ k > l \tag{4.12}$$

$$\mathbf{\Phi}[l,l] = \mathbf{I} \tag{4.13}$$

For a time-invariant system, the state transition matrix is:

$$\mathbf{\Phi}[k,l] = \mathbf{A}^{k-l} \tag{4.14}$$

Therefore, if $u[i]$ is increased by δu_i, output at i is increased by:

$$y^{new}[i] - y^{old}[i] = d_u \delta u_i \tag{4.15}$$

and the output of the period immediately after is increased by:

$$y^{new}[i+1] - y^{old}[i+1] = \mathbf{CIB_u}\delta u_i \tag{4.16}$$

where $\mathbf{I}$ is the identity matrix. Similarly, the output at period j where $j > i$ is increased by:

$$y^{new}[j] - y^{old}[j] = \mathbf{CA}^{j-i-1}\mathbf{B_u}\delta u_i \tag{4.17}$$

Based on the results of Chapter 3 on the properties of the coefficient matrices, the right-hand side of the above equation is positive. Also, as the distance of $j - i$ is increased in time, the value of the output change is decreased asymptotically to zero. Because the output change is positive, we can decrease the input at time j by:

$$\delta u_j = \frac{y^{new}[j] - y^{old}[j]}{d_u} \tag{4.18}$$

which forces the output to return to its original value $y^{old}[j]$. Block representation, presented next, makes these arguments easier to understand.

4.5.3 Reformulation using block representation

A better understanding of these concepts can be gained by studying the *block representation* of the system. It is also the formulation on which the algorithm is based. Block representation is obtained by recursively substituting for the state variables in the state-space equations. As a result, state and output variables at any period are described in terms of the initial state variables and the stream of all the previous input variables starting with the initial period. Block description of the system is:

$$\mathbf{X}[k+1] = \mathcal{A}_k \mathbf{X}[1] + \mathcal{B}_{u_k} \mathcal{U}_k + \mathcal{B}_{w_k} \mathcal{W}_k \tag{4.19}$$

$$\mathcal{Y}_k = \mathcal{C}_k \mathbf{X}[1] + \mathcal{D}_{u_k} \mathcal{U}_k + \mathcal{D}_{w_k} \mathcal{W}_k \tag{4.20}$$

where

$$\mathcal{A}_k = \mathbf{A}[k]\mathbf{A}[k-1]\cdots\mathbf{A}[1] = \mathbf{A}^k \tag{4.21}$$

$$\mathcal{B}_{u_k} = \left(\mathbf{A}^{k-1}\mathbf{B}_{\mathrm{u}} \mid \mathbf{A}^{k-2}\mathbf{B}_{\mathrm{u}} \mid \cdots \mid \mathbf{B}_{\mathrm{u}}\right)$$

$$\mathcal{B}_{w_k} = \left(\mathbf{A}^{k-1}\mathbf{B}_{\mathrm{w}} \mid \mathbf{A}^{k-2}\mathbf{B}_{\mathrm{w}} \mid \cdots \mid \mathbf{B}_{\mathrm{w}}\right)$$

(4.22)

and the input blocks are represented as:

$$\mathcal{U}_k = \begin{bmatrix} u[1] \\ u[2] \\ \vdots \\ u[k] \end{bmatrix} \tag{4.23}$$

and

$$\mathcal{W}_k = \begin{bmatrix} \mathbf{W}[1] \\ \mathbf{W}[2] \\ \vdots \\ \mathbf{W}[k] \end{bmatrix} \tag{4.24}$$

where elements of $\mathcal{U}_k$, represented by lower case terms, indicate the fact that this is a single-input system. Similarly, the output vector is:

$$\mathcal{Y}_k = \begin{bmatrix} y[1] \\ y[2] \\ \vdots \\ y[k] \end{bmatrix} \tag{4.25}$$

Other coefficient matrices are:

$$\mathcal{C}_k = \begin{pmatrix} \mathbf{C} \\ \mathbf{CA} \\ \vdots \\ \mathbf{CA}^{k-1} \end{pmatrix} \tag{4.26}$$

and

$$\mathcal{D}_u = \begin{pmatrix} d_u & 0 & 0 & \cdots & 0 \\ \mathbf{CB_u} & d_u & 0 & \cdots & 0 \\ \mathbf{CAB_u} & \mathbf{CB_u} & d_u & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{CA}^{k-2}\mathbf{B_u} & \mathbf{CA}^{k-3}\mathbf{B_u} & \cdots & \cdots & d_u \end{pmatrix} \tag{4.27}$$

where $\mathcal{D}_w$ is similarly defined.

In this manner, the effect of one unit increase of the input on the outputs of the future periods can be easily computed. For example the first column of $\mathcal{D}_u$ is the response of all the future outputs to a unit increase in the input of the first period.

In the problem considered here, the total number of time steps is N, and the subscripts in the block vectors and matrices indicating the size, can be dropped without further loss of generality. Thus, for example, the vector $\mathcal{U}$ is a $N \times 1$ column vector.

4.5.4 Mathematical description of the algorithm

Problem statement in block representation: The optimization problem can be restated using the block notation:

$$\min_{u[k]} Z = \min_{u[k]} \sum_{k=1}^{N} p[k] \cdot u[k] \tag{4.28}$$

subject to the following constraints for $k = 1$ to N :

$$\mathcal{D}_u\, \mathcal{U} \le \mathcal{Y}_{max} - \mathcal{C}\ \mathbf{X}[1] - \mathcal{D}_w\, \mathcal{W} \tag{4.29}$$

$$\mathcal{D}_u \mathcal{U} \ge \mathcal{Y}_{min} - \mathcal{C}\ \mathbf{X}[1] - \mathcal{D}_w\, \mathcal{W} \tag{4.30}$$

$$\mathcal{U}_{min} \le \mathcal{U} \le \mathcal{U}_{max} \tag{4.31}$$

The constraints are written in the form of vector inequalities, and they consist of $4N$ scaler inequalities. Care must be taken to avoid confusing a block vector such as $\mathcal{W}$ with its associated variable vector $\mathbf{W}[k]$.

In the latest formulation, the only variables to be determined are the N elements of the vector $\mathcal{U}$. All the other variables and coefficients are given. Again, the optimization problem is set up in the familiar linear programming format.

A particular feasible initial solution: A feasible initial solution is found by ignoring the changes in prices, i.e. assuming a flat price pattern, and finding the set of control inputs which minimize the total sum of control inputs for the time horizon. This is done be keeping the output variable at the lowest level possible, i.e. at $y[k]_{min}$, for each period k, as long as the control inputs are not forced to violate their bounds . Solutions found this way are actually the minimum value of the sum of the inputs $u[k]$ necessary to keep the output in the feasible region.

The initial feasible solution set is found by solving for $\mathcal{U}$ which forces the output to $\mathcal{Y}_{min}$. This is done by solving the difference equations recursively for u_k for $k = 1, 2, ...N$, given the initial values and all the exogenous inputs $\mathbf{W}[k]$, and by setting all $y[k]$ equal to y_{min} for all k. This may be impossible for certain systems. Finding such a feasible solution depends on the *reachability* of the system. If the system is reachable, $\mathcal{U}^{init}$ is found by solving the state equations recursively for each $u[k]$ given that each $y[k]$ is substituted by y_{min}.

Given this initial solution set, no input can be decreased unilaterally without violating at least one of the output constraints. The only way to decrease an input at any period is to simultaneously increase the input of a previous period. Thus, the feasible movements from the initial solution points are restricted to changes in inputs where the first input in time subject to change must be increased in magnitude, i.e. *increase must occur prior to the decrease.* This means that any change from the initial solution is only possible if the input is increased at a period with subsequent decreases of inputs at later periods, rather than the reverse.

Impulse response vectors: To find a suitable feasible direction vector, let's look at the effect of a change in the input of a period on the outputs of all of the periods. Block representation is in fact equivalent to the state transition method.

Denoting the ith column of $\mathcal{D}_u$ by $\mathbf{I}^i$, for the first column we have:

$$\mathbf{I}^1 = \begin{bmatrix} d_u \\ \mathbf{CB}_u \\ \mathbf{CAB}_u \\ \vdots \\ \mathbf{CA}^{N-2}\mathbf{B}_u \end{bmatrix} \tag{4.32}$$

and similarly, for the ith column is :

$$\mathbf{I}^i = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ d_u \\ \mathbf{CB_u} \\ \mathbf{CAB_u} \\ \vdots \\ \mathbf{CA}^{N-i-1}\mathbf{B_u} \end{bmatrix} \tag{4.33}$$

This last vector is the change in the output response of the system for all periods to a unit increase in the input at the ith period.

Taking the previous set of N outputs to be $\mathcal{Y}^{old}$, then if at period i the input $u[i]$ is increased by δu_i, the new output block becomes:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \mathbf{I}^i \delta u_i \tag{4.34}$$

Or alternatively, for each period we can write:

$$\begin{aligned} y^{new}[i] - y^{old}[i] &= d_u \, \delta u_i \\ y^{new}[i+1] - y^{old}[i+1] &= \mathbf{CB_u} \delta u_i \\ y^{new}[i+2] - y^{old}[i+2] &= \mathbf{CAB_u} \delta u_i \\ &\cdots \\ y^{new}[i+k] - y^{old}[i+k] &= \mathbf{CA}^{k-1}\mathbf{B_u} \delta u_i \end{aligned}$$

and for $j < i$:

$$y^{new}[j] - y^{old}[j] = 0 \tag{4.35}$$

Results of Chapter 3 show that all the coefficient matrices are nonnegative. Therefore, increasing the input at period i increases most of the future outputs above their previous level. Then at any future period, say period j, the input $u[j]$ can be decreased until the output level at j is back to its previous level. The amount of decrease in the input is:

$$\delta u_j = u^{old}[j] - u^{new}[j] = \frac{1}{d_u}(y^{new}[j] - y^{old}[j]) \tag{4.36}$$

It was also shown in Chapter 3 that if at some period i the input $u[i]$ is increased, the input of some future period j can be decreased until the output of that period $y[j]$ is brought down to its initial level, and still have some residual outputs remaining in the periods after j. In other words, when the output of some period is brought down to a level by a decrease in the input of the same period, the outputs of future periods are also brought down but not necessarily to their previous level.

Suppose the input at period i is increased by δu, and the question is by how much can the input at period $i+1$ be decreased without violating any of the constraints?

To answer this question lets look at the changes in the output block:

$$\mathcal{Y}^{new_1} - \mathcal{Y}^{old} = \mathbf{I}^i \delta u_i \tag{4.37}$$

If the input at period $i+1$ is now decreased by δu_{i+1} then the total changes in the output block become:

$$\mathcal{Y}^{new_2} - \mathcal{Y}^{old} = \mathbf{I}^i \delta u_i - \mathbf{I}^{i+1} \delta u_{i+1} \tag{4.38}$$

Given the value of δu_i we can solve for the value of δu_{i+1} which brings back at least one of the outputs back to its old value. If we denote the jth element of $\mathbf{I}^k$ by i^k_j, then δu_{i+1} is found from:

$$\delta u_{i+1} = \min\left\{ \frac{i^i_{i+1}}{i^{i+1}_{i+1}}, \frac{i^i_{i+2}}{i^{i+1}_{i+2}}, \cdots, \frac{i^i - N}{i^{i+1}_N} \right\} \tag{4.39}$$

Results of Chapter 3 ensure that in the above equation, the minimum is attained by the *first term* inside the braces. Hence, after decreasing the input at $i+1$ by δu_{i+1} the outputs at periods $k > i+1$ would still be above their initial levels. Consequently, the inputs at the following periods can be decreased in the above manner until all the future outputs are back at their initial level. Thus, it is possible to reschedule inputs of more than one period with the increase in the input of a single period.

Elemental direction vectors: The preceding development suggests one possible way to determine a feasible direction vector. An elemental direction vector is determined first, and then it is used to determine feasible direction methods at any solution point. The elemental direction vector associated with period i is denoted by $\mathbf{J_u}^i$. The kth element of this vector is denoted by j^i_k. The following steps determine the elemental direction vector associated with period i:

- Set all the elements of j^i_k to zero for $k < i$.
- Set j^i_i equal to one. This is equivalent to increasing the input at period k by one. As a result the output block will be increased by $\mathbf{I}^i$.
- Determine the amount by which the input of period $i+1$ can be decreased without taking any of the future outputs below their old values. With the increase of the input at period i by one unit we get:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \mathbf{I}^i \tag{4.40}$$

  Now if the input of the second period is to be decreased by j^i_{i+1}, the new output block becomes:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \mathbf{I}^i - \mathbf{I}^{i+1} j^i_{i+1} \tag{4.41}$$

  Therefore, the maximum value of the input decrease at $i+1$ without taking any future output below its old value is:

$$j^i_{i+1} = \min\left\{ \frac{i^i_{i+1}}{i^{i+1}_{i+1}}, \frac{i^i_{i+2}}{i^{i+1}_{i+2}}, \ldots, \frac{i^i_N}{i^{i+1}_N} \right\} \tag{4.42}$$

  where i^i_{i+m} is the $(i+m)$th element of $\mathbf{I}^i$,the impulse response vector associated with period i. *For the thermal systems the minimum is attained by the first term in the braces*, and there is no need to consider the other terms.

- Next find the amount by which the input of the period $i+2$ can be decreased without taking any future outputs below their values. For the ease of computation, define a new vector $\mathbf{L}$ as:

$$\mathbf{L} = \mathbf{I}^i - \mathbf{I}^{i+1} j^i_{i+1} \tag{4.43}$$

Then, if the input at period $i+2$ is decreased by j^i_{i+2}, the new output block becomes:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \mathbf{L} - \mathbf{I}^{i+2} j^i_{i+2} \tag{4.44}$$

Therefore, the value of the maximum value of the input decrease at $i+2$ is determined from:

$$j^i_{i+2} = \min \left\{ \frac{l_{i+2}}{i^{i+2}_{i+2}}, \frac{l_{i+3}}{i^{i+2}_{i+3}}, \cdots, \frac{l_N}{i^{i+2}_N} \right\} \tag{4.45}$$

Again, for the thermal systems, the minimum is attained by the first term in the braces, and there is no need to consider the other terms.

- Redefine $\mathbf{L}$ as:

$$\mathbf{L} = \mathbf{L} - \mathbf{I}^{i+2} j^i_{i+2} \tag{4.46}$$

  and find j^i_{i+3} is the same manner as above.

- Continue, until all the elements of $\mathbf{J_u}^i$ are found. The input at $i+m$ can be decreased by:

$$j^i_{i+m} = \min \left\{ \frac{l_{i+m}}{i^{i+m}_{i+m}}, \frac{l_{i+m+1}}{i^{i+m}_{i+m+1}}, \cdots, \frac{l_N}{i^{i+m}_N} \right\} \tag{4.47}$$

  where the minimum is achieved by the first term in the brackets.

Now, the elemental vector $\mathbf{J_u}^i$ represents a unity increase in the input of the period i, and the maximum sequential decreases in the inputs of future periods without pushing any of the outputs below their old values.

In most situations, direct substitution of the elements of the impulse response vector (in terms of the system coefficient matrices) may result in a simple expression for the elemental direction vector. Indeed, in the case study of the heating of a 2R1C building, the elemental direction vector as expressed in terms of the system coefficients, has a very simple form.

It is important to note that the fact that some future input may be at its minimum value to begin with, was not considered here. Thus although the elemental vector $\mathbf{J_u}^i$ may be considered a direction vector along which the input block $\mathcal{U}$ may be moved, however, such movement may not be feasible.

Properties of the elemental direction vector: The special usefulness of the $\mathbf{J_u}$ vector is that if the input of some period is increased by some value, the decreases in subsequent periods can be found by multiplying the $\mathbf{J_u}$ vector by the scaler value of the increase in the input of the first period. However, there are a few things which must be considered before proceeding as above. The corresponding $\mathbf{J_u}^i$ vector for a period i other than the first can be found from the original vector by shifting the elements of the original vector in time, and substituting zeroes for periods before i. The elemental vector for period i is:

$$\mathbf{J_u}^i = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ j_{i+1}^i \\ j_{i+2}^i \\ \vdots \\ j_N^i \end{bmatrix} \tag{4.48}$$

where the elements are found using the methods of the previous section.

Another property is that any feasible movement can be described as a combination of movements along the elemental directions. For example, if the input of the first period is to be increased, but the rescheduling is to be started with the third period, then the resulting direction vector is:

$$\mathbf{F} = \mathbf{J_u}^1 - j_2^1 \mathbf{J_u}^2 \tag{4.49}$$

where j_2^1, the second element of $\mathbf{J_u}^1$ is negative, and in some cases, zero. Extension to more complex directions follows directly.

In most general problems, once the feasible direction vectors are found, there would be no need to retain the elementary direction vectors, since all future updating of movement direction are based on the feasible direction vectors. However, in simple first order problems, all future feasible direction vectors can be constructed from the elementary direction vectors. These special cases will be described when we discuss applications to space heating/cooling and water heating. Thus, for the sake of generality, the distinction between the elementary and feasible direction vectors is maintained here, even if we do not mention the elementary direction vectors in this chapter again.

Feasible direction vectors: Movement along an elementary direction vector is not necessarily feasible. At some periods, the input may be already at a minimum. Therefore, care must be taken so that the input of that period is not decreased further. This requires modification of the $\mathbf{J_u}$ vectors. Suppose we want to construct the elementary direction vector for period i. However, at period j where $j > i$ the input is already at $u[j]_{min}$. Therefore, the jth element of vector $\mathbf{J_u}^i$ must be kept at zero. The modified vector is:

$$\mathbf{F} = \mathbf{J_u}^i - j_j^i \mathbf{J_u}^j \tag{4.50}$$

If the elemental vector for period i is completely modified by considering all the future periods at which inputs are at their minimums, then the final modified elemental direction vector also becomes a feasible direction vector, denoted by $\mathbf{F_u}^i$, and corresponding to the period i. To emphasize that this vector represents a feasible change in the values of the inputs, its elements are represented as:

$$\mathbf{F_u}^i = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ \delta u_{i+1}^i \\ \delta u_{i+2}^i \\ \vdots \\ \delta u_N^i \end{bmatrix} \tag{4.51}$$

Thus, given an old solution $\mathcal{Y}^{old}$, the algorithm selects a charge period k_c, by a rule described next, and then finds the corresponding feasible direction vector $\mathbf{F_u}^{k_c}$. Then it finds how much it can move in the direction of the $\mathbf{F_u}^{k_c}$ before it meets a new constraint. In other words, it finds the value of α, the proportion of the length of the feasible direction vector, which lies between the old and the new feasible solution. The new solution set is found by:

$$\mathcal{U}^{new} = \mathcal{U}^{old} + \alpha \mathbf{F_u}^{k_c} \tag{4.52}$$

and the new set of outputs is:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \mathcal{D}_u \alpha \mathbf{F_u}^{k_c} \tag{4.53}$$

A feasible direction vector for the output block can be defined based on the last equation:

$$\mathbf{F_y}^{k_c} = \mathcal{D}_u \mathbf{F_u}^{k_c} \tag{4.54}$$

Therefore, the new output block is determined from:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \alpha \mathbf{F_y}^{k_c} \tag{4.55}$$

Again, for emphasis, the elements are denoted as shown below:

$$\mathbf{F_y}^{i} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ \delta y^i_{i+1} \\ \delta y^i_{i+2} \\ \vdots \\ \delta y^i_N \end{bmatrix} \tag{4.56}$$

Thus if the input block is changed by $\mathbf{F_u}$ the output block will change by $\mathbf{F_y}$.

Marginal savings: After finding a new solution point, the algorithm has to decide the new feasible direction for its next move. It has N choices of periods to select as the *charge period* k_c. This is the period at which the input is going to be increased for subsequent rescheduling inputs at later periods. For each period, the algorithm can construct a feasible direction vector. To chose a particular period as the charge period, the algorithm has to evaluate the *marginal savings* s_k for each period. It is equal to the net savings due to the increase of the input at each possible charge period by one unit. Marginal savings of all periods are represented by the *savings* column vector **S** whose kth elements is s_k.

Any movement from any charge period will be proportional to the corresponding feasible direction vector, and the total savings along that direction will remain proportional to the marginal savings of the charge period. Therefore, to determine the marginal savings of each period it suffices to compute the savings and costs associated with the feasible direction vector. Suppose we want to find the marginal savings for period i. Then we have:

$$costs = \delta u_i^i p[i] = p[i] \tag{4.57}$$

$$savings = -\sum_{k=i+1}^{N} \delta u_k^i p[k] \tag{4.58}$$

$$s_i = savings - costs \tag{4.59}$$

where δu_k^i is the kth element of the feasible direction vector $\mathbf{F_u}^i$.

At each new solution point the feasible direction vectors are going to change, since new constraints are encountered. As a result, at each new solution point, the algorithm has to find the new value of s_k for each period k, and to construct a new **S** vector.

The marginal savings defined above is a measure of net savings per unit of increase in the input of the charge candidate. It is possible to define other measures of cost reduction associated with linear movements in the feasible space. One is based on the ratio of savings to costs. This is equivalent to net savings per unit cost of the increase in the heat input of the charge candidate. Another measure is based on the net savings per unit linear movement in the direction of the feasible direction vector. All these measures play the same role in the algorithm and the overall result should be independent of the measure chosen. However, the choice may affect the efficiency of the algorithm.

Charge period at each iteration: To find the charge period at the start of each iteration, the savings vector **S** is searched and the period with the highest marginal savings s_k is designated as the charge period k_c. If no element in the savings vector is greater than one, then there are no further movements that can decrease the total cost. If that is the case, the algorithm terminates.

The periods at which the input or the output is at the maximum, should not be considered. At these periods, inputs can not be increased further without violating the input or the output constraint of that period.

In addition, if at some periods the input is at its minimum and at the same time, the output is at its maximum, then some previous periods have to be excluded from being candidates for the charge period. Suppose at period j the input is at its minimum and the output is at its maximum. Now take period i where $i < j$. If i is taken as the charge period, the input at i is going to be increased, and at the same time the inputs of some future periods are going to be decreased in the direction of $\mathbf{F_u}^i$. However, the input at j can not be decreased further, and the algorithm has already excluded that period from further input reduction. This may result in an increase in the output level at j since the output increase at that period, is no longer going to be brought back to zero. The variable that determines whether the output at j is going to be increased, is the jth element of the feasible direction vector for the output block $\mathbf{F_y}^i$, namely δy^i_j. If the value of this variable is greater than zero, then the period i should not be considered as a possible charge candidate. For certain dynamic systems, the result is that all periods before j are excluded from consideration.

How far to move: After finding the charge period k_c and the associated feasible direction vector $\mathbf{F_u}^{k_c}$ the algorithm must decide the length of the feasible movement. A new solution point is reached when the movement in the feasible direction reaches a new constraint. Since movement is proportional to the length of the feasible direction vector, all needed is to find the proportionality factor α

which takes the movement out of the feasible region.

To find the first constraint that is encountered, an actual movement is made from the old solution point in the feasible direction vector, and the minimum change in each input and output that reaches a lower or an upper limit is taken as desired length of movement. The four possibilities are:

- Control input at the charge period reaches the maximum level.
- Output variable at the charge period reaches the maximum level.
- A control input at some future period reaches the minimum level.
- An output variable at some of the future periods reaches the maximum level.

The last case may happen for those future periods at which the control input is already zero and no further reduction in the control input possible. This causes an increase in the output of that period.

The new solution set: When the proportionality factor α is found, the new inputs and outputs are:

$$\mathcal{U}^{new} = \mathcal{U}^{old} + \alpha \mathbf{F_u}^{k_c} \tag{4.60}$$

and the new set of outputs is:

$$\mathcal{Y}^{new} = \mathcal{Y}^{old} + \mathcal{D}_u \alpha \mathbf{F_u}^{k_c} \tag{4.61}$$

The new solution set sits on a new constraint point in the multi-dimensional space. Additional feasible movement in the direction of the last feasible direction vector is no longer possible. Therefore, it is necessary to determine new feasible direction vectors and new savings to costs ratios associated with each period.

New set of feasible direction vectors: New constraints are met at the new solution point. Therefore, the feasible direction vectors associated with each period in the time horizon must be modified. However, such modification is only necessary if the new constraint is the result of an input reaching its minimum value. This is described in the following section.

Other new constraints, such as an output or an input reaching their maximum level, do not necessitate a modification in the feasible direction vectors. Their effect is exclusion of some periods from consideration as charge periods in the next round of iterations.

Excluding a period from input reduction: If at a new solution point the input at some period reaches its minimum value, then at the next iterations, that period must be excluded from input reduction. Also, when a period is selected as charge period, meaning that it has the highest marginal savings, then it must be excluded from input reduction, since it has already been deemed economical to have its input increased. This necessitates modification of all the feasible direction vectors.

Suppose all the previous sets of feasible direction vectors are given. Also suppose that at period j the new input is at its minimum, or that it was the charge period at the last iteration. The new feasible direction vectors can be considered to be a combination of movements utilizing the old direction vectors. For example, if a movement in the $\mathbf{F_u}^i$ direction is made, where $i < j$, then the input at j will be decreased by δu_j^i. To adjust, a compensating movement in the $\mathbf{F_u}^j$ must be made. Thus, the new feasible direction vector associated with period i is:

$$\mathbf{F_u}^i = \mathbf{F_u}^i + \delta u_j^i \mathbf{F_u}^j \tag{4.62}$$

This adjustment must be done for all the feasible direction vectors associated with periods before j. For periods after j the feasible direction vectors will not change, since $\delta u_j^i = 0$ if $i > j$.

If at the new solution point more than one period reaches its minimum levels, then the exclusions must be carried for all such periods.

Including a period for input reduction: Suppose the output at period j, which has already been excluded from input reduction, reaches its maximum level. In that case, many periods before period j can not be considered as a possible charge period, since without input reduction the output at j may increase. However, if period j was one of the periods which was excluded because it was a charge period, then some savings may be possible if the input at that period can be decreased in order to let the prior periods to be considered as charge candidates.

Thus, when the output at a previous charge period reaches the maximum, it is included in the input reduction schemes. To find the new set of feasible direction vectors, the reverse of the exclusion methodology is employed. For example the new feasible direction for period i becomes:

$$\mathbf{F_u}^i = \mathbf{F_u}^i - \delta u_j^i \mathbf{F_u}^j \tag{4.63}$$

Note that δu_j^i is either zero or negative for $j > i$.

Again, this adjustment must be done for all the feasible direction vectors associated with periods before j. For periods after j the feasible direction vectors will not change, since $\delta u_j^i = 0$ if $i > j$.

New marginal savings: If the set of feasible direction vectors were modified after the finding of the new solution set, then the associated marginal savings have to be modified as well.

Suppose at the last iteration the input at period j is brought to its minimum level. This would necessitate a modification of the feasible direction vectors, as described in the last section. Then, all that is necessary is to evaluate the savings and costs of the modification for each feasible direction vector. This can be done at the same time that each feasible direction vector is being modified. For example, given the old version of the feasible direction vector $\mathbf{F_u}^i$ associated with period i where $i < j$, the new version is found by moving in the old direction and then compensating for violation at period j as shown previously. The compensation is equivalent to moving in the direction of $\mathbf{F_u}^j$ by the amount of δu^i_j. Thus, the old savings to costs ratio for period i is modified according to:

$$s_i^{new} = s_i^{old} - \delta u^i_j(-p[j] + \sum_{k=j+1}^{N} p[k]\delta u^j_k) \tag{4.64}$$

The modifying term is subtracted if period j is being excluded. The operation must be changed to addition if period j is being included back. The expression in the parenthesis is equivalent to the change in the savings associated with a unit movement in the $\mathbf{F_u}^j$ direction and does not depend on the period i. Therefore, in modifying the marginal savings for all i, where $i < j$, the only expression which requires reevaluation is the fractional expression. Also, for $i > j$ marginal savings do not change, since the feasible direction vectors do not change for i equal or greater than j.

4.6 Properties of the Algorithm and Its Relation to the Simplex Method

The following list summarizes the basic properties of the algorithm and how it differs from the simplex method:

- The new algorithm is applicable to only a subclass of linear programming problems, namely, those with structures similar to asymptotically stable positive dynamic systems. No attempt has been made to extend the algorithm to include a more general class of problems.

- The new algorithm utilizes the superposition property and the simplicity of impulse response vector.

- At each iteration, a movement in the feasible space is made only if the total cost can be reduced.

- The linear movements are made within the feasible space, and thus, the feasibility is always maintained.

- The algorithm terminates when no further cost saving scheduling is possible. Due to the convexity of the feasible space, the local optimum is global.

- For first-order systems, optimality of the algorithm can be proven theoretically by the analysis of the equivalent network problem. The final solution results in a *spanning tree* which can not be improved further by any method.

- All the coded algorithms used in the case studies give results which are identical or equivalent (in case of multiple optimal solutions) to the final results of the simplex method.

- Based on the selection criteria for each iteration, the new algorithm chooses the movement with the highest cost reduction.

- For all the classes of problems used in the case studies, both the new algorithm and the simplex algorithm were written in the same language (**APL**) and run on the same environment and computer. Comparisons indicate that the simplex method requires a higher number of iterations for the optimal solution. The difference becomes more significant as the size of the problem is increased.

- For a time horizon of N periods, the simplex method, using the block formulation, consists of a tableau of size $4N \times 5N$. This is because of the $4N$ inequality constraints, and the inclusion of the associated slack and surplus variables. The original state-space formulation will result in an even larger tableau. In comparison, the new algorithm works with vectors of size N. If no simple formulation exist for feasible direction vectors, then at most, the algorithm works with N vectors of size N.

- The simplex method requires many scalar computations in each iteration. In contrast, the new algorithm uses vector operations. For time-invariant systems, all of the feasible direction vectors are based on the original elemental direction vector, and at each iteration, they are easily updated.

- For all the models tested, the algorithm performs faster than the equivalent simplex algorithm written by the author, and also faster than a commercially available (LP83) linear programming package.

- The simplex algorithm approaches the efficiency of the new algorithm for shorter time horizons.

- For a time horizon of 168 periods and under various parameter values, the time used by the 0R1C algorithm (water heater case study) varied between 4 to 6 minutes. In comparison, the simplex algorithm took between 70 to 125 minutes.

4.7 Convergence and Optimality

This section provides an itemized list of arguments which show the feasibility and convergence of solution points at each iteration, together with the optimality of the final solution. All the arguments are based on the results of the preceding mathematical presentation.

- All the intermediate solutions are feasible. The algorithm moves in the feasible space only.

- At each iteration total costs are reduced. No movement with increasing costs are allowed.

- Each iteration is based on a linear movement in the feasible space. There are a finite number of pre-defined elemental movements.

- Any other linear movement in the feasible space is a linear combination of the elemental movements.
- Each new iteration is based on a linear movement with the highest possible reduction in the total costs.
- Each new solution point is found when a new constraint becomes active.
- There is a finite number of combinations for intermediate solution points.
- If no further cost reducing movement is possible, then the last solution is optimal.

Chapter 5

SPACE HEATING AND COOLING UNDER SPOT PRICES

5.1 Case Study: A 2R1C House

5.1.1 The model

The 2R1C (two resistances and one capacitance) model of a house assumes that the thermal masses of the air inside the house and the external shell are negligible compared to the internal walls and the furniture. The last two are represented by a single lump of thermal mass. The heating system heats the inside air, and the resulting heat is transferred to the internal mass of the house, and alsc to the outside through the external shell and filtration. Figure .1 shows an equivalent 2R1C circuit model of the house discussed here.

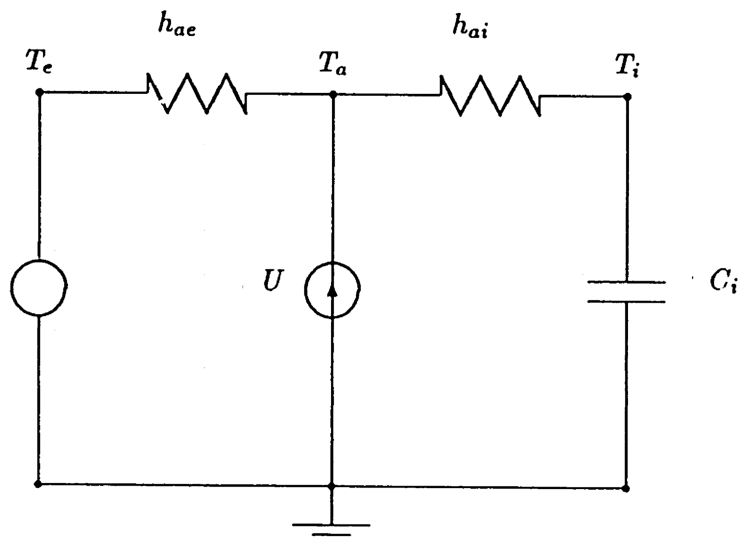

T_e	External Temperature
T_a	Air Temperature
T_i	Internal Mass Temperature
U	Heat Input
C_i	Internal Thermal Mass
h_{ae}	Heat Transfer between air and External Temperature
h_{ai}	Heat Transfer between air and Internal Thermal Mass

Figure 5.1: 2R1C Heat Transfer Model of a House

In addition to its simplicity, this special model was chosen for the case study because of the availability of data for its parameters. Sondregger [S8] used parameter estimation techniques to determine the parameters for a 2R1C model of a townhouse in Twin Rivers, New Jersey. Wilson, et all [W3] used a similar technique in estimating the parameters for a 2R1C model of a gas-heated house in Windsor, Ontario. In the case study here, the values of the parameters were chosen to be of the same order of the magnitude as the values reported in the studies. The single storage model can be easily extended to include solar heating and direct interaction between the storage and the ambience, without any major modifications in the arguments of this chapter. However, lack of reliable data on the additional parameters, makes it wise to make the model less complicated than the model of the studies mentioned above.

For the 2R1C model, the energy balance equations in discrete time are:

$$U[k] = h_{ai}\Delta t(T_a[k] - T_i[k]) + h_{ae}\Delta t(T_a[k] - T_e[k]) \tag{5.1}$$

$$0 = C_i(T_i[k+1] - T_i[k]) + h_{ai}\Delta t(T_i[k] - T_a[k]) \tag{5.2}$$

Rearranging the first equation and substituting in the second equation results in the following set of equations:

$$T_i[k+1] = aT_i[k] + b_u U[k] + b_w T_e[k] \tag{5.3}$$

$$T_a[k] = cT_i[k] + d_u U[k] + d_w T_e[k] \tag{5.4}$$

where the coefficients are:

$$a = 1 - \frac{h_{ai}\Delta t}{C_i} + \frac{h_{ai}^2 \Delta t}{C_i(h_{ai}+h_{ae})} \tag{5.5}$$

$$b_u = \frac{h_{ai}}{C_i(h_{ai}+h_{ae})} \tag{5.6}$$

$$b_w = \frac{h_{ai}h_{ae}\Delta t}{C_i(h_{ai}+h_{ae})} \tag{5.7}$$

$$c = \frac{h_{ai}}{h_{ai}+h_{ae}} \tag{5.8}$$

$$d_u = \frac{1}{\Delta t(h_{ai}+h_{ae})} \tag{5.9}$$

$$d_w = \frac{h_{ae}}{h_{ai}+h_{ae}} \tag{5.10}$$

This formulation conforms to the set up that was used in algorithm of Chapter 4. This model ignores all the other heat inputs such as solar and non-electrical outlays. Also ignored are the inside humidity and outside non-temperature weather variables such as wind. It is possible to include other heat sources if they are deterministic. Also, it is possible to represent the outside non-temperature weather variables with equivalent temperature potentials.

The above model also assumes time-invariant coefficients. This, and the low order of the dynamic system, can be used in simplifying the cost optimization algorithm. It is important to note that the controlled output (inside air temperature) is not a state variable itself. If it were otherwise, the algorithm would be simplified even further, and made more efficient as it is the case in the electric water heater study of the next chapter. The algorithm used in the case study here is presented and explained in Appendix B.

5.2 Statement of the Optimization Problem

A formal statement of the optimization problem results in a linear programming formulation. Suppose the time horizon of interest consists of N periods of equal times. The period length must be taken to be smaller than the thermal time constant of the house; and the time horizon must be long enough so that most of the savings potential due to the on-off scheduling can be realized. If the time horizon is not long enough then the storage will not have sufficient time to return most of the stored heat. Using the terminology developed so far the problem can be stated as follows:

$$\min_{U[k]} \sum_{k=1}^{N} p_k \times U[k] \tag{5.11}$$

subject to the following for all k:

$$T_i[k+1] = aT_i[k] + b_u U[k] + b_w T_e[k] \tag{5.12}$$

$$T_a[k] = cT_i[k] + d_u U[k] + d_w T_e[k] \tag{5.13}$$

$$T_{min}[k] \leq T_a[k] \leq T_{max}[k] \tag{5.14}$$

$$U_{min}[k] \leq U[k] \leq U_{max}[k] \tag{5.15}$$

In the Appendix B, the optimization problem is reformulated using block variables which gives insights into ways of simplifying the general algorithm. Before we describe the case study and the algorithm, let us address the issue of parameter estimation for the 2R1C model.

5.3 Parameter Estimation for the 2R1C Model

For practical purposes, the storage temperature must be eliminated from the equations.

From the output equation we have:

$$T_i[k] = \frac{1}{c}T_a[k] - \frac{d_u}{c}U[k] - \frac{d_w}{c}T_e[k] \tag{5.16}$$

Then by substituting for $T_i[k+1]$ and $T_i[k]$ in the state equations we get:

$$T_a[k+1] = aT_a[k] + d_u U[k+1] + (cb_u - ad_u)U[k] + d_w T_e[k+1] + (cb_w - ad_w)T_e[k] \tag{5.17}$$

or, using a simpler notation:

$$T_a[k+1] = \theta_1 T_a[k] + \theta_2 U[k+1] + \theta_3 U[k] + \theta_4 T_e[k+1] + \theta_5 T_e[k] \tag{5.18}$$

where, these coefficients in terms of the original heat transfer parameters are:

$$\theta_1 = a = 1 - \frac{h_{ai} h_{ae} \Delta t}{C_i (h_{ai} + h_{ae})} \tag{5.19}$$

$$\theta_2 = d_u = \frac{1}{\Delta t\,(h_{ai} + h_{ae})} \tag{5.20}$$

$$\theta_3 = cb_u - ad_u = \frac{\Delta t\, h_{ai} - C_i}{\Delta t\, C_i (h_{ai} + h_{ae})} \tag{5.21}$$

$$\theta_4 = d_w = \frac{h_{ae}}{h_{ai} + h_{ae}} \tag{5.22}$$

$$\theta_5 = cb_w - ad_w = \frac{h_{ae}(\Delta t\, h_{ai} - C_i)}{C_i (h_{ai} + h_{ae})} \tag{5.23}$$

By making numerous observations on $T_a[k]$, $U[k]$, and $T_e[k]$ for many consecutive periods, an over-constrained system of linear equations is formed. A simple *least square method* can then be applied to find the optimal estimate of the θs. It is then necessary to use these estimates in turn, to estimate the values of original coefficients from the relationships given above. As can be seen, these relationships are not linear, and some type of iterative method is required. Also, for Tactical Control it is not necessary to have values for the original heat transfer parameters. What is essential, is to have values for the coefficients of the state and output equations. In other words, the algorithm requires values for a, b_u, b_w, c, d_u, and d_w. However, to understand the results of the case study in physical terms, and to relate them to the thermal properties of the building, it is necessary to know the values of C_i, h_{ai}, and h_{ae} as well.

5.4 The Case Study

This section presents the results of various numerical simulations and the optimal response of a 2R1C building.

Base Case Parameters

The coefficients chosen for the 2R1C building are the intermediate values between those reported by Sondregger [S8] and those reported by Wilson, et all [W3]. These values are shown in Table 5.1. The values reported for the control heat input in the two papers are 6 and 30 kW respectively. Thus, our choice of 6 kW for control heat input is merely representational. The objective is to study the optimal behavior under spot prices. In addition to ignoring humidity, solar heat input, and stochasticity, there are other short-comings in the study, which have to be taken into account for an actual economic assessment of scheduling under spot prices. The most important is the simplicity of the model, and also, the choice of the time step (1 hour), which may be too long for the discrete time representation of the building's thermal dynamics. See Chapter 3 for a discussion of the choice

thermal capacity C_i	2.00 kWh/C
heat transfer coef. h_{ai}	0.5 kW/C
heat transfer coef. h_{ae}	0.3 kW/C

Table 5.1: Values of Coefficients for the 2R1C model

of time step for the discretization of continuous time systems. The choice of a physically meaningful time step for a 2R1C building is discussed in the following section.

The coefficients h_{ai} and h_{ae} are the heat transfer rates between the inside air and the internal storage, and the inside air and the ambience, respectively. Other factors such as internal heat sources due to people and appliances, and also the heating due to solar incidence are ignored here. All the weather variables are represented by the deterministic exogenous variable T_e. Other base case parameters are shown in Table 5.2.

In actual control, the initial storage temperature must be evaluated from previous observed data. This is only possible if the system is observable in the mathematical sense. The maximum heat input per period depends on the power rating and the efficiency of the heating system. The bounds on the inside air temperature should be supplied by the residents. The choice of time step depends on the frequency of price and exogenous variable changes. It must also be short enough for the model to be physically meaningful. The choice of time horizon depends on the frequency of the price schedule update, which is assumed to be every 24 hours.

Choice of Time Step

For the discrete time 2R1C model to be physically meaningful, the coefficient a must satisfy the following inequality:

$$0 < 1 - \frac{h_{ai}\Delta t}{C_i} + \frac{h_{ai}^2\Delta t}{C_i(h_{ai} + h_{ae})} < 1 \tag{5.24}$$

period length, time step Δt	1 hour
number of periods, time horizon N	24 hours
constant external temperature T_e	12 C
maximum heat input U_{max}	6.0 kWh/period
maximum inside air temperature T_{max}	22 C
minimum inside air temperature T_{min}	18 C
initial storage temperature	18 C

Table 5.2: Values of Base Case Parameters Used in the Study

Then, solving for Δt results in:

$$0 < \Delta t < \frac{C_i(h_{ai} + h_{ae})}{h_{ae}h_{ai}} \tag{5.25}$$

or

$$0 < \Delta t < \frac{C_i}{h_{ai}} + \frac{C_i}{h_{ae}} \tag{5.26}$$

This result can be interpreted as requiring the time step to be less than the sum of system time constants. Also, the time step must be smaller than the individual system time constants. For our choice of coefficients, the time step of 1 hour seems to be well within the accepted limits.

System Parameters vs Thermal Coefficients

In the case studies based on various price patterns, the values of thermal coefficients C_i, h_{ai}, and h_{ae} are varied around their base values in order to investigate their impact on savings. However, from a systems dynamics point of view, the system behavior is best described in terms of system coefficients and parameters which are algebraic functions of these thermal coefficients. Therefore, for future reference, the variation of system parameters with the thermal coefficients is reported here. The system parameters are divided into three groups as shown in Table 5.3.

The state and output coefficients have been defined before, and are the coefficients of the state and output equations.

state coefficients	a, b_u, and b_w
output coefficients	c, d_u, and d_w
time constants	τ, μ

Table 5.3: Grouping of System Coefficients and Parameters

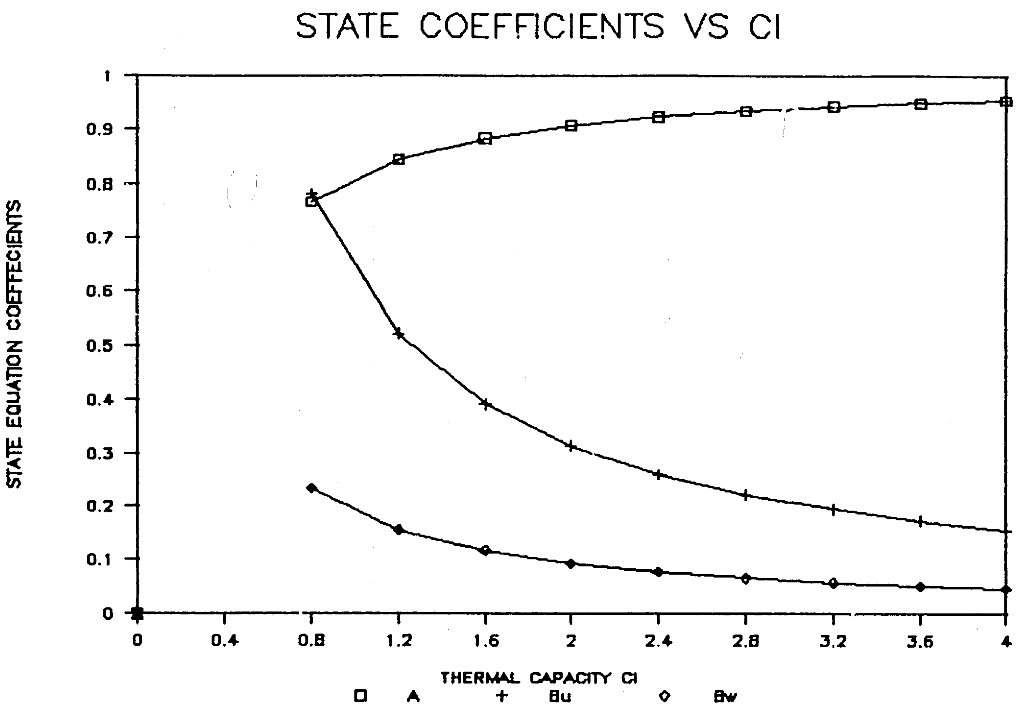

Figure 5.2: State Coefficients vs Thermal Capacity

Figures 5.2, 5.3, and 5.4 depict the variations in system coefficients and parameters in terms of the variations in the thermal capacity C_i.

There are three time constants reported. These are:

$$T_1 = \frac{C_i}{h_{ai}} + \frac{C_i}{h_{ae}} \tag{5.27}$$

$$T_2 = \frac{C_i}{h_{ai}} \tag{5.28}$$

$$T_3 = \frac{C_i}{h_{ae}} \tag{5.29}$$

Figures 5.5, 5.6, and 5.7 depict the variations in system coefficients and parameters in terms of the variations in the heat transfer coefficient h_{ai}.

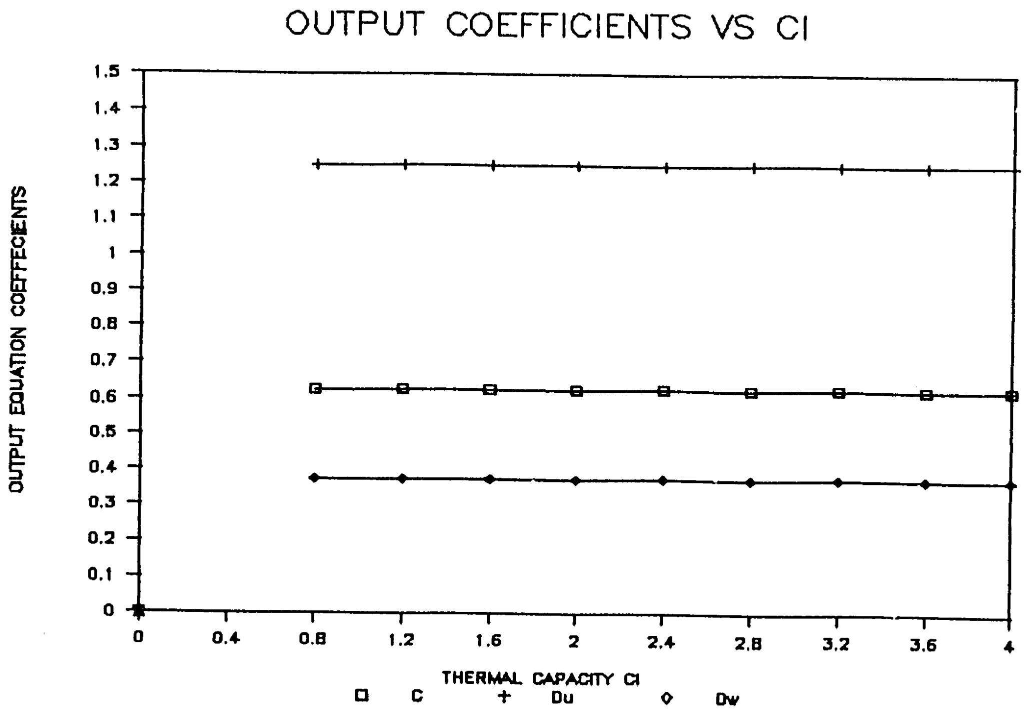

Figure 5.3: Output Coefficients vs Thermal Capacity

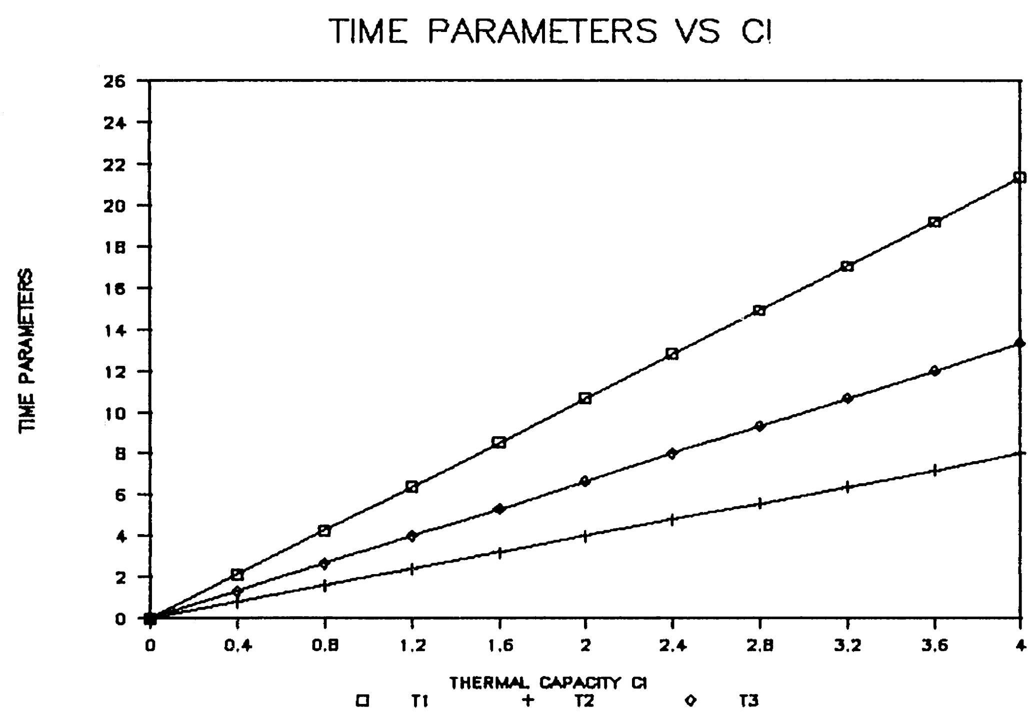

Figure 5.4: Time Constants vs Thermal Capacity

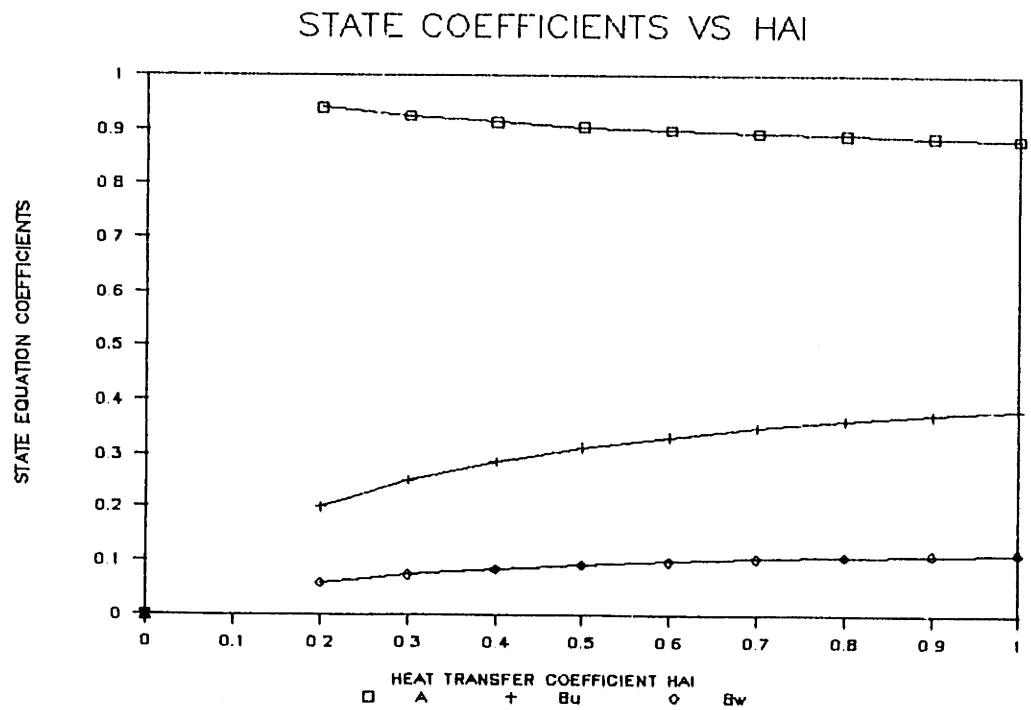

Figure 5.5: State Coefficients vs Heat Trans. Coefficient h_{ai}

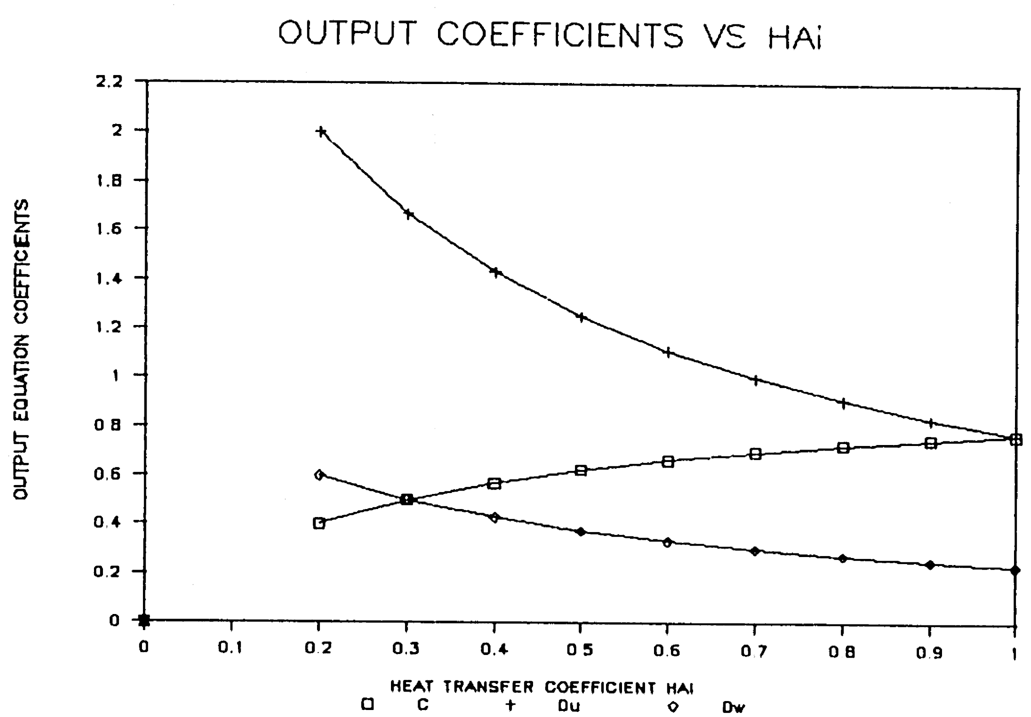

Figure 5.6: Output Coefficients vs Heat Trans. Coefficient h_{ai}

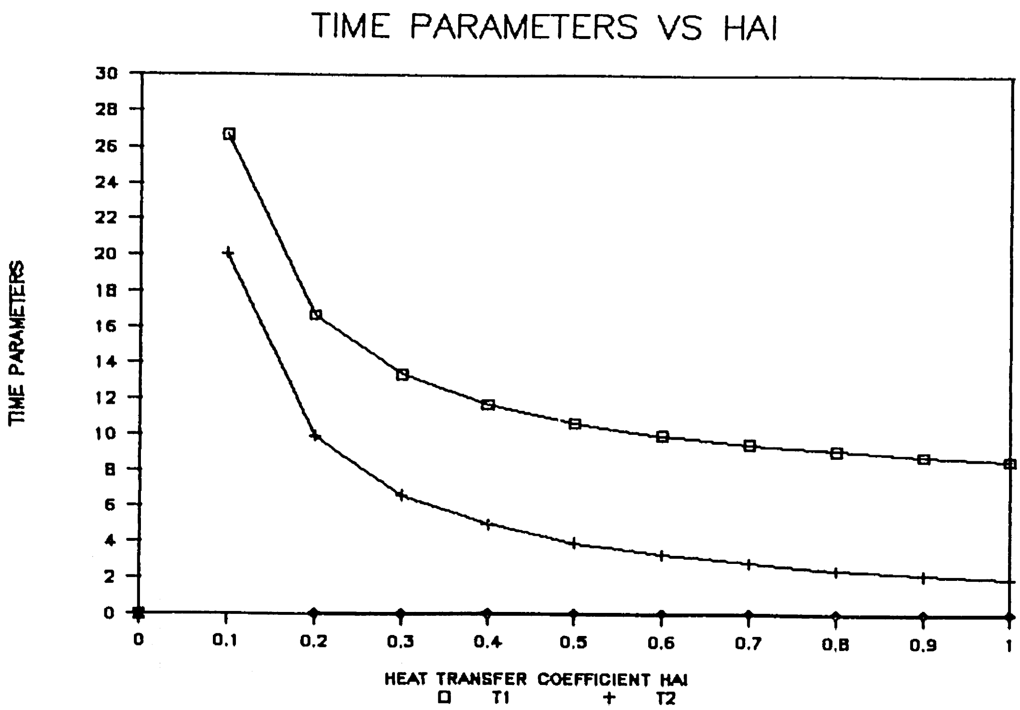

Figure 5.7: Time Constants vs Heat Trans. Coefficient h_{ai}

Figures 5.8, 5.9, and 5.10 depict the variations in system coefficients and parameters in terms of the variations in the heat transfer coefficient h_{ae}.

Definition of Savings

To evaluate the benefits of optimal scheduling, the resulting total costs are compared to the case where no scheduling occurs. The optimal costs are based on the total cost of optimally scheduled electricity usage.

If the prices are assumed constant, then the optimal schedule of electricity usage is the result of an optimization problem where the total electricity usage is minimized. This solution corresponds to a situation where at each period electricity is used only to the extend where the minimum service requirements are met. This is equivalent to solving the output equation at each period k for the heat input $u[k]$, by setting the air temperature of period k at T_{min}. This result corresponds to the *initial* solution of the algorithm used here. It is, as mentioned before, the optimal solution if prices are taken to be flat. Thus, this *reference*

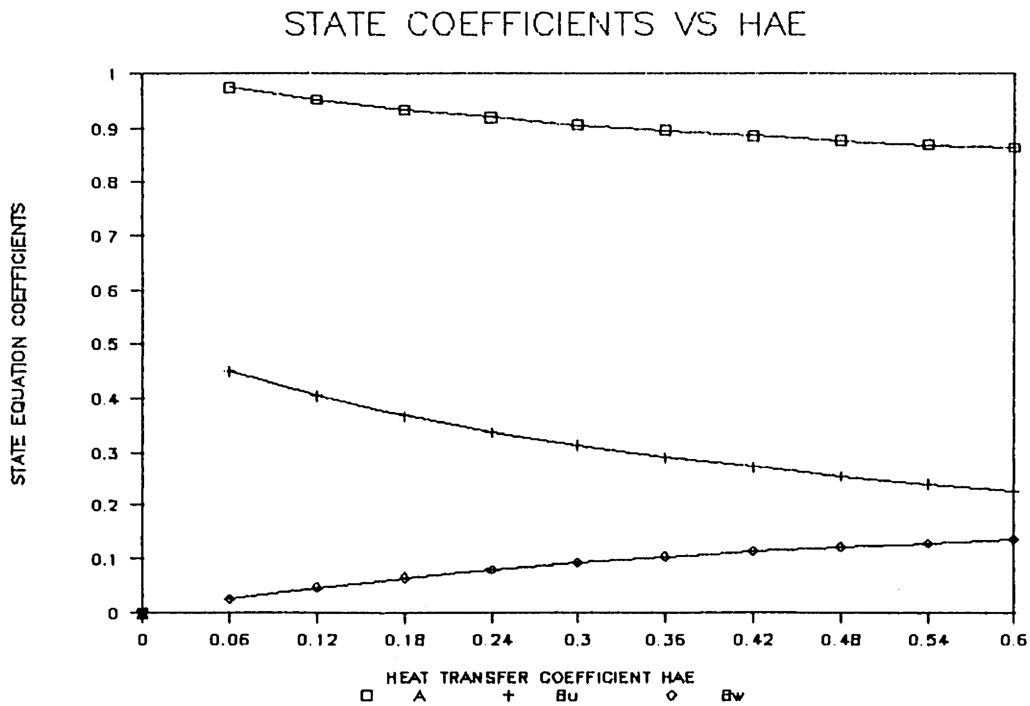

Figure 5.8: State Coefficients vs Heat Trans. Coefficient h_{ae}

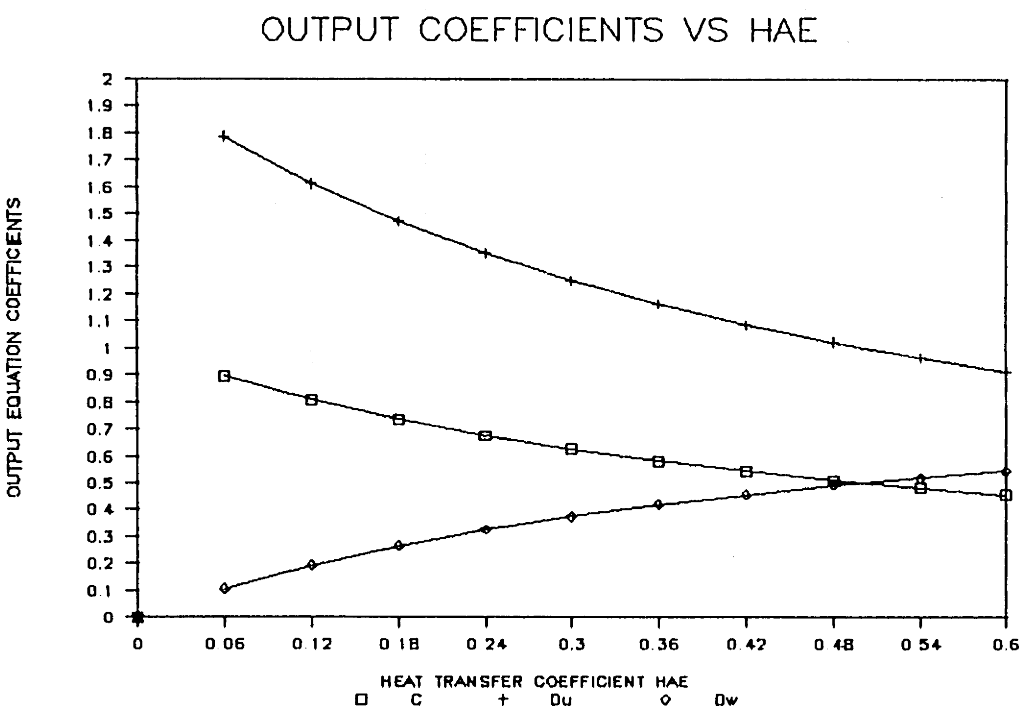

Figure 5.9: Output Coefficient vs Heat Trans. Coefficient h_{ae}

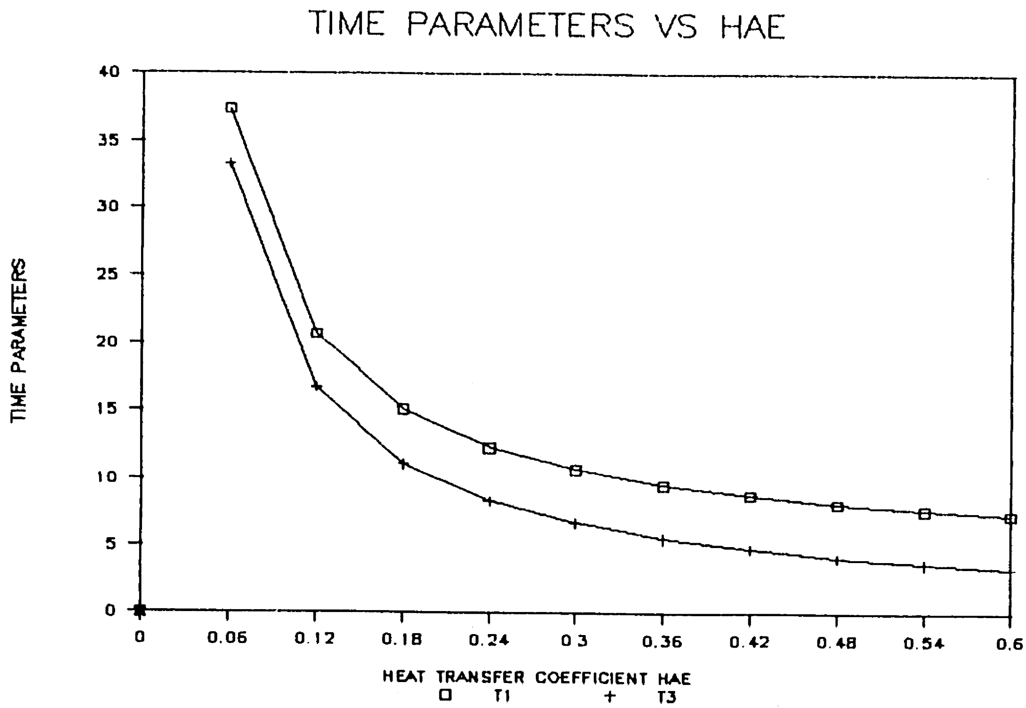

Figure 5.10: Time Constants vs Heat Trans. Coefficient h_{ae}

solution and the corresponding cost provide a basis for the comparison of the optimal performance. Both the reference solution and the final optimal solution can be determined by the application of any linear programming algorithm, and as such are independent of the algorithm being used. Furthermore, for different price patterns, the *base* solutions of the electricity usage will be identical, because for each case the prices are ignored and are assumed to be flat. However, the base cost are evaluated from the price pattern used in determination of the optimal schedule.

Then, according to the above description, the savings rate is defined as the difference between the cost of the optimal solution and the reference solution over the cost of the reference solution. The procedure is as follows:

1. For a given price pattern, ignore the prices and assume any flat price pattern.
2. Determine the *reference solution* which corresponds to the total minimum (or the non-scheduled) electricity usage during the time horizon. Call these

$U_{ref}[k]$.

3. Evaluate the *reference cost* using the price pattern being studied. Thus:

$$\text{reference cost} = \sum_{k=1}^{N} p[k]U_{ref}[k]$$

4. Determine the *optimal solution* given the price pattern. Call these $U_{opt}[k]$.

5. Evaluate the *optimal cost* using the price pattern:

$$\text{optimal cost} = \sum_{k=1}^{N} p[k]U_{opt}[k]$$

6. Net savings is evaluated from:

$$\text{net savings} = \text{reference cost - optimal cost}$$

7. Savings rate is evaluated from:

$$\text{savings rate} = \text{net savings / reference cost}$$

For a given thermal system, the *reference solution* will be the same no matter what the price pattern is, but the *reference cost* will depend on the given price pattern.

Since the reference solution found by the algorithm ignores the prices and is based on the minimum of total heat input during the time horizon, it is independent of the given price pattern. Figure 5.11 shows the initial *reference* solution for the heat input. Figure 5.12 shows the resulting inside air temperature. All the optimal solutions for the price patterns studied are compared to this reference solution.

As can be seen, the initial heat input and inside air temperature solutions are constant. This is partially due to the fact that only a *constant* outside temperature is considered. Another reason is that the reference solution tries to keep the inside air temperature at the lowest level allowable, and the initial storage temperature was also set at the T_{min}. If the initial storage temperature is set at a different value, then the reference heat input will approach a constant level in the steady state.

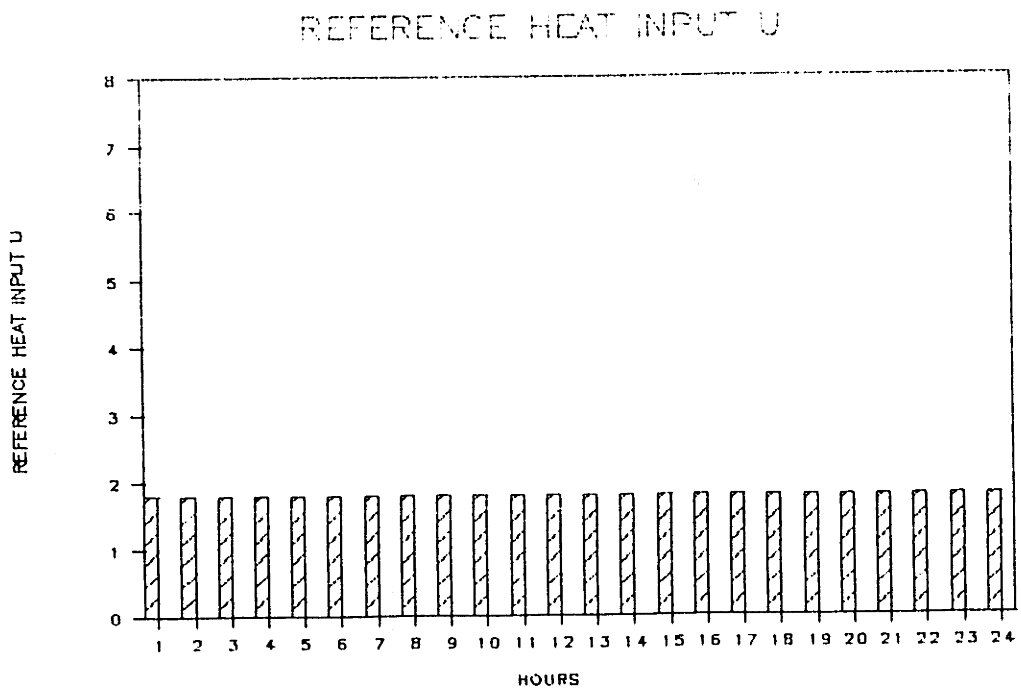

Figure 5.11: Reference Heat Input

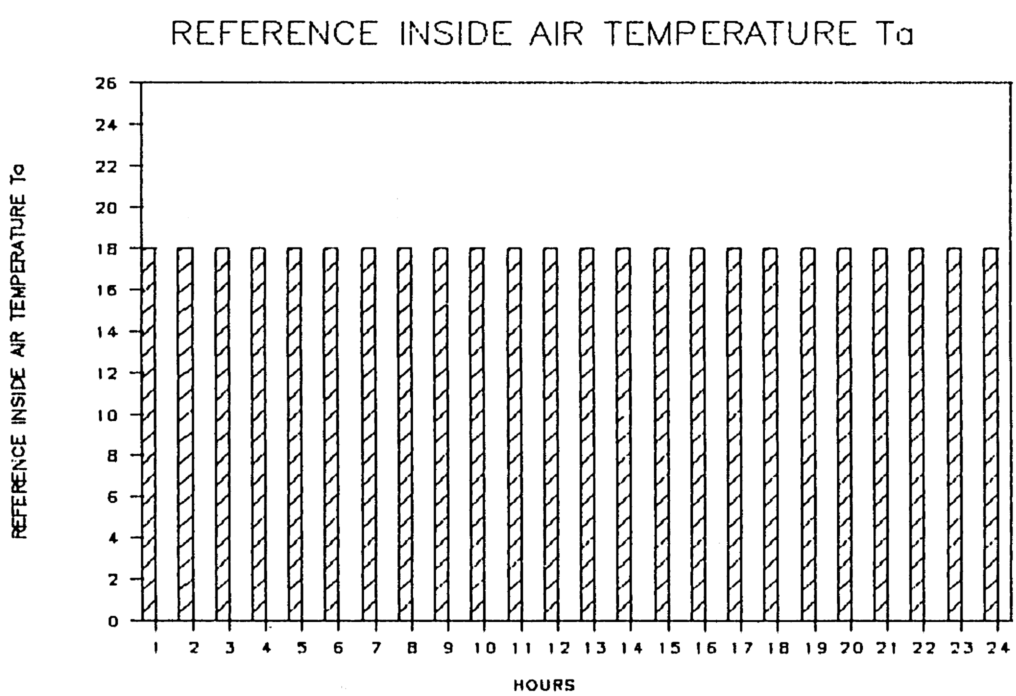

Figure 5.12: Reference Inside Air Temperature

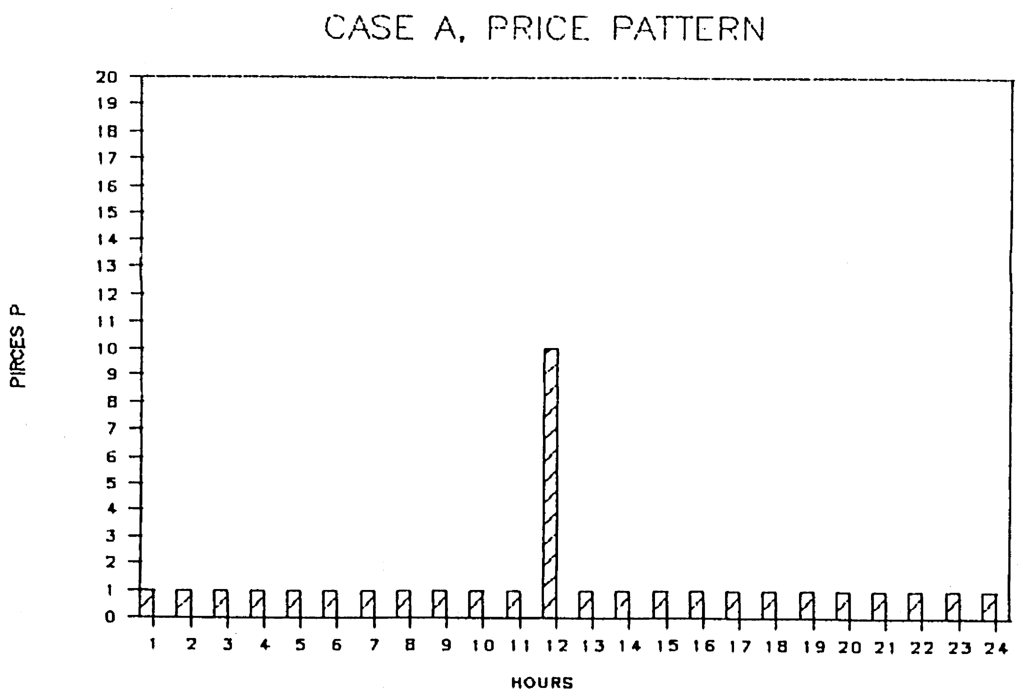

Figure 5.13: Case A, Impulse Price Pattern

5.4.1 Case study A: impulse price pattern

The first study is concerned with the optimal response to an impulse type price pattern, where prices are constant at every hour except one during the time horizon. An example of an impulse price pattern in shown in Figure 5.13.

Figures 5.14 and 5.15 show the results for the optimal solution. The price impulse happens at the hour 12. Thus, the heat input is scheduled to be higher than its initial values at earlier hours in order to produce sufficient stored heat to substitute for the heat input at hour 12.

Heat input levels at hours after 12 are also decreased below their reference levels. This is due to the slow dissipation of the stored heat that is carried away beyond the hour 12. Thus, it is possible to decrease heat input of some future hours in order to push the inside air temperature down to its minimum level. Another interesting behavior is the behavior of heat inputs at hours before 12.

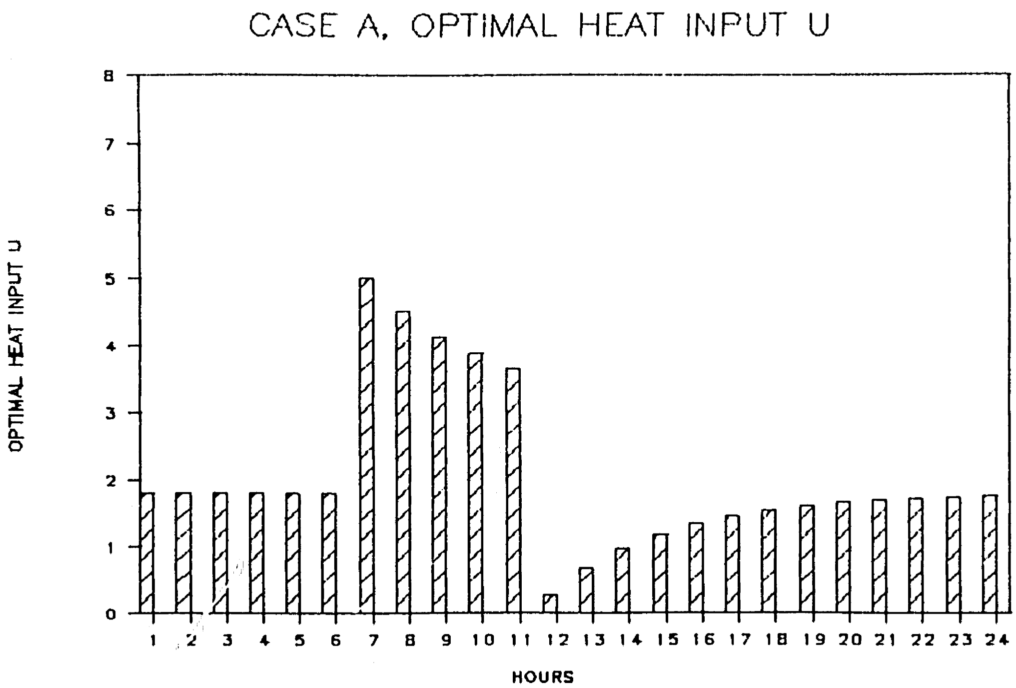

Figure 5.14: Case A, Optimal Heat Input

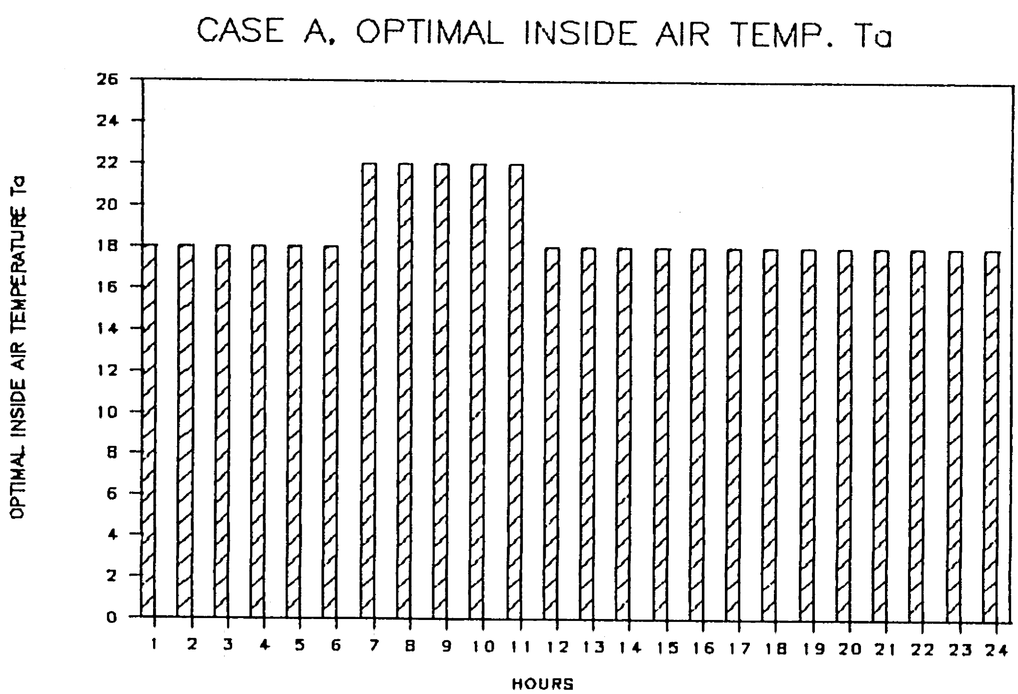

Figure 5.15: Case A, Optimal Inside Air Temperature T_a

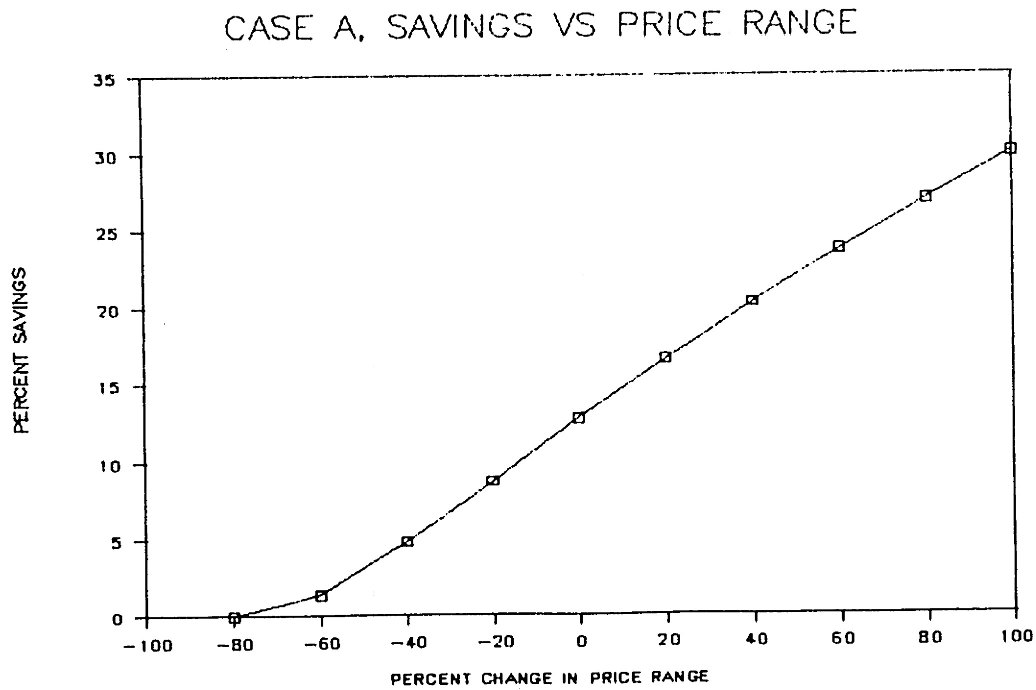

Figure 5.16: Case A, Savings vs Size of the Price Range

The algorithm, at its first iteration increases the heat input at hour 11 as much as possible, which is more than the final level shown in Figure 5.14. In this case, the air temperature at hour 11 reaches the maximum before the heat input has had the chance to reach U_{max}. However, the stored heat, at the first solution is not enough to completely reduce the heat input at hour 12. Thus, at the next solution, the heat input at hour 10 is also increased. However, to do so, the heat input at hour 11 must be lowered just a little bit so that the inside air temperature at hour 11 does not go beyond T_{max}. This process is continued until no additional heat input increases become economical. As can be seen, the heat input at hour 12 is no eliminated completely.

Savings vs Size of the Price Impulse

Savings increase with the size of the price impulse. Figure 5.16 demonstrates this fact. The location of price impulse was kept at hour 12, but the range of the price jump was varied. At price ratios below 3 no savings were observed.

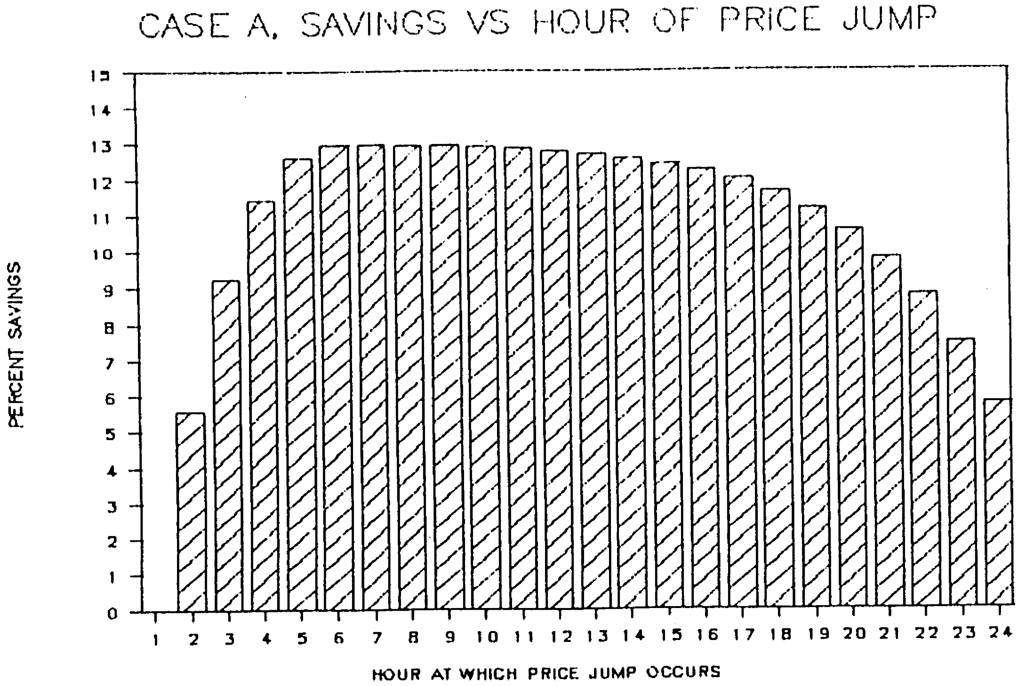

Figure 5.17: Case A, Savings vs Hour of Price Impulse

Savings vs the Hour of Price Impulse

The period within the time horizon at which the price impulse occurs affects the rate of savings. This is due to the cutoff conditions imposed by the length of the time horizon. Figure 5.17 illustrates the results.

As expected there is no savings if the price jump occurs at the first period. There is no chance to heat up the storage. The maximum savings results if the price jump occurs at hour 6. By this hour there have been enough time to economically heat up the storage and reduce the heat input at the hour of price impulse. However, there are also some savings associated with the input reductions of the future hours made possible by the residual heat remaining in the storage. As the location of the price impulse is moved closer to the end of the time horizon, then the size of these residual savings are decreased. It must be noted that the savings beyond the time horizon are ignored. This is a major characteristic of finite time horizon problems. Figure 5.16 also indicates that the size of residual savings

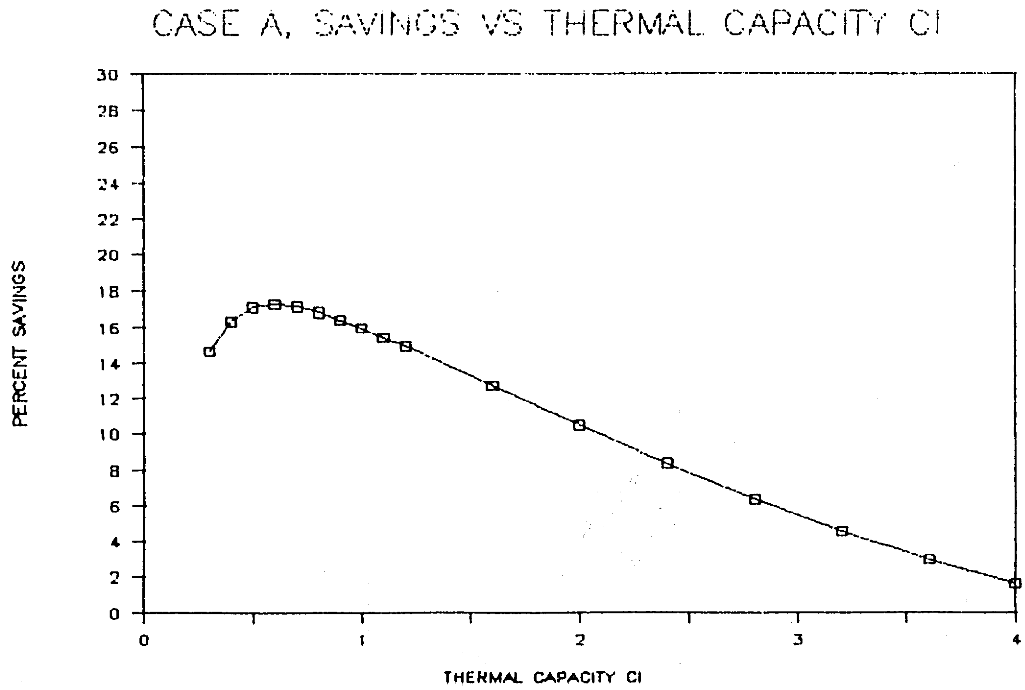

Figure 5.18: Case A, Savings vs Thermal Capacity C_i

can be significant, and in this case they may account for almost half of the total savings.

In practice, the problem of finite time horizon is dealt with by applying the algorithm at every hour as new price information becomes available. However, it is preferable to extend the range of the time horizon if additional price and weather information are made available.

Savings vs Thermal Capacity C_i

Figure 5.18 depicts the variations in the savings rate with the variation in the thermal capacity. For an impulse price pattern, a small thermal capacity provides the better savings rate. This is due to the fact that the thermal storage must be charged with as little heat input as possible, since it is going to be used mainly to displace the heat input of only one period.

It is interesting to note that for the other price patterns, the reverse is true,

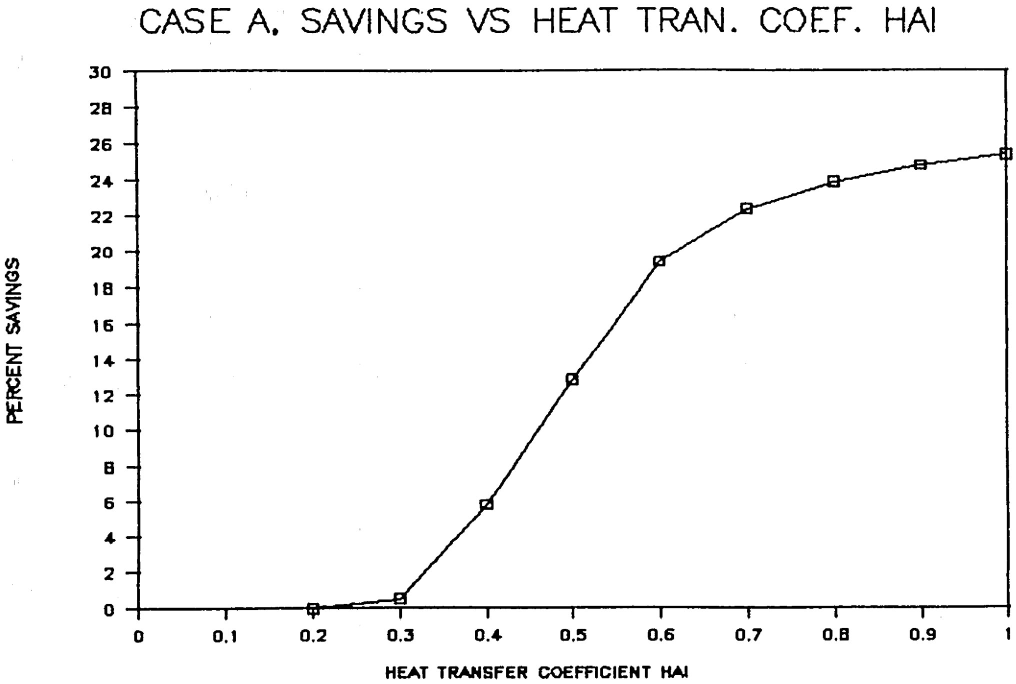

Figure 5.19: Case A, Savings vs Heat Transfer Coefficient h_{ai}

and a larger thermal capacity provides a higher savings rate. As can be seen there is an optimal value for the thermal capacity corresponding to each particular price pattern. As such, the algorithm can be used as a design tool in the evaluation of the optimal thermal storage size for a given price pattern.

Savings vs Heat Transfer Coefficient h_{ai}

Figure 5.19 shows the variation of the savings rate with the variation of the inside air-storage heat transfer coefficient h_{ai}.

As expected, savings improve as the heat transfer rate between the inside air and the storage medium is increased.

Savings vs Heat Transfer Coefficient h_{ae}

Figure 5.20 shows the variation in the savings rate with respect to the variation in the heat transfer rate between the inside air and the ambience.

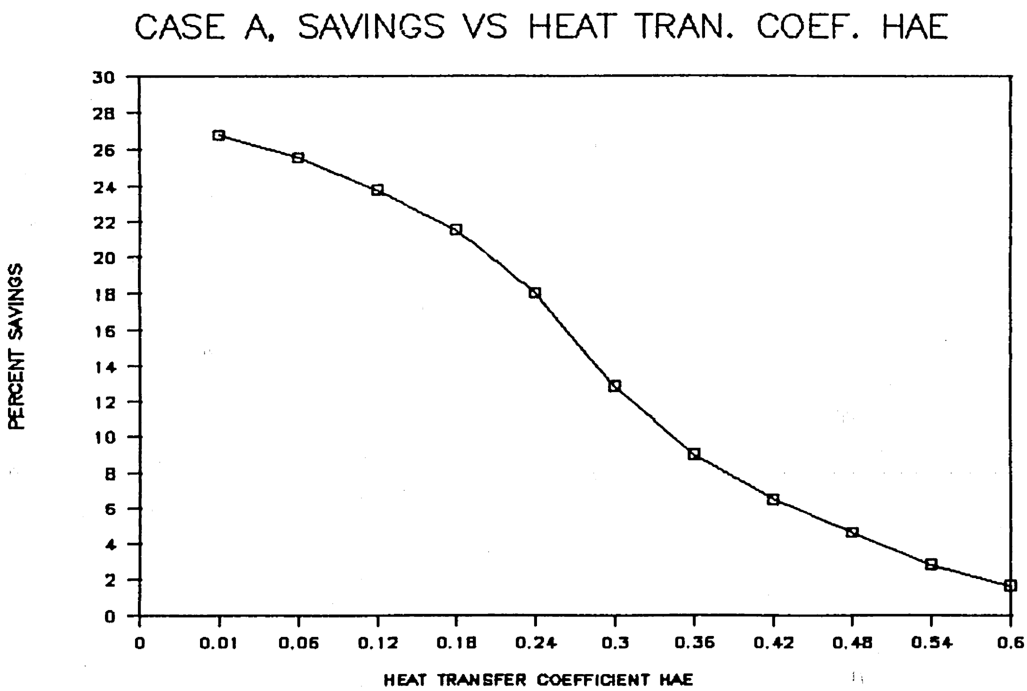

Figure 5.20: Case A, Savings vs Heat Transfer Coefficient h_{ae}

As the heat transfer rate between the outside and the inside increases the savings rate decreases. This result illustrates the impact of insulation on the savings. Figure 5.20 also indicates that even with near perfect insulation, savings account for only 25% of the reference total cost. However, it must be noted that the reference total cost is also reduced with better insulation.

Savings vs Constant Outside Temperature

If all the other parameters are kept at their base values, the rate of savings is increased if the outside temperature is closer to the minimum acceptable inside temperature. This fact is shown in Figure 5.21.

However, it must be noted that the total initial and final cost is lower when the outside temperature is closer to T_{min}. Also, the result only concerns the optimal behavior under the base values of parameters. Care must be taken in generalizing the results for other system coefficient or parameter values.

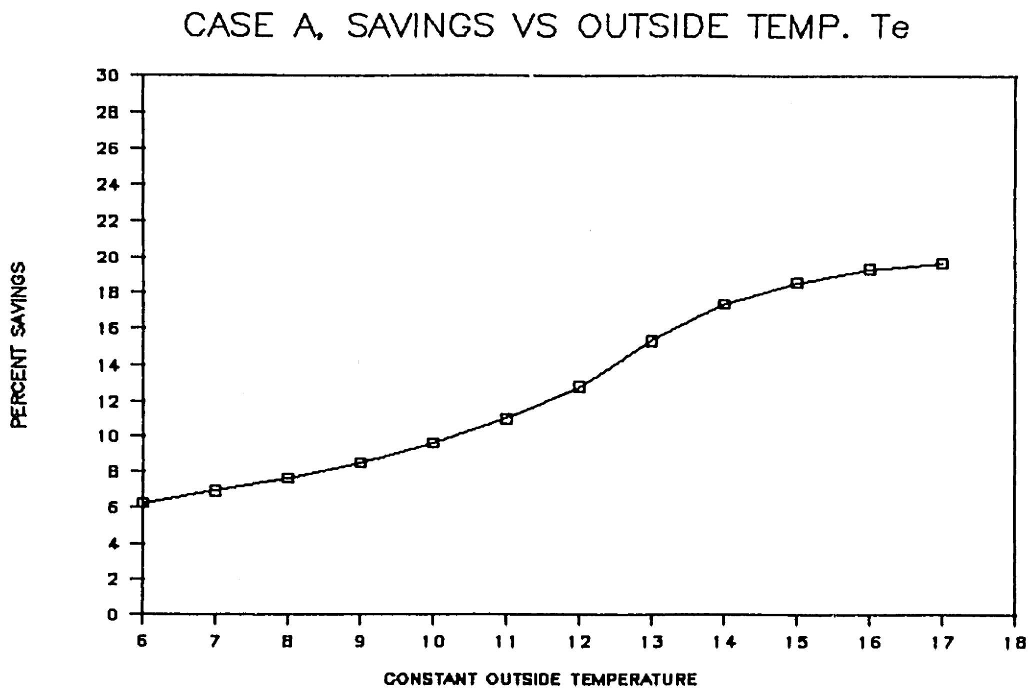

Figure 5.21: Case A, Savings vs Outside Temperature T_e

Savings vs the Maximum Heat Input U_{max}

As the bound on the maximum heat input is increased, it is expected that the rate of savings will also increase. Figure 5.22 verifies that expectation.

One interesting result is that the rate of savings levels off as U_{max} is increased beyond a certain point. The explanation is that at those higher values, the heat input is not a constraining factor. In other words, the heating system never gets to be utilized at its full capacity. Doing so, would result in violation of the maximum temperature constraint.

Savings vs the Lower Temperature Bound

As the minimum temperature bound T_{min} is lowered, the savings rate increases. Figure 5.23 also shows that there is some leveling off as this bound is lowered closer to the outside temperature.

At the other end, as expected, the savings rate decreases to zero when the

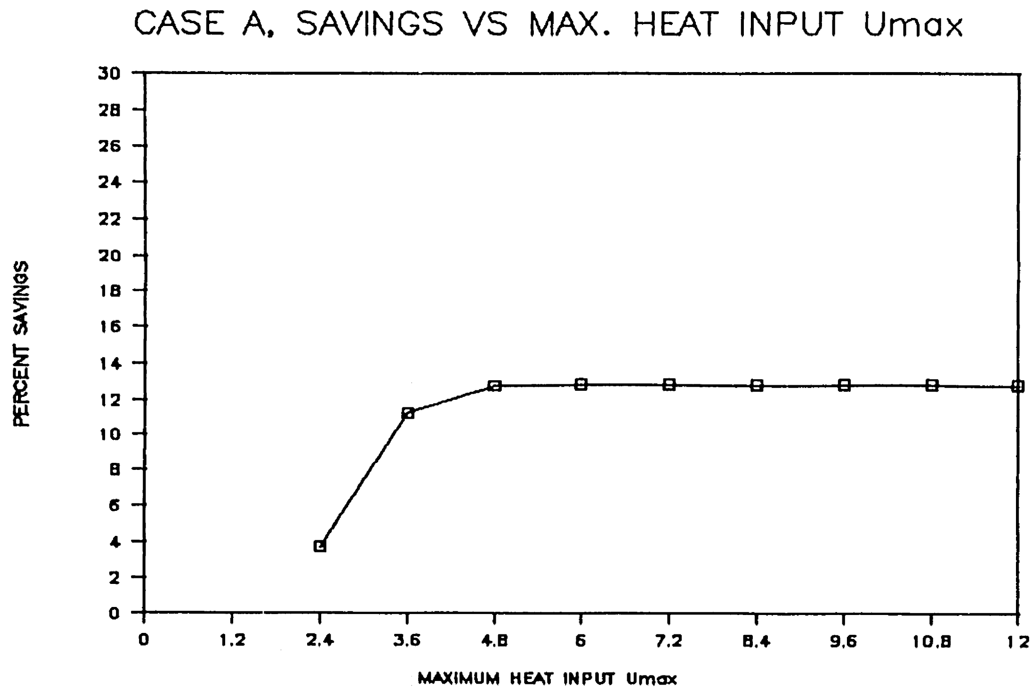

Figure 5.22: Case A, Savings vs Maximum Heat Input U_{max}

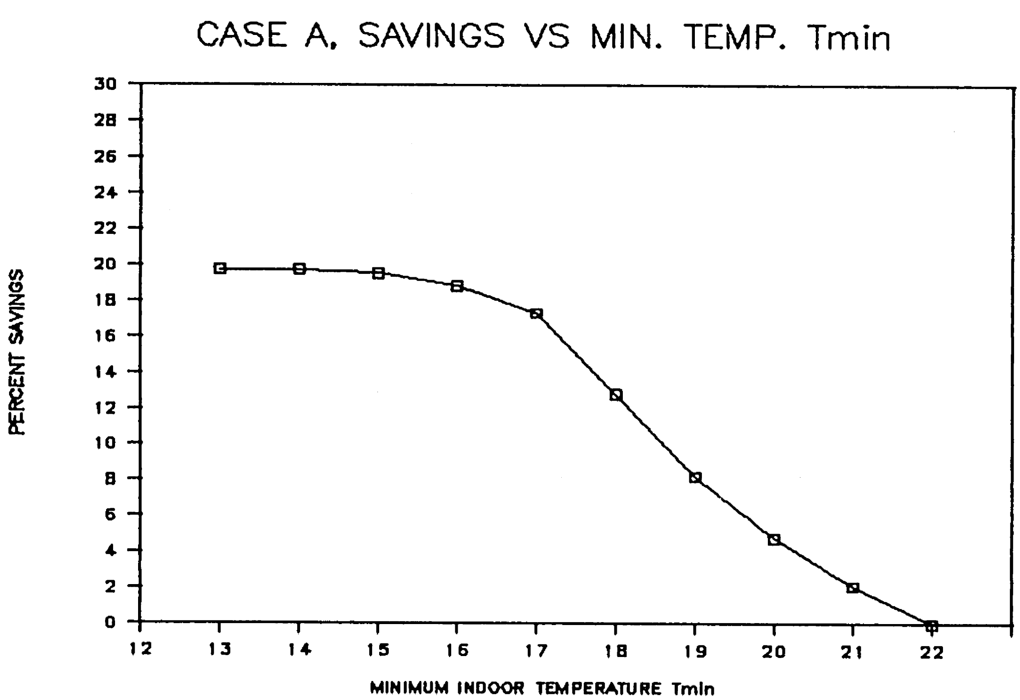

Figure 5.23: Case A, Savings vs the Lower Temperature Bound

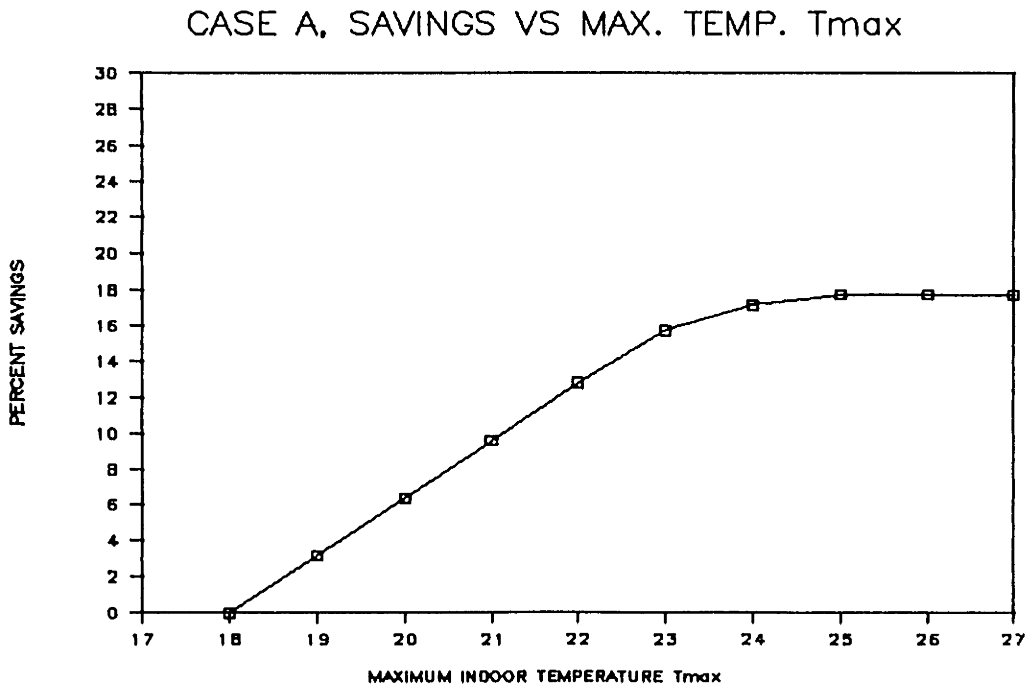

Figure 5.24: Case A, Savings vs the Upper Temperature Bound

comfortable temperature range is decreased to zero.

Savings vs the Upper Temperature Bound

Figure 5.24 shows the effect of increasing T_{max} on the savings rate. As expected, the saving rate increases as the upper temperature bound is increased. However, beyond a certain level, the savings rate stays constant. Again, the reason is that at those levels, the upper temperature bound is not a constraining factor, and in fact, in optimal operation, the inside air temperature will not reach the upper bound. This is true for the given level of the maximum heat input. Higher rates of savings will be possible if the constraint on the maximum heat input, is relaxed.

5.4.2 Case study B: Step price pattern

This study is concerned with the optimal response to step type price pattern. This is a simplified version of the time of use rates. Here, only a single step of price

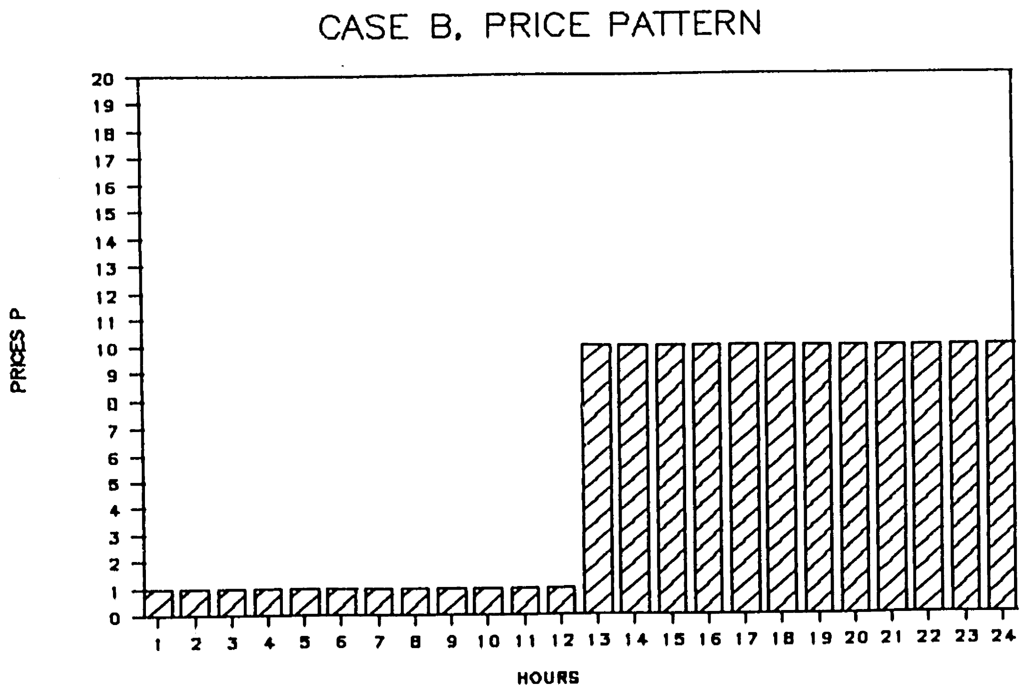

Figure 5.25: Case B, Step Price Pattern

changes is considered. An example of step price pattern in shown in Figure 5.25.

Figures 5.26 and 5.27 show the results for the optimal solution. The step change happens at the hour 13. Thus, the heat input is scheduled to be higher than its reference values at earlier hours in order to produce sufficient stored heat to substitute for the heat inputs of the later half of the time horizon.

Again, as in case study A, the heat input levels at hours after 13 are also decreased below their reference levels. This is due to the slow dissipation of the stored heat that is carried away beyond the hour 13. Thus, it is possible to decrease heat input of some future hours and allow the inside air temperature down to its minimum level. In the case step price pattern, this results in substantial savings.

Savings vs Size of the Price Range

Savings increase with the size of the price change. Figure 5.28 demonstrates this fact. The distance between the high and low prices was varied between -100%

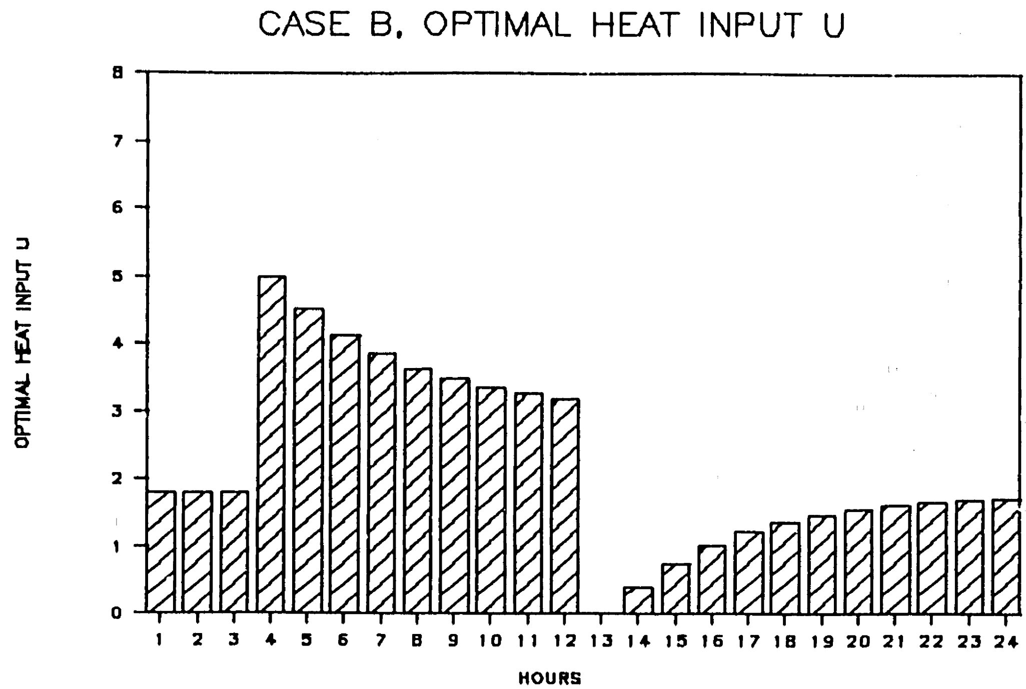

Figure 5.26: Case B, Optimal Heat Input U

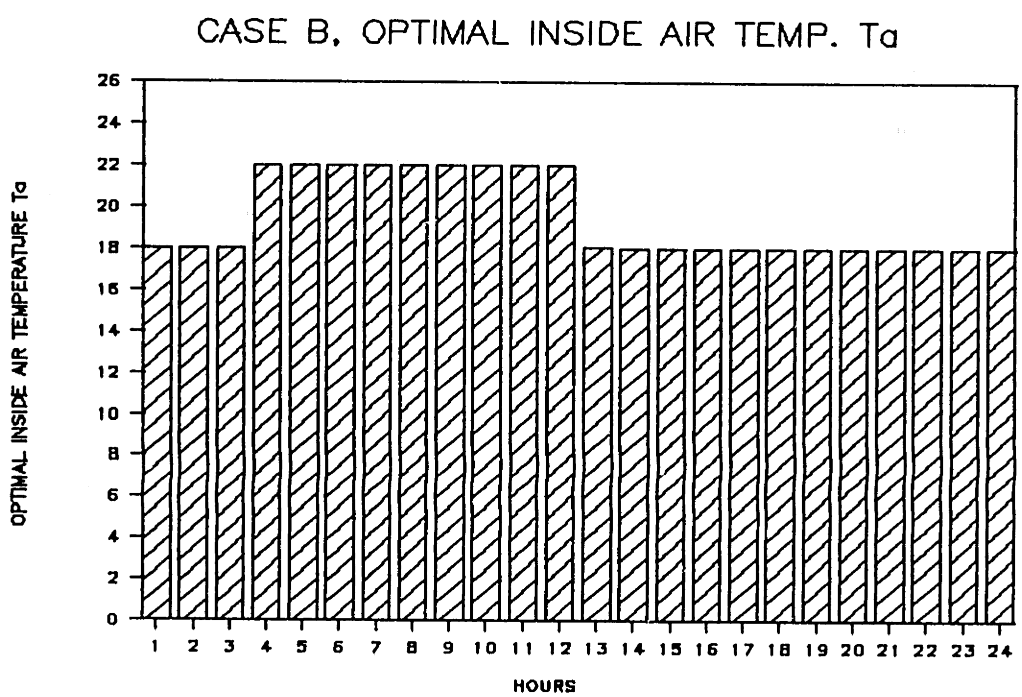

Figure 5.27: Case B, Optimal Inside Air Temperature T_a

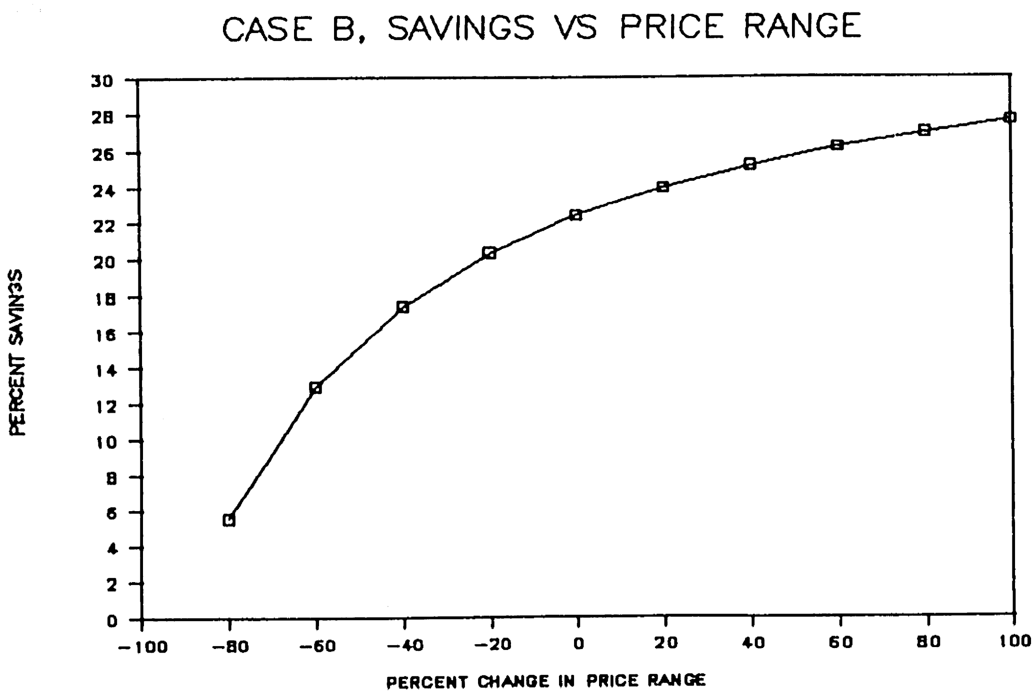

Figure 5.28: Case B, Savings vs Size of the Price Range

and +100%. In contrast to the case study A, savings are possible even when the ratio of the price jump is as low as 2.

Savings vs Thermal Capacity C_i

Figure 5.29 depicts the variations in the savings rate with the variation in the thermal capacity. In contrast to the case of impulse price pattern, savings rate increases with the increase in the size of the thermal capacity from its base value. This is due to the fact that the thermal storage must carry enough energy in order to substitute the heat inputs of many hours. Again, a maximum is achieved, but at a higher value of the thermal capacity compared to that of case study A.

Savings vs Heat Transfer Coefficient h_{ai}

Figure 5.30 shows the variation of the savings rate with the variation of the inside air-storage heat transfer coefficient h_{ai}.

As expected, savings improve as the heat transfer rate between the inside air

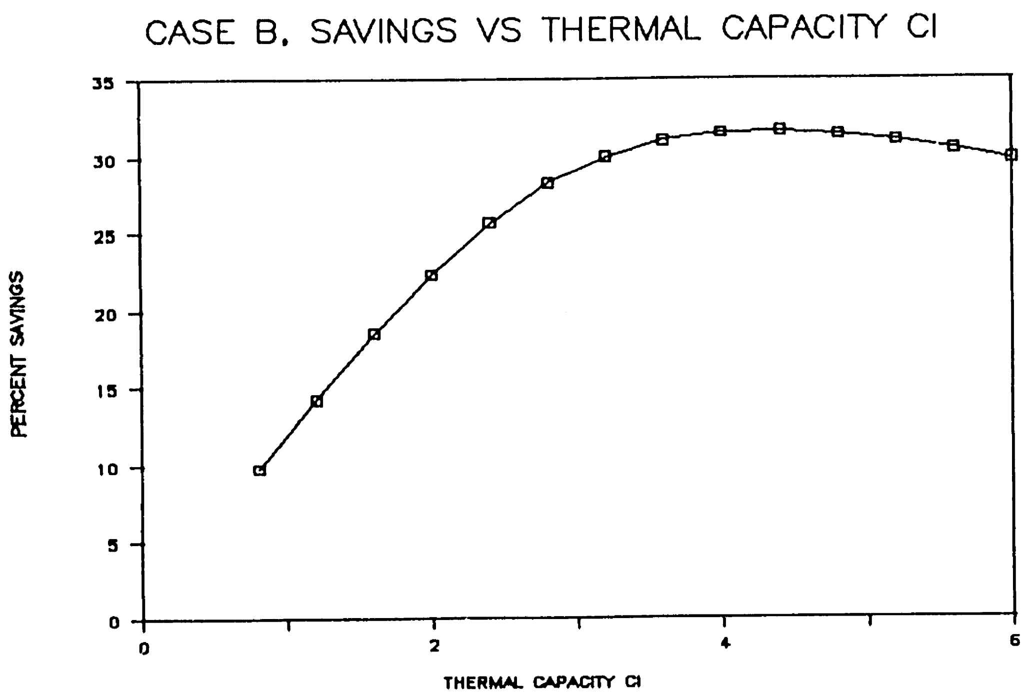

Figure 5.29: Case B, Savings vs Thermal Capacity C_i

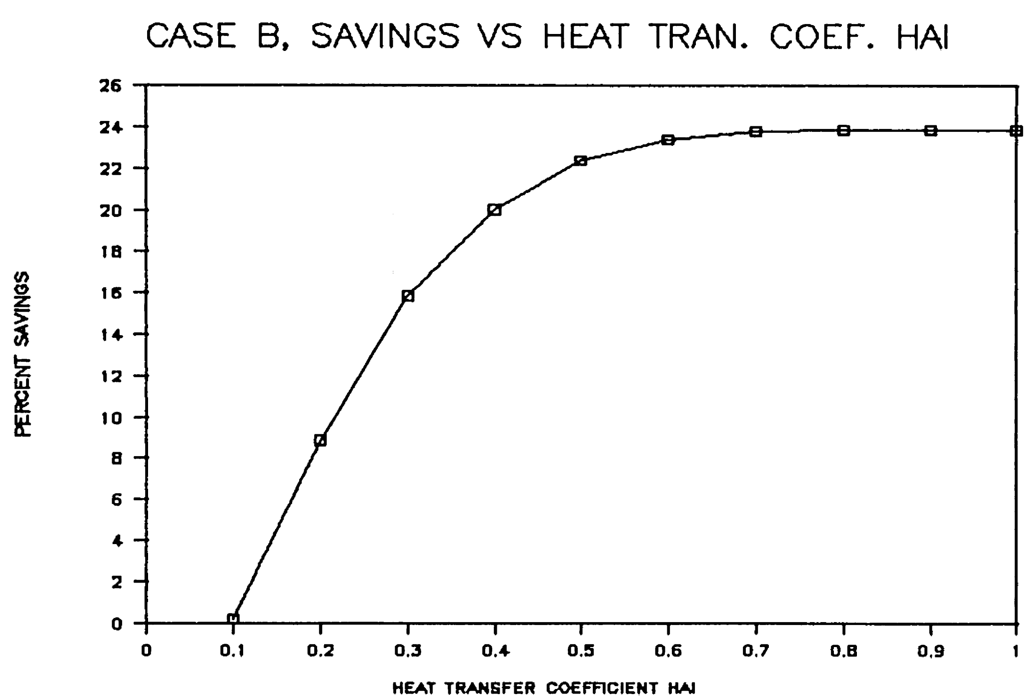

Figure 5.30: Case B, Savings vs Heat Transfer Coefficient h_{ai}

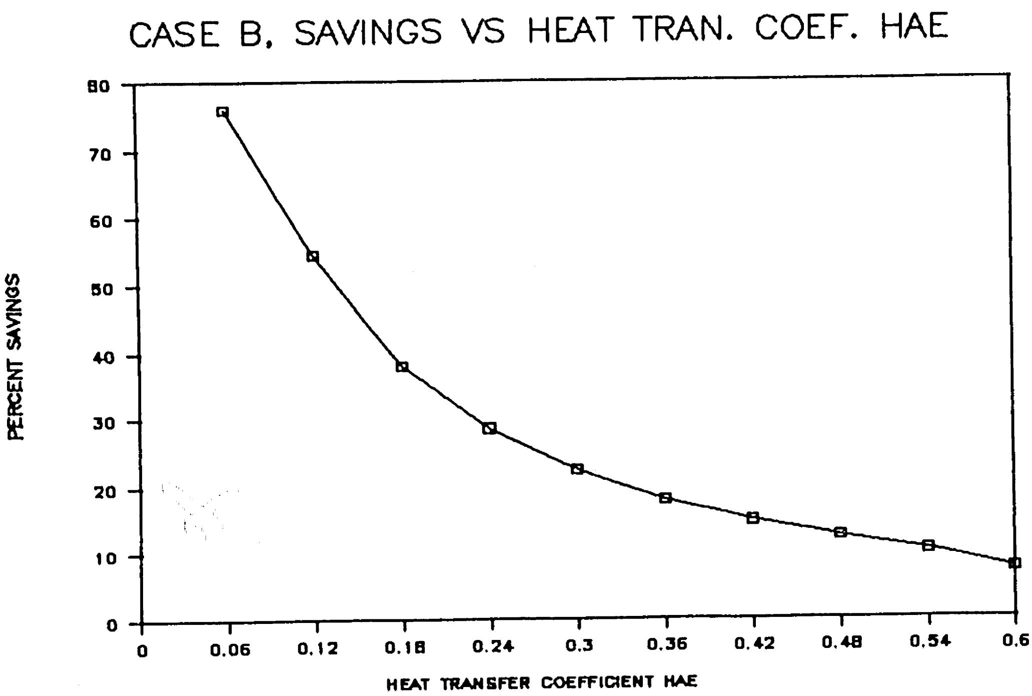

Figure 5.31: Case B, Savings vs Heat Transfer Coefficient h_{ae}

and the storage medium is increased.

Savings vs Heat Transfer Coefficient h_{ae}

Figure 5.31 shows the variation in the savings rate with respect to the variation in the heat transfer rate between the inside air and the ambience.

As the heat transfer rate between the outside and the inside increases the savings rate decreases. As before, this result illustrates the impact of insulation on the savings. Figure 5.31 also indicates that savings rate are substantially higher for this case compared to case A for near perfect insulation. It goes as high as 70%.

Savings vs Constant Outside Temperature

If all the other parameters are kept at their base values, the rate of savings is increased if the outside temperature is closer to the minimum acceptable inside temperature. This is shown in Figure 5.32.

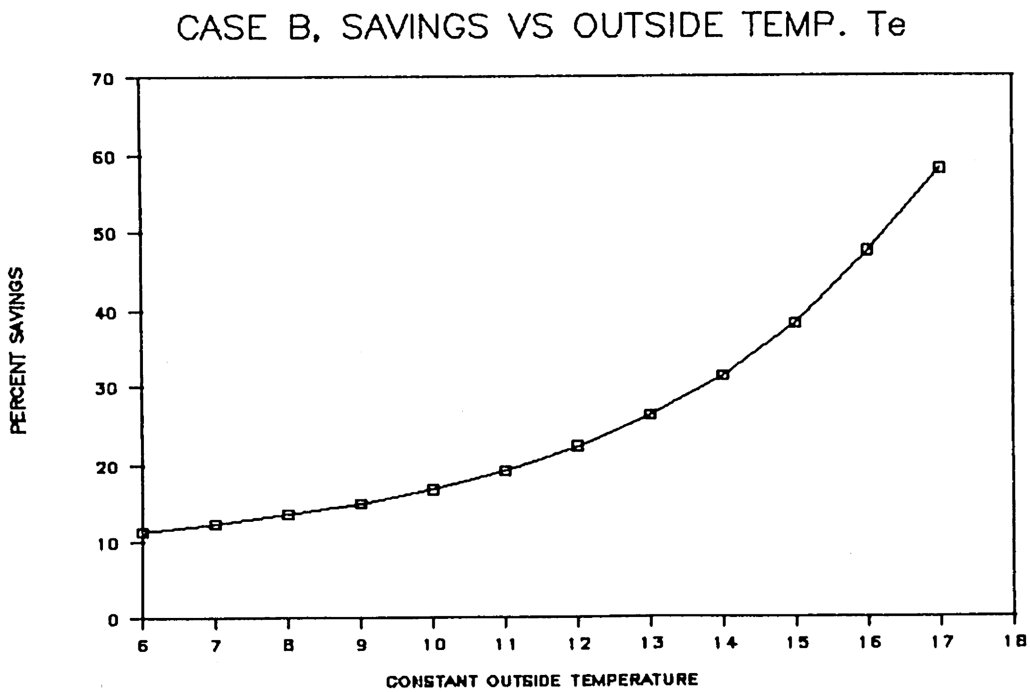

Figure 5.32: Case B, Savings vs Outside Temperature T_e

However, it must be noted that the total initial and final cost is lower when the outside temperature is closer to T_{min}. Also, the result only concerns the optimal behavior under the base values of parameters.

Savings vs the Maximum Heat Input U_{max}

As the bound on the maximum heat input is increased, it is expected that the rate of savings will also increase. Figure 5.22 verifies that expectation.

Again, the rate of savings levels off as U_{max} is increased beyond a certain point. The explanation is that at those higher values, the heat input is not a constraining factor. In other words, the heating system never gets to be utilized at its full capacity. Doing so, would result in violation of the maximum temperature constraint.

Savings vs the Lower Temperature Bound

As the minimum temperature bound T_{min} is lowered, the savings rate increases.

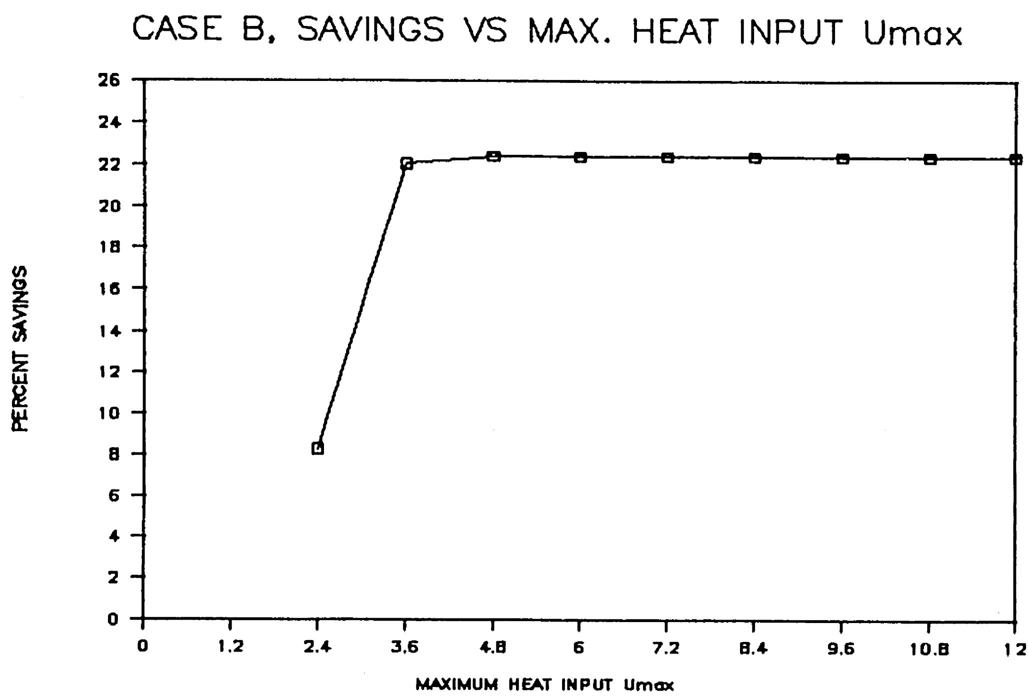

Figure 5.33: Case B, Savings vs Maximum Heat Input U_{max}

Figure 5.34 also shows that when there is some leveling off as this bound is lowered closer to the outside temperature.

At the other end, as expected, the savings rate decreases to zero when the comfortable temperature range is decreased to zero.

Savings vs the Upper Temperature Bound

Figure 5.35 shows the effect of increasing T_{max} on the savings rate.

As expected, the saving rate increases as the upper temperature bound is increased. However, beyond a certain level, the savings rate stays constant. Again, the reason is that at those levels, the upper temperature bound is not a constraining factor, and in fact, in optimal operation, the inside air temperature will not reach the upper bound. In such a case, higher rates of savings will be possible if the maximum heat input is increased.

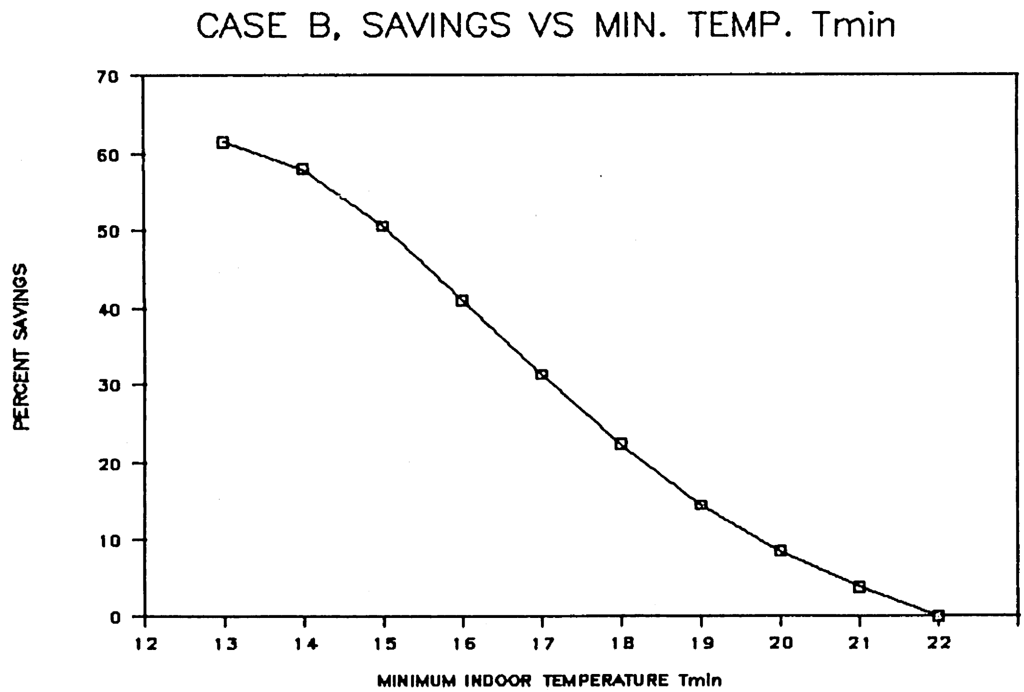

Figure 5.34: Case B, Savings vs the Lower Temperature Bound

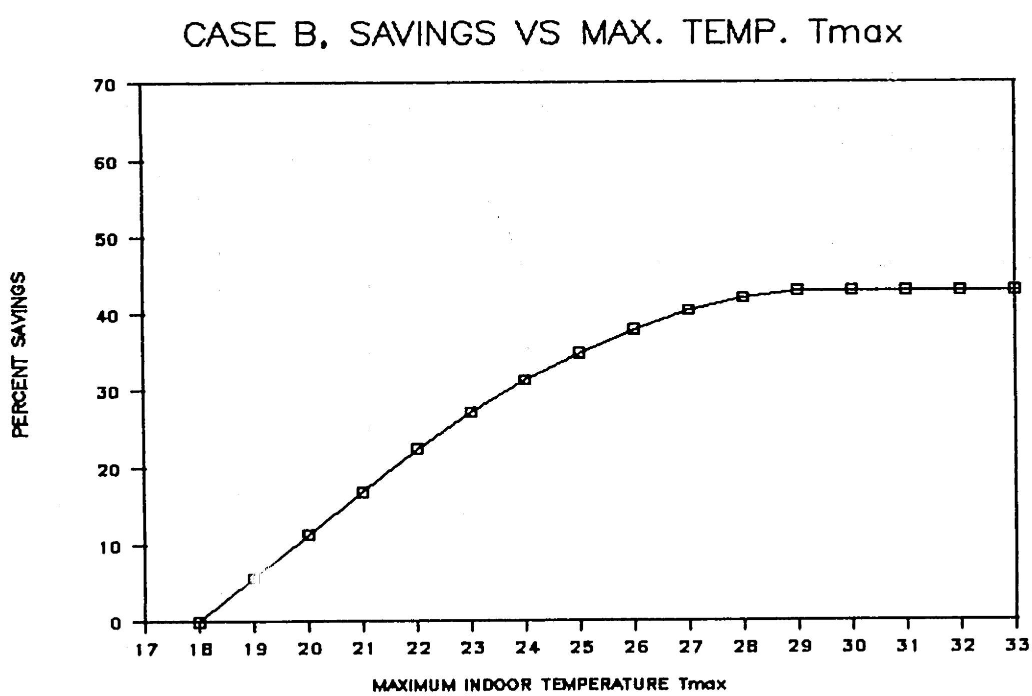

Figure 5.35: Case B, Savings vs the Upper Temperature Bound

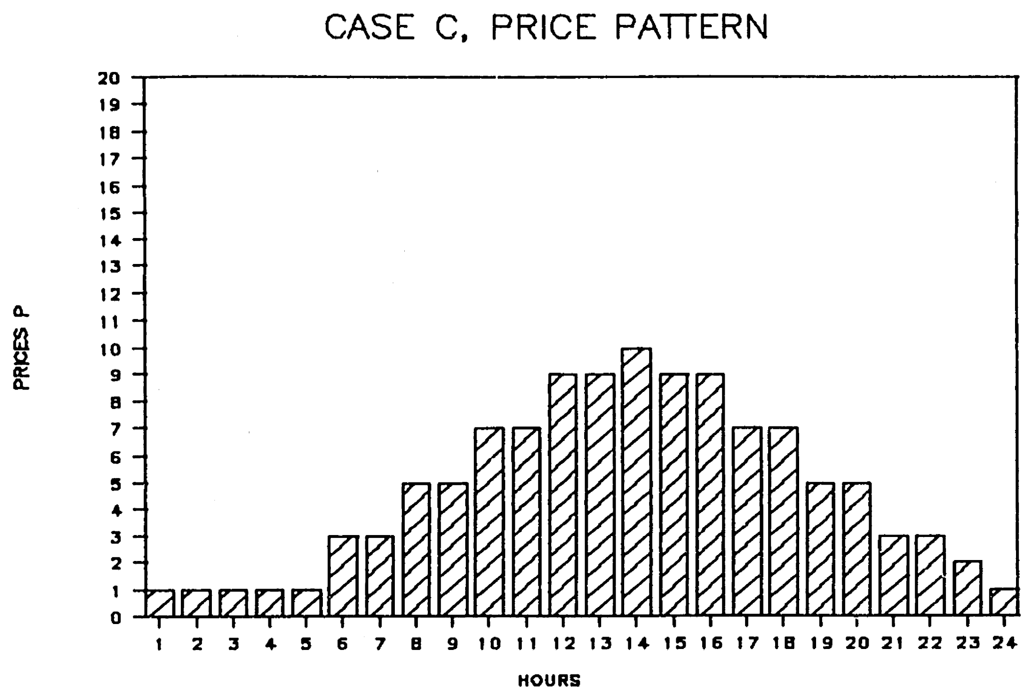

Figure 5.36: Case C, Single Peak Price Pattern

5.4.3 Case study C: single peak price pattern

The study is concerned with the optimal response to a single peak type price Pattern. An example is shown in Figure 5.36.

Figures 5.37 and 5.38 show the results for the optimal solution.

Savings vs Size of the Price Range

Savings increase with the size of the price changes. Figure 5.39 demonstrates this fact.

Savings vs Thermal Capacity C_i

Figure 5.40 depicts the variations in the savings rate with the variation in the thermal capacity. Again, it is shown that there is a unique optimal design value for the size of the thermal mass.

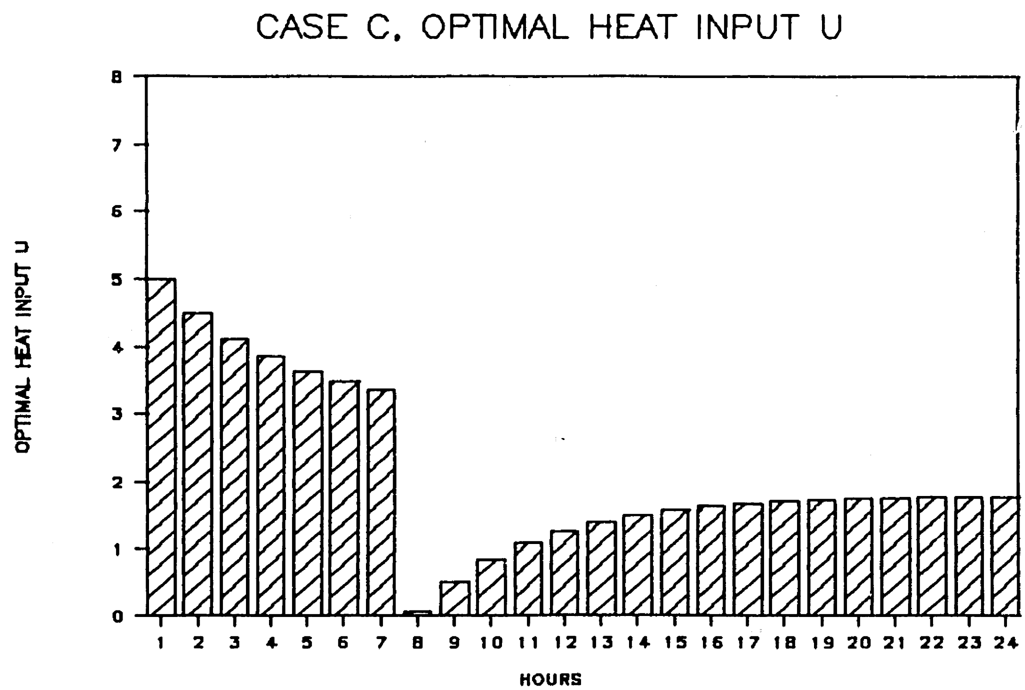

Figure 5.37: Case C, Optimal Heat Input U

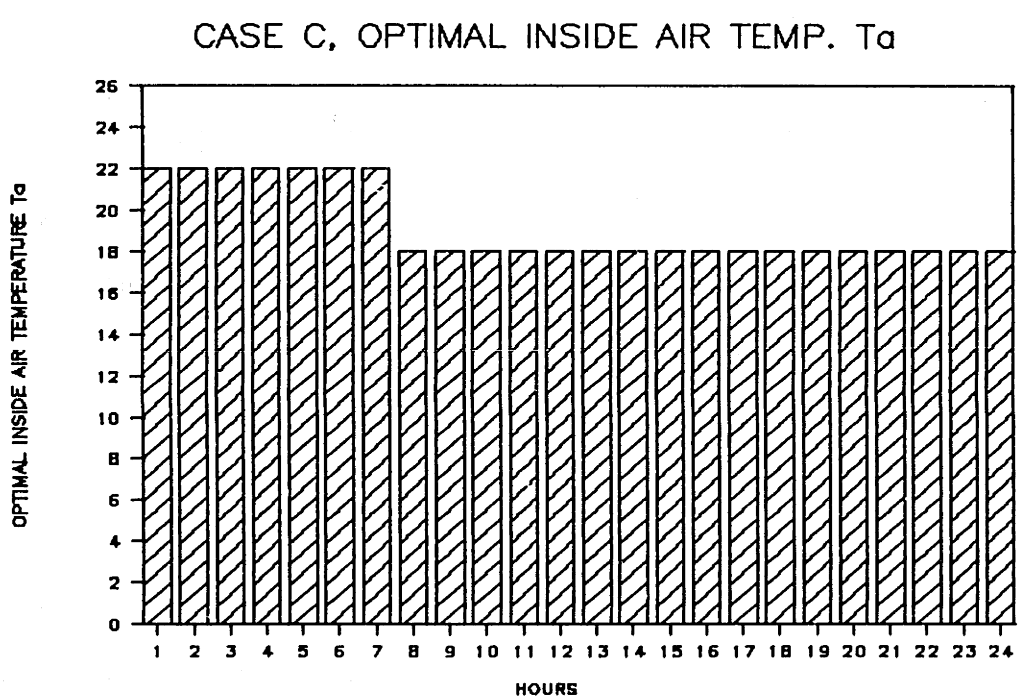

Figure 5.38: Case C, Optimal Inside Air Temperature T_a

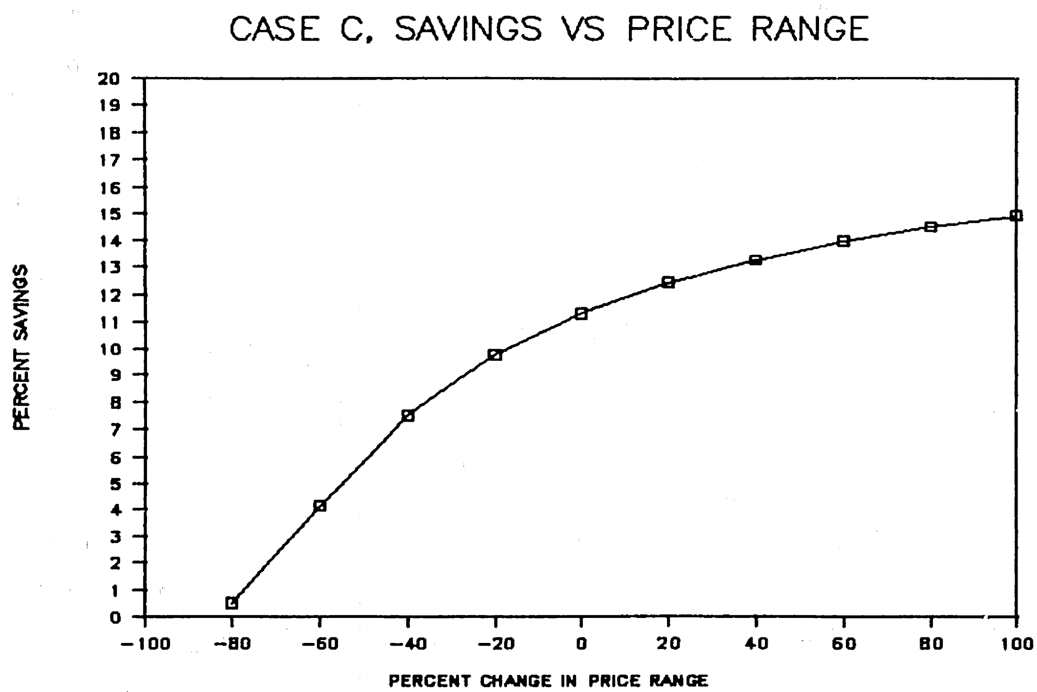

Figure 5.39: Case C, Savings vs Size of the Price Range

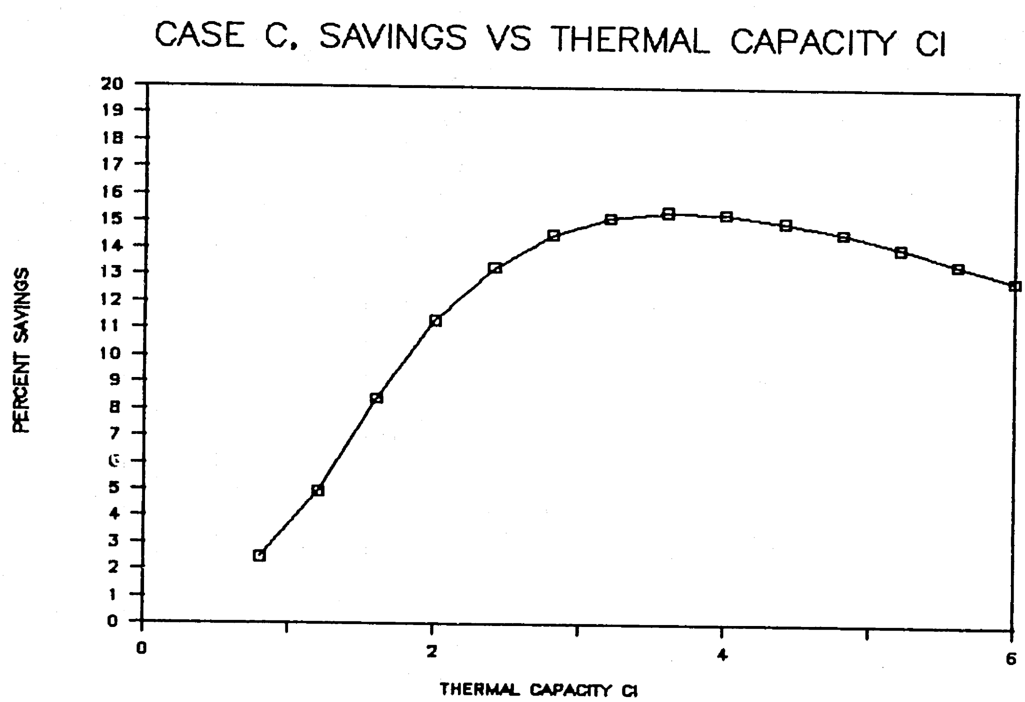

Figure 5.40: Case C, Savings vs Thermal Capacity C_i

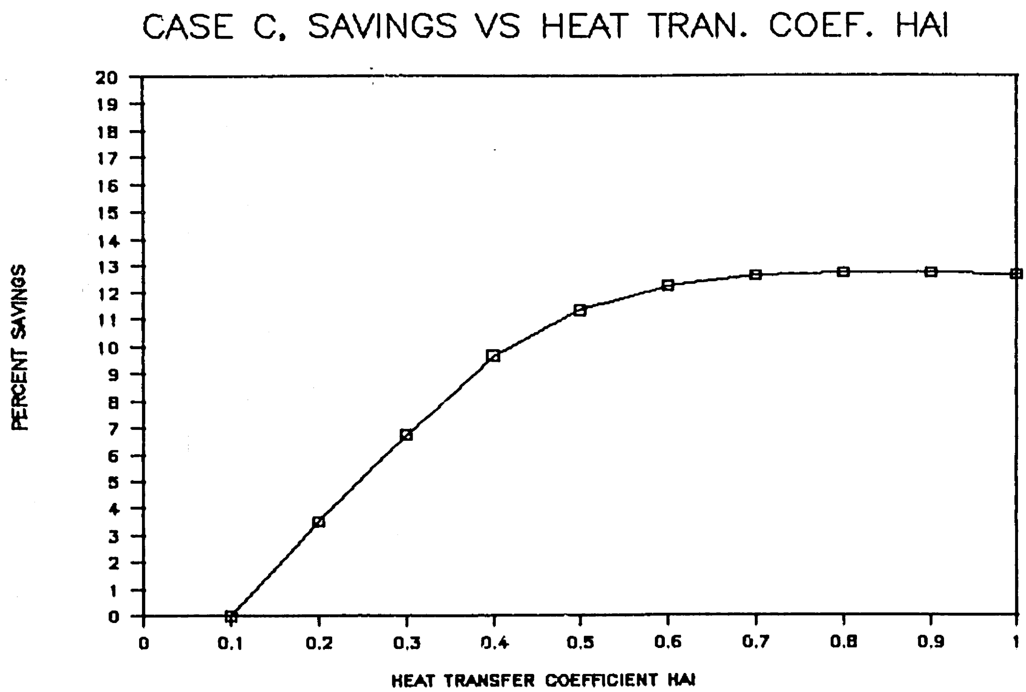

Figure 5.41: Case C, Savings vs Heat Transfer Coefficient h_{ai}

Savings vs Heat Transfer Coefficient h_{ai}

Figure 5.41 shows the variation of the savings rate with the variation of the inside air-storage heat transfer coefficient h_{ai}.

As expected, savings improve as the heat transfer rate between the inside air and the storage medium is increased.

Savings vs Heat Transfer Coefficient h_{ae}

Figure 5.42 shows the variation in the savings rate with respect to the variation in the heat transfer rate between the inside air and the ambience.

As the heat transfer rate between the outside and the inside increases, the savings rate decreases.

Savings vs Constant Outside Temperature

If all the other parameters are kept at their base values, the rate of savings is

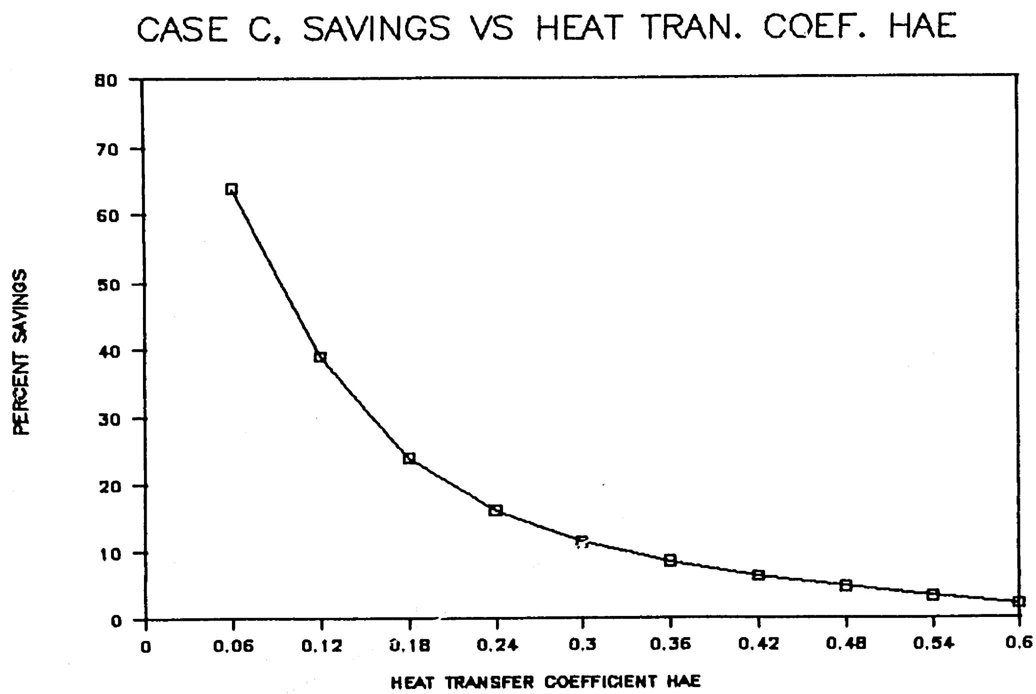

Figure 5.42: Case C, Savings vs Heat Transfer Coefficient h_{ae}

increased if the outside temperature is closer to the minimum acceptable inside temperature. This is shown in Figure 5.43.

Savings vs the Maximum Heat Input U_{max}

As the bound on the maximum heat input is increased, it is expected that the rate of savings will also increase. This is shown in Figure 5.44.

Again, the rate of savings levels off as U_{max} is increased beyond a certain point. The heating system never gets to be utilized at its full capacity because of maximum temperature constraint.

Savings vs the Lower Temperature Bound

As the minimum temperature bound T_{min} is lowered, the savings rate increases. Figure 5.45 also shows that when there is some leveling off as this bound is lowered closer to the outside temperature.

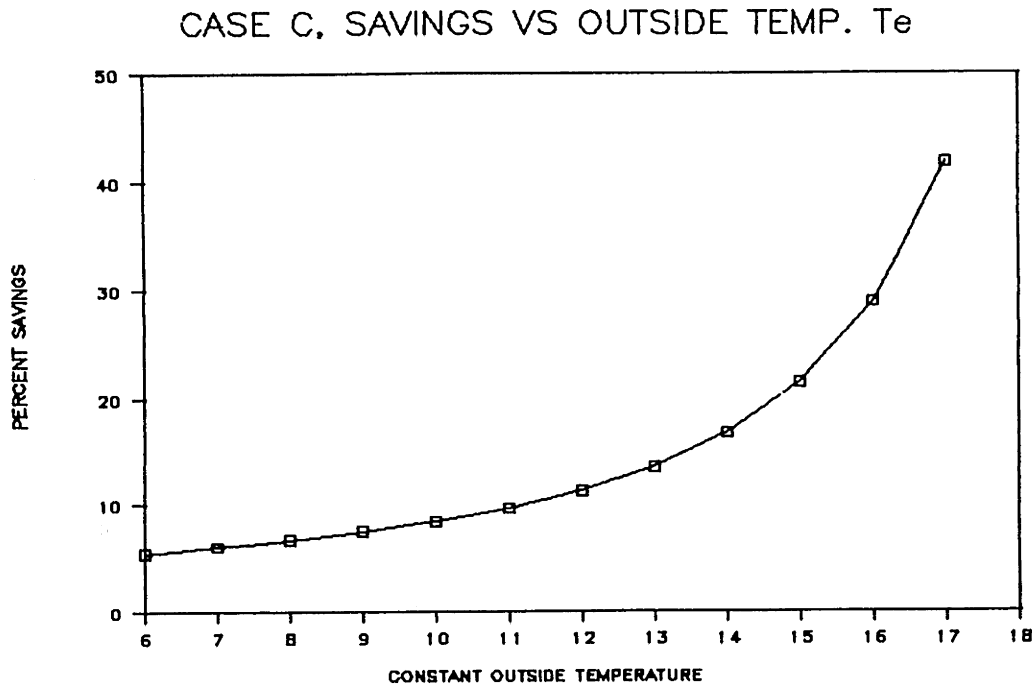

Figure 5.43: Case C, Savings vs Outside Temperature T_e

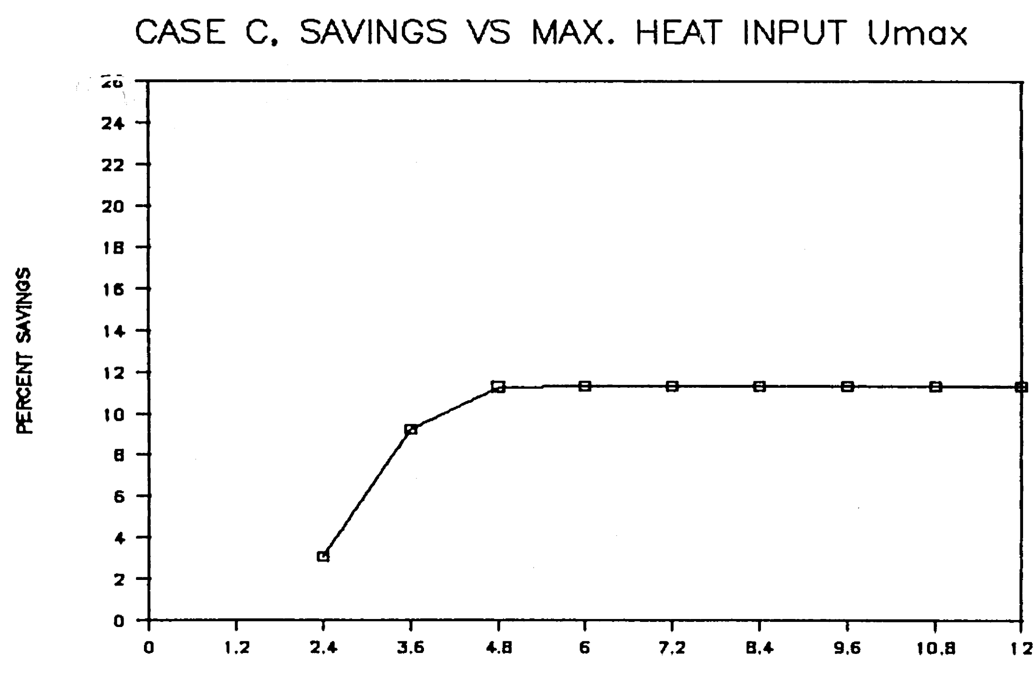

Figure 5.44: Case C, Savings vs Maximum Heat Input U_{max}

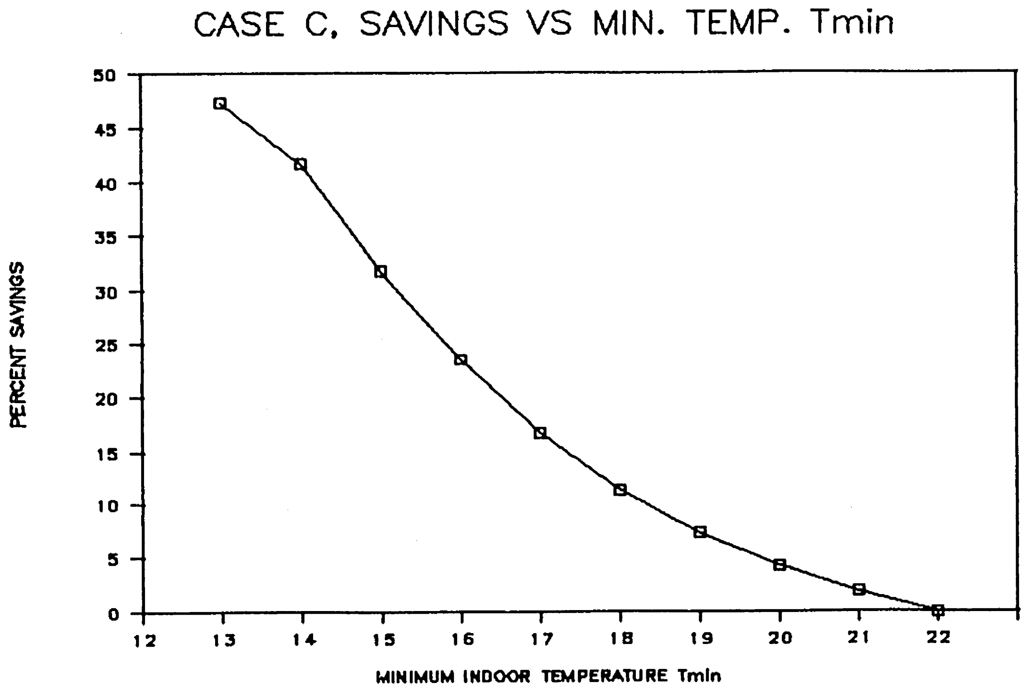

Figure 5.45: Case C, Savings vs the Lower Temperature Bound

At the other end, as expected, the savings rate decreases to zero when the comfortable temperature range is decreased to zero.

Savings vs the Upper Temperature Bound

Figure 5.46 shows the effect of increasing T_{max} on the savings rate.

As expected, the saving rate increases as the upper temperature bound is increased.

5.4.4 Case study D: double peak price pattern

The study is concerned with the optimal response to double peak type price pattern. An example is shown in Figure 5.47.

Figures 5.48 and 5.49 show the results for the optimal solution.

Savings vs Size of the Price Range

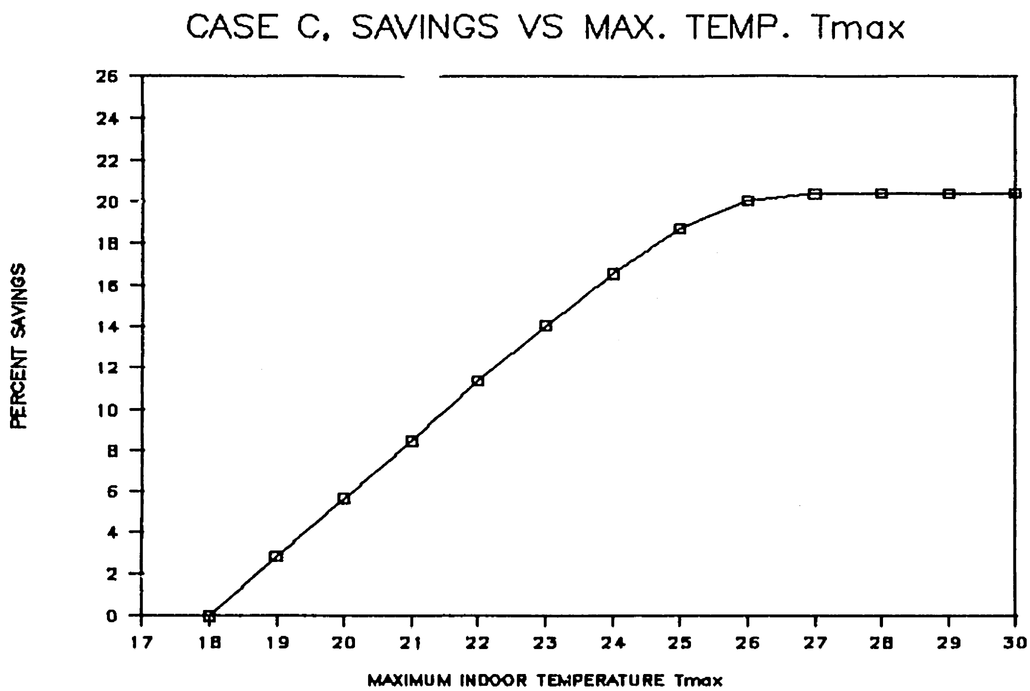

Figure 5.46: Case C, Savings vs the Upper Temperature Bound

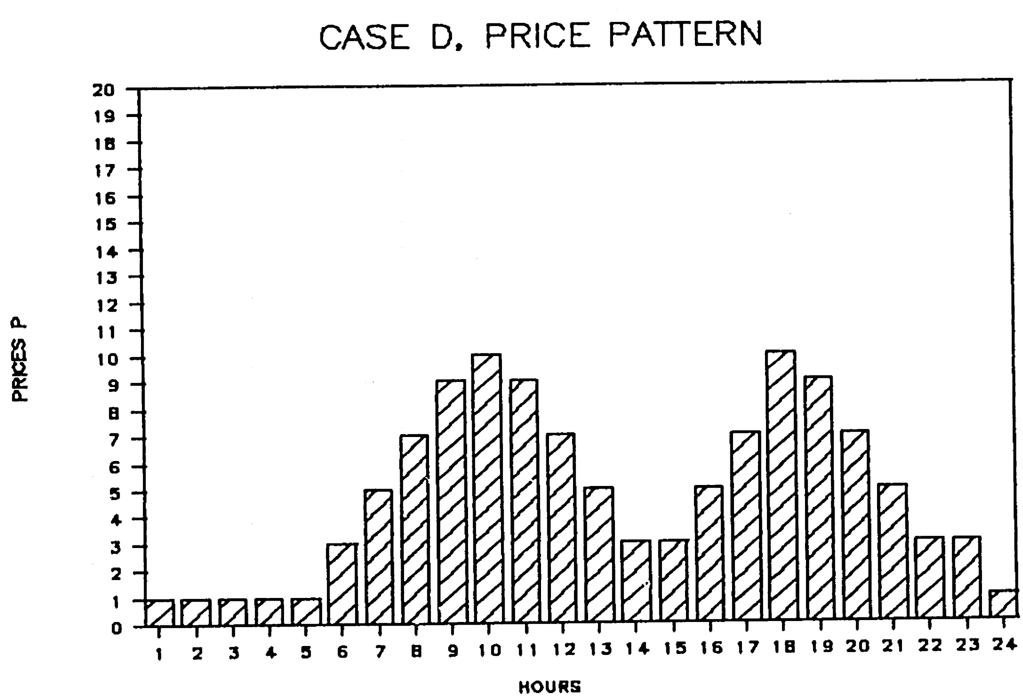

Figure 5.47: Case D, Double Peak Price Pattern

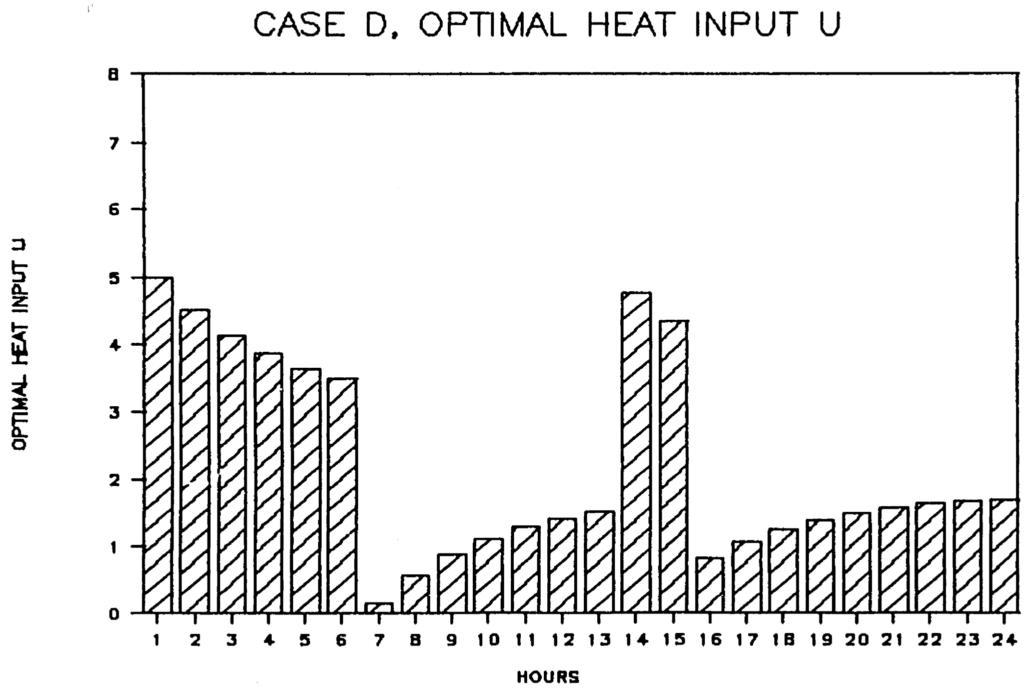

Figure 5.48: Case D, Optimal Heat Input U

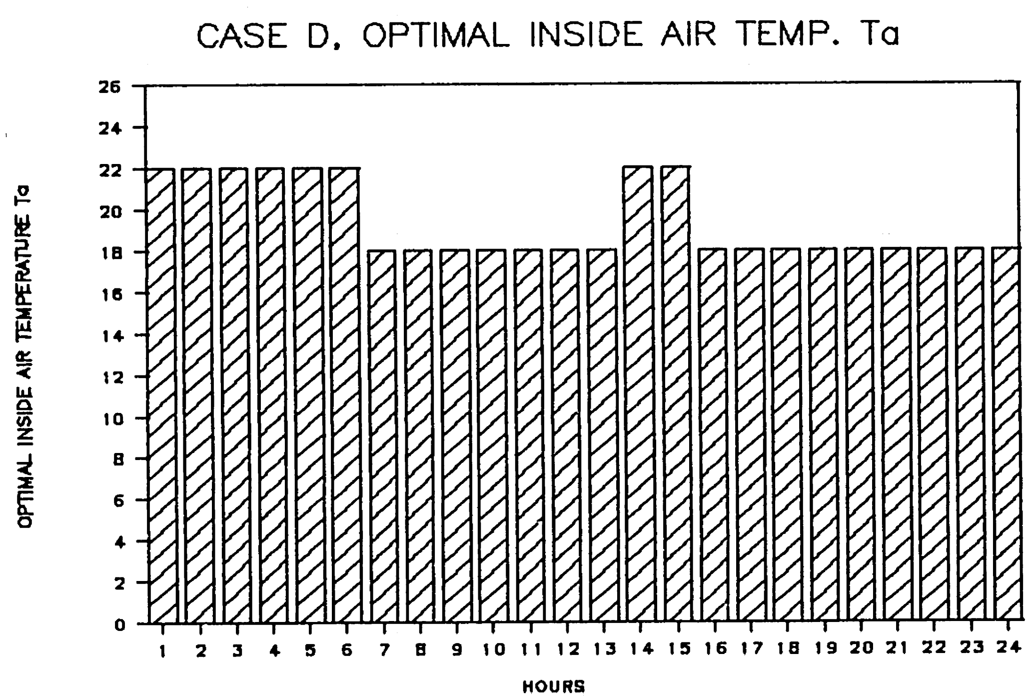

Figure 5.49: Case D, Optimal Inside Air Temperature T_a

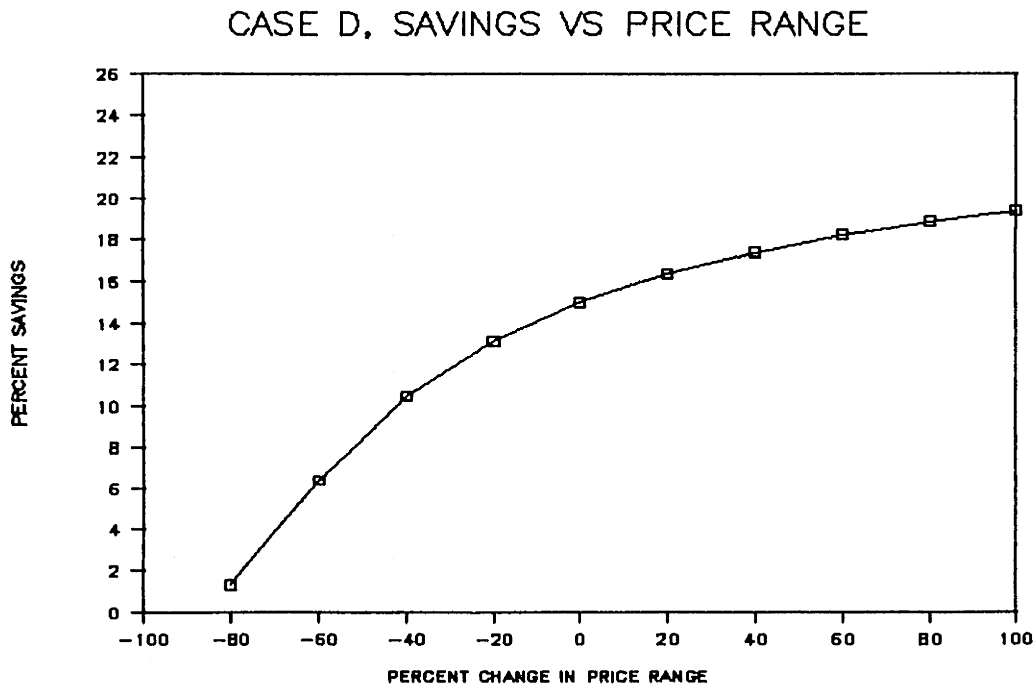

Figure 5.50: Case D, Savings vs Size of Price Range

Savings increase with the size of the price levels. This is shown in Figure 5.50.

Savings vs Thermal Capacity C_i

Figure 5.51 depicts the variations in the savings rate with the variation in the thermal capacity.

Savings vs Heat Transfer Coefficient h_{ai}

Figure 5.52 shows the variation of the savings rate with the variation of the inside air-storage heat transfer coefficient h_{ai}.

As expected, savings improve as the heat transfer rate between the inside air and the storage medium is increased.

Savings vs Heat Transfer Coefficient $h_{a\epsilon}$

Figure 5.53 shows the variation in the savings rate with respect to the variation in the heat transfer rate between the inside air and the ambience.

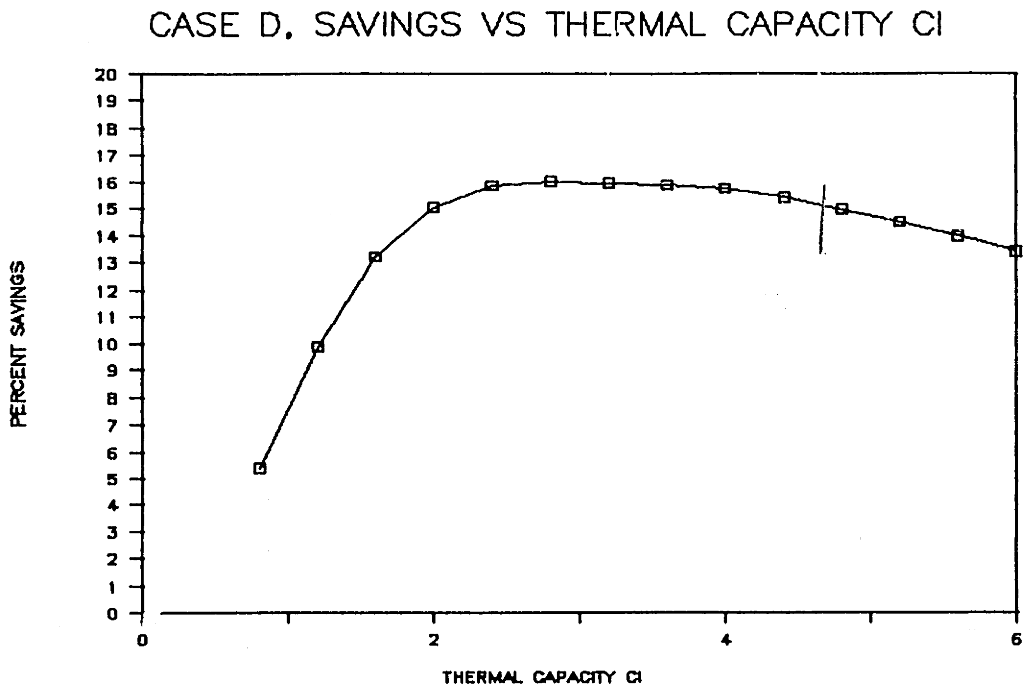

Figure 5.51: Case D, Savings vs Thermal Capacity C_i

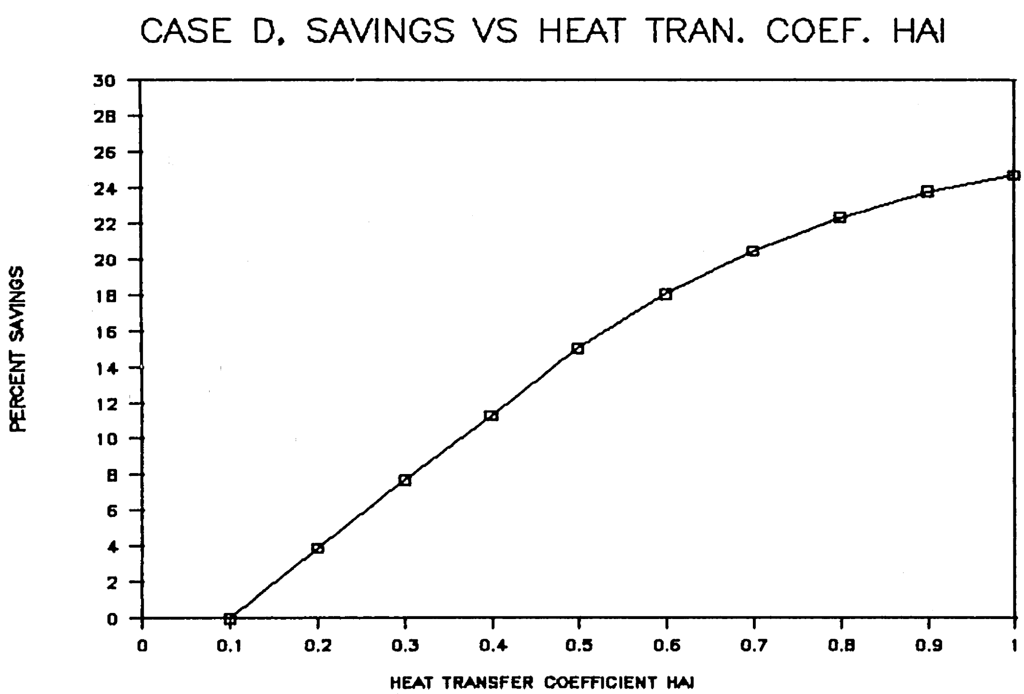

Figure 5.52: Case D, Savings vs Heat Transfer Coefficient h_{ai}

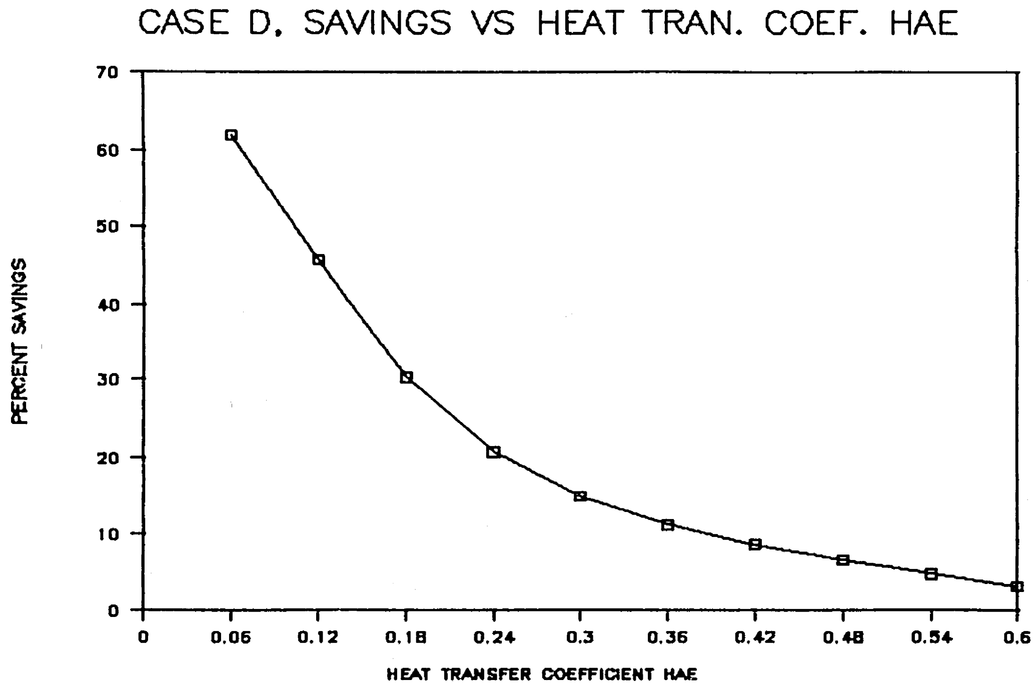

Figure 5.53: Case D, Savings vs Heat Transfer Coefficient h_{ae}

As the heat transfer rate between the outside and the inside increases the savings rate decreases.

Savings vs Constant Outside Temperature

If all the other parameters are kept at their base values, the rate of savings is increased if the outside temperature is closer to the minimum acceptable inside temperature. This fact is shown in Figure 5.54.

Savings vs the Maximum Heat Input U_{max}

As the bound on the maximum heat input is increased, it is expected that the rate of savings will also increase. This is shown in Figure 5.55.

Again, the behavior is similar to the previous cases.

Savings vs the Lower Temperature Bound

As before, as the minimum temperature bound T_{min} is lowered, the savings

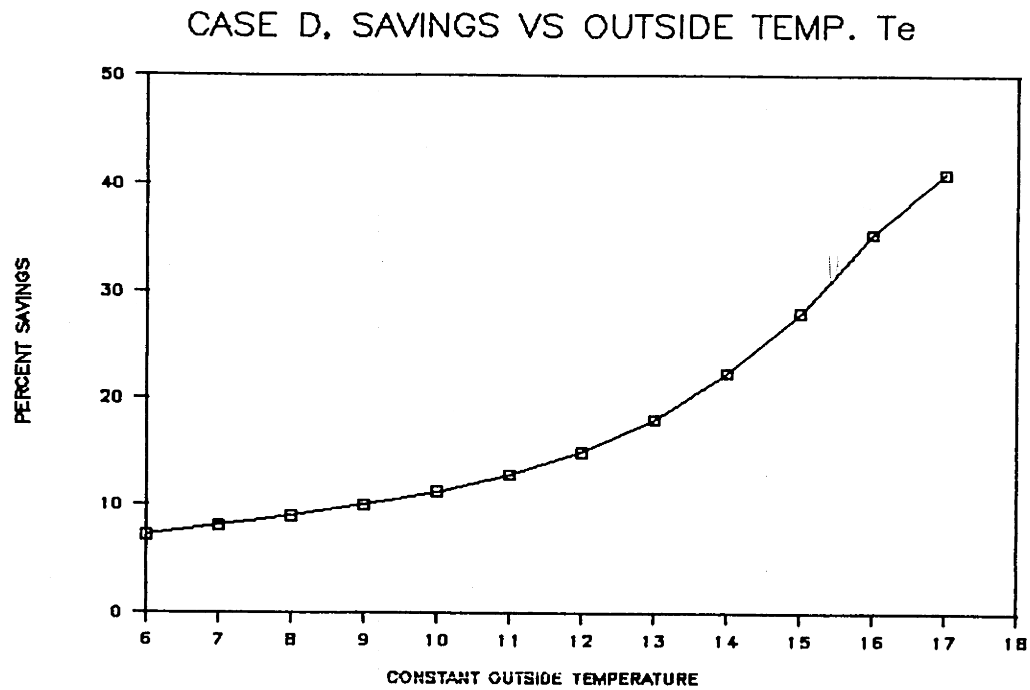

Figure 5.54: Case D, Savings vs Outside Temperature T_e

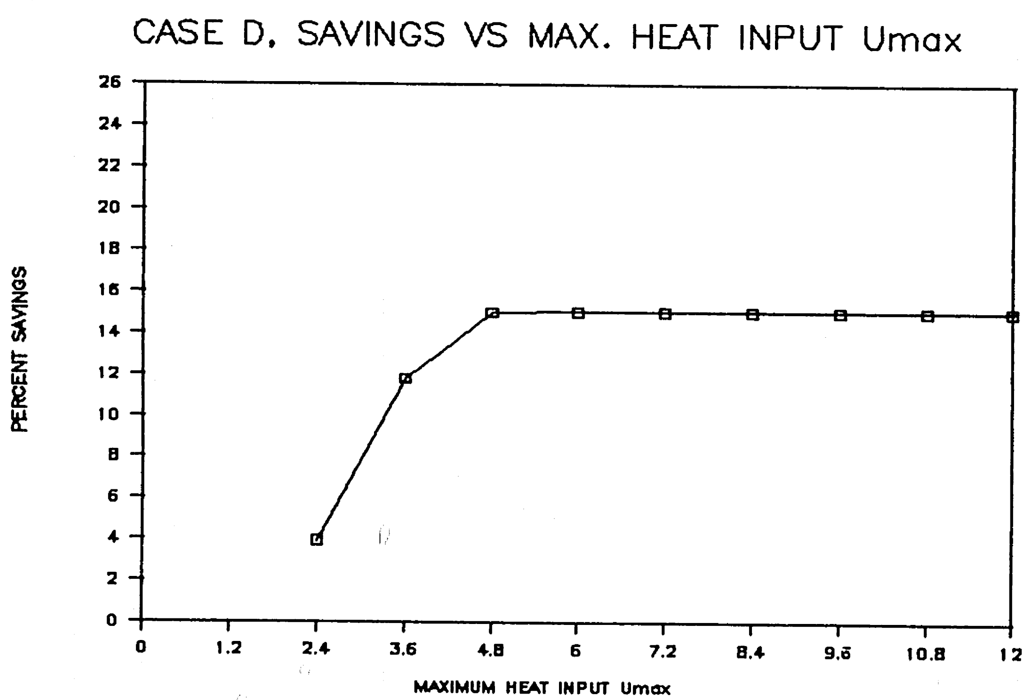

Figure 5.55: Case D, Savings vs Maximum Heat Input U_{max}

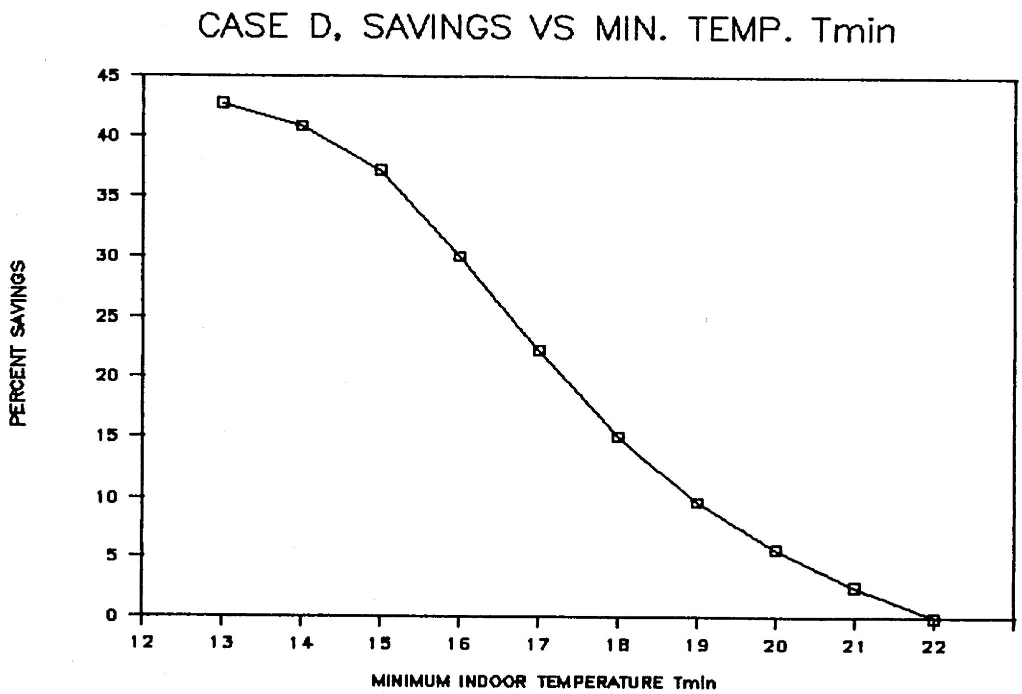

Figure 5.56: Case D, Savings vs the Lower Temperature Bound

rate increases. Figure 5.56 also shows that when there is some leveling off as this bound is lowered closer to the outside temperature.

At the other end, as expected, the savings rate decreases to zero when the comfortable temperature range is decreased to zero.

<u>Savings vs the Upper Temperature Bound</u>

Figure 5.57 shows the effect of increasing T_{max} on the savings rate. The result is similar to the previous cases studied.

5.5 Observations

Results clearly indicate that the savings depend on:

- price patterns and price variations

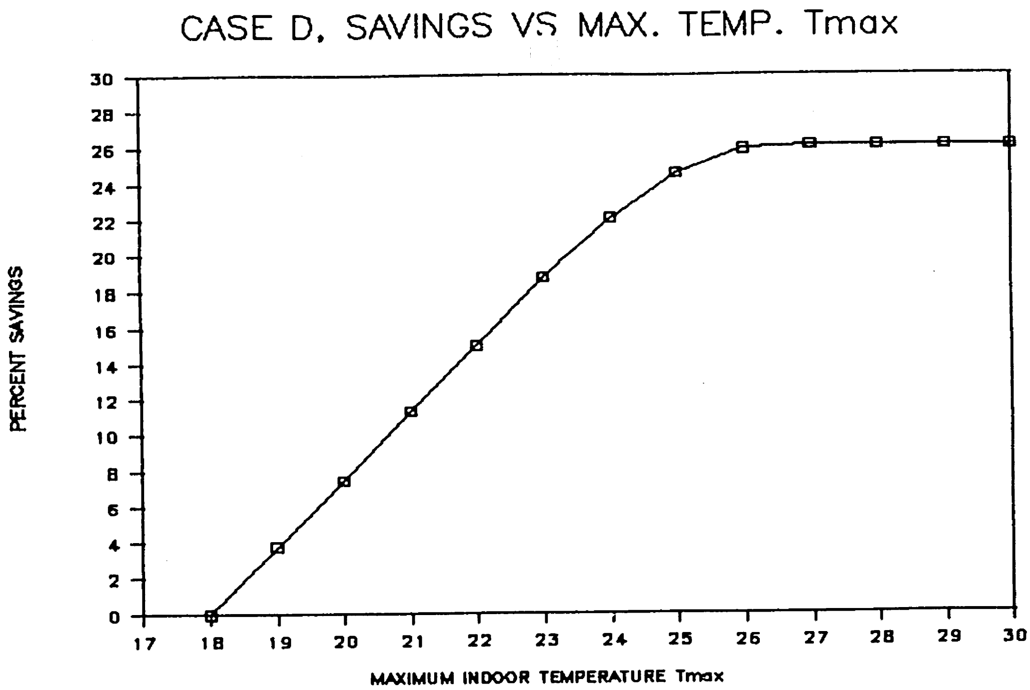

Figure 5.57: Case D, Savings vs the Upper Temperature Bound

- building thermal characteristics
- outside temperature
- heat input capacity
- temperature comfort range

However, the behavior observed under the various price patterns depend on the particular values of the parameters used. For instance, in our study, the heating input is never utilized at its maximum capacity. Hence, the results of the study should be interpreted as the behavior of the particular building under consideration. In addition to the choice of the parameters, another influencing factor is the type of the building being studied. A 3R2C building may behave differently compared to a 2R1C building.

The numerical results on savings must be interpreted in the light of a number of limiting factors such as the choice of a relatively long time step, assumption of the predictability of weather data, and ommission of other heat sources in the case studies. However, a few general observations can be made, the most important of which are:

- The algorithm can be used as a control tool for planning of optimal response to spot prices.
- There is a unique optimal thermal capacity size associated with each price pattern. This result has important implications for the use of the algorithm in the design of optimal auxiliary storage under various price patterns.
- The algorithm can also be used to assess the value of savings due to investments in heating capacity, insulation, storage, etc. In addition, It can be used to evaluate the trade-off between comfort and savings for individual households.
- The algorithm can be used by the utilities in the evaluation of the load shifts caused by customer response to spot prices.

Chapter 6

ELECTRIC WATER HEATING UNDER SPOT PRICES

6.1 Introduction

This chapter introduces a practical methodology for the optimal operation of dual element electric water heaters under spot pricing. Control of the home appliance electricity consumption, specially electric water heaters, constitute an important aspect of demand-side management [E2]. The most studied f rm of electric water heater control has been the indirect and remote control of the electricity use by the utilities [B3,H1,L1]. The principal drawback of these methods has been the lack of control by the actual users and owners of the controlled water heaters. A similar study on water heater control was done by Wilber [W2], but the presented algorithm was not structurally defined.

Ordinary storage type residential electric water heaters are designed to maintain an adequate storage of hot water throughout the day. Heating elements are turned on as soon as the thermostats sense a decrease in the temperature of the water, which is caused by the cold water from below replacing the hot water drawn at the top. Water heating may occur at any time throughout the day as determined by the hot water demand pattern for the day and the heat loss to the environment.

The electricity consumption pattern follows the demand and loss pattern except for some sluggishness due to the hysteresis properties of the thermostats.

Spot pricing of electricity provides the motivation for a more economical operation of electric water heaters. The storage capability of residential water heaters suggests the possibility of hot water production at the times of low electricity prices and its storage for use at the times of high electricity prices.

In the context of residential electric water heaters, the two most important physical characteristics which determine the scheduling potential are storage capacity and water-heating (or production) capacity of the water heaters. If the demand pattern for the hot water for the next 24 hours in question is somehow known, then it is possible to optimize electricity usage (minimize the cost) by choosing appropriate low cost times for heating of water and taking advantage of storage capability in order to have enough hot water to meet the demand at high cost times. Implicit in the concept of scheduling is the notion of not keeping a full storage of hot water all the time, which is in contrast to the full storage maintenance of the normal water heater operation.

The scheduling methodology consists of two parts. One is to predict hot water demand, and the other is to find the optimal schedule of electricity use subject to the given schedule of spot prices and the predicted schedule of hot water demand.

An important constraining factor in this study is the assumption that there would be no alteration to the basic design of ordinary dual element residential electric water heaters other than electrical connections to the heating elements and thermostats for the data acquisition and control purposes.

6.2 Dual Element Electric Water Heaters

The basic physical system of dual element electric water heaters, as shown in Figure 6.1, consists of a storage tank with a free cold water inlet near the bottom and a hot water outlet at the top. The two heating elements and the associated thermostats are located at the bottom and at more than two third of the way up near the top of the storage tank. Commercially available residential water heaters usually include additional features such as foam-insulated casing for energy conservation, and an anode rod for protection from chemical actions in the water due to electrolysis.

A typical water heater comes in standard tank sizes which may vary from 30 to 80 gallons, the most common being 52 gallons. The heating elements are resistor type metallic tubing with protective coating, and are usually bent in the shape of **U** or **J**. A typical power rating for each element may vary from 4.5 to 6 kWh. Ordinarily, the upper and lower elements both have the same rating. The upper and the lower thermostats are usually set at the same setting for the desired hot water temperature, which is usually around 150 degrees fahrenheit. Due to the hysteresis nature of the thermostats, water temperature drops a few degrees below 150 before the thermostats start activating the heating elements. An important feature of dual element electric water heaters is the interlock between the two heating elements which curbs the operation of the lower element unless the upper element is off.

During the normal operation, a full storage of hot water is maintained. Whenever there is a hot water draw, the cold water enters at the bottom, and being heavier than hot water, it remains at the bottom. Usually, there is little mixing, and the water in the storage tank remains stratified by temperature. As more hot water is drawn, the cold water moves upward; and as soon as the lower thermostat senses the colder temperature (which as mentioned must be more than a few degrees below 150) the lower element is turned on and water is heated until the temperature reaches 150 degrees. If the rate of hot water draw is greater than the production rate of the lower element for a sufficiently long period of time, the hot water depletes beyond the level of the upper heating element and upper thermostat, and consequently, the upper element is also turned on. However, due to the interlock between the two elements, the lower element turns off until after the

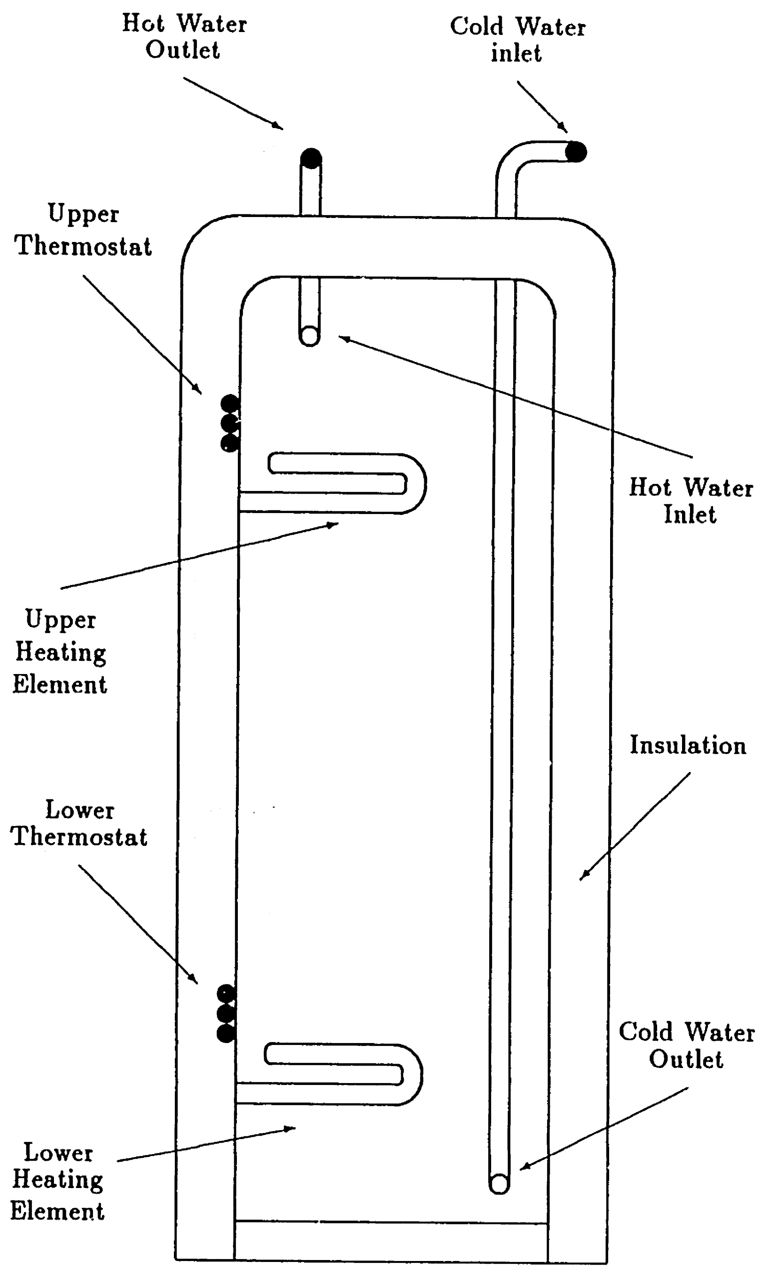

Figure 6.1: Dual Element Electric Water Heater

water temperature around the upper element increases beyond the 150 degrees. In a typical household the upper element does not come on line very often.

6.3 Demand Data

As mentioned before, it is assumed that no basic physical changes other than electrical connections to the heating elements and thermostats are allowed. One exception can be the installation of a flow meter, which would provide data on the time and the volume of water drawn from the water heater. In the absence of a flow meter, the only types of information that can be gathered are the intervals at which the two thermostats and the two elements are on and off. Note that due to the interlock, the lower thermostat and the lower heating element would not necessarily be on at the same time.

The information provided by the heating elements are simply the electricity consumption for the periods during which they are on. Since the elements use electricity at a constant rate, it is possible to calculate the average or total electricity use for the periods at which the upper or lower elements are on. However the electricity usage values do not directly relate to hot water demand values since some of the energy increases the hot water stored, and some is lost through heat transfer to the environment.

The information provided by the upper and lower thermostats relate to the amount of hot water in the storage. In computing the hot water level in the storage tank, the effect of hysteresis of the thermostats should be taken into consideration. In other words, a thermostat turns on its corresponding heating element at a certain water temperature (or hot water level) and it turns it off at a different temperature.

However, *during the interval when a thermostat turns off its corresponding heating element and the next time that it does it again, none of the energy would* *go into the net storage of the hot water since the level of the hot water*, i.e. storage, would be back at where it started. Therefore, any amount of electricity used during such an interval goes either into producing hot water which is subsequently withdrawn, or it is lost to the environment through heat transfer.

Therefore, the combination of the streams of information from the thermostats together with those from the heating elements provide data that can be used as the basis of calculations for the prediction of demand and optimal scheduling of electricity usage.

A computer program would identify the time when the upper and lower thermostats go off, and it would evaluate the total electricity consumption for all the intervals between such periods separately, since the two thermostats do not turn on or off at the same time. The interim results are two different sets of time intervals during which the total electricity use is known, corresponding to the intervals of consecutive series of on-off information provided by the two thermostats. Due to the stochastic nature of demand, not only the two time series differ, but also the time intervals within each set are unequal in length; and furthermore, they change from day to day.

The next step is to translate these values into demand plus loss (or demand for short) patterns. As a start, we can assume uniform demand during each interval, and therefore, demand at each unit of time (period) is simply the ratio of the electricity consumption during each interval to the length of that interval. A different approach is to assume independence of demand at each period during each interval. This assumption results in a binomial probability distribution for the probability of discrete demand levels at each period within an interval. Interestingly enough, the expected value of demand at each period is equal to the average uniform demand for the same interval. In this way, the expected value of demand for all the periods of the day can be evaluated.

It must be kept in mind that a different calculation would be carried out for each of the two sets of data. These two sets of data can be considered as representing the same random variable and can be combined in some appropriate manner in order to determine the expected value of hot water use at each time unit (period) of the day.

The next step is to predict future demand from the evaluated demand patterns of the previous days. A sensible way of doing this would be to separate the results by the day of the week. In other words, to predict demand for a monday, only the data from previous mondays would be used. Then, the changing pattern of demand must be taken into account. One way of incorporating the notion of demand pattern evolution is to assign a higher weight to the more recent data. For example, if W_i denotes the vector of demand for the i'th monday to be predicted, then

$$W_i = 0.4W_{i-1} + 0.3W_{i-2} + 0.2W_{i-3} + 0.1W_{i-4} \tag{6.1}$$

An interesting feature of the above formulation is the effective cut-off of the data of the previous month, which is appropriate in the sense that it takes into account the seasonal (or monthly) variation in the demand pattern.

6.4 Mathematical Model of the Operation

The problem can be modeled in discrete time as a general form of an inventory problem where the storage at any period is equal to the production minus demand and minus loss at that period plus the storage from the previous period.

$$X[k+1] = a\, X[k] + U[k] - W[k] \tag{6.2}$$

where

k time period

$X[k]$	storage at the start of period k
$U[k]$	electricity consumption during period k
$W[k]$	demand during period k
a	storage loss coefficient

The values of W are stochastic and can not be determined before hand in a deterministic fashion unless some simplifying assumptions are made. The values of a depend on the physical characteristics (heat transfer) of the water heater and the environment. Values of U are variables that must be determined such that the total cost of electricity consumption is minimized subject to the capacity constraints and non-negativity constraints for both storage and production capacities.

Values of W can be predicted only through the monitoring of information obtained from the thermostats and the heating elements. An underlying assumption is that hot water usage pattern is not totally chaotic, and it is presumed possible to establish some sort of approximate pattern for daily use of hot water in a particular household by continues analysis of information obtained from the thermostats and the heating elements. Of course, special care must be taken to take into account the evolution of the hot water usage patterns, and the influence of factors such as the time of the year and the day of the week.

6.5 Optimization Methodology

Assuming that on each day the previous electricity usage data is processed and some form of demand pattern for the next day is established, then it is possible to treat the demand pattern as deterministic and solve the deterministic optimization problem for minimizing electricity consumption costs given the forecasted demand and spot prices, and subject to the capacity and non- negativity constraints for the storage and production capacities of the water heater.

The optimization algorithms to be used here, are based on the simplification of the general algorithm described in Chapter 4. If losses are significant (the 1R1C model), then the algorithm becomes similar to the one given by Tsitsiklis [T2].

However, the information acquisition methodology lumps the storage losses as a part of the demand, and therefore, there is no need to worry about the storage losses at this point. Algorithm for single storage without losses (the 0R1C model) takes a specially simple form. It is a special case of the Tsitsiklis algorithm, however, it was developed independently by this author [D1] for an earlier project that led to the present work. All these algorithms are special cases of the general algorithm of Chapter 4. These algorithms are described in the Appendix C.

The basic technique of the algorithm is first to establish an initial solution or electricity usage pattern based on uniform prices. This is the simple pattern of electricity use in the absence of spot prices which corresponds to the demand pattern. The next step is to reschedule the electricity use of the most expensive period with the least expensive period subject to storage and production capacity constraints. The algorithm reiterates this operation for the next pair of the most and the least expensive periods until all the periods have been considered. It has been proved [D1] that each period needs to be considered only once, and that the final solution is optimal and corresponds to the exact result of a linear program, and yet it does it faster.

6.6 Optimization and Demand Prediction

Before applying the algorithm to find the optimal electricity consumption scheduling for the water heater, two distinct problems must be considered first. One is the question of ensuring that the water heater would still generate hot water even if the actual hot water demand does not correspond to the predicted hot water demand. The other is the question of information generation.

Fortunately, the basic design of the dual element water heater lends itself to the solution of these problems. The approach chosen here is to treat the volume of hot water above the upper heating element as a *safety feature* and to *disconnect the upper element from the control signal.* Consequently, the optimal algorithm would consider only the volume between the lower and upper elements to be the depletable storage, and the decision variable whether to turn on or turn off the electricity would only apply to the lower element. Therefore, the upper element would behave in a manner similar to the normal situations, and would turn on whenever it is signaled by the upper thermostat to do so. If the actual demand is exactly equal to the predicted demand, then the upper element would never turn on since the reserve storage would never be used. Only when the actual demand deviates from the predicted demand and the reserve storage above the upper element gets used, would the upper element start generating hot water. The lower heating element, however, would not act as it would under normal operating procedures. The *purpose* of optimal scheduling is to *do away with the maintenance of a permanently full storage*, and that implies that the main storage would usually be less than maximum. Output of the optimal algorithm, which is the optimal production schedule, can be interpreted as schedule for the times when the lower heating element would be allowed to be on. In this manner, the control methodology is combined with the normal functions of the thermostats and the heating elements of the dual element water heater in order to optimize the use of the storage between the two elements, and at the same time to let the upper heating element to behave normally and keep a reserve storage for the times when the main storage is depleted beyond what the optimal algorithm considers to be the minimum level.

An interesting feature of this methodology is that if the actual demand is exactly equal to the predicted demand, then the upper thermostat or the upper element would never come on and no information would be produced that could relate electricity consumption to hot water demand for intervals shorter than 24 hours. Only if the actual demand deviates from the predicted values, would the upper thermostat generate the signals that can be used in the computation and

Heating elements ratings	4.5 kWh
Storage tank capacity	52 gallons
Upper element placement	13 gallons
Hysteresis effect equivalence	15% of storage above the thermostat
Hot water temperature	150 degrees F
Cold water temperature	60 degrees F
Initial storage (for simulation)	25 gallons

Table 6.1: Information for Water Heater Case Study

estimation of hot water demand. There are various options available which need further study. Under any circumstances the information obtained as outlined above would indicate over supply or under supply of electricity; and thus, the major goal of additional work would be to determine those control strategies which result in more information at the cost of less savings, i. e. control strategies which involve occasional and intentional distortion or alteration of the predicted demand in determining the optimal electricity use scheduling.

6.7 Simulation

A simulation of a simple model of a dual electric water heater has been implemented on a personal computer, using periods of 6 minutes as the unit of the discrete time. The simulation runs can be based on either the normal operation or the optimal operation based on a decision variable which is determined by the optimal algorithm. The water heater model characteristics are based on the most widely used residential water heaters. The most important features are given in Table 6.1.

Based on these results the optimal algorithm takes the working storage capacity to be 52 -13 = 39 gallons. The data used are based on the data of a typical household in New Mexico (taken from reference [E2]). Figure 6.2 presents the

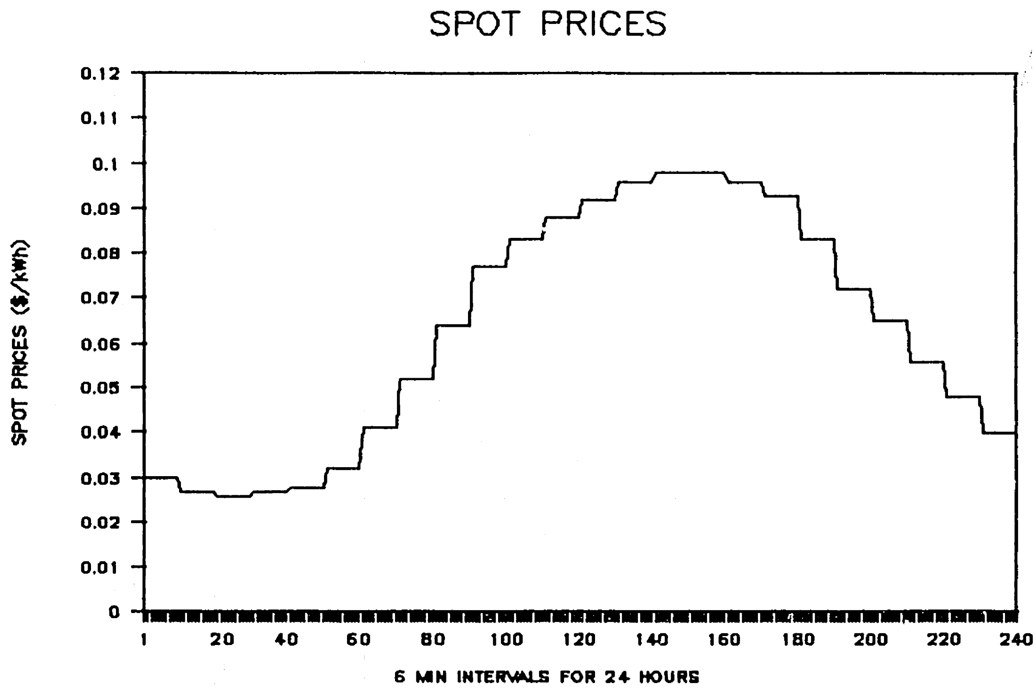

Figure 6.2: Spot Prices Used in Water Heater Study

spot price values used in the optimization.

Figure 6.3 is the demand for hot water in gal/period. Under normal (non-scheduled) operation, electricity is used when hot water is drawn from the water heater. Therefore, the reference electricity usage is same as the hot water demand profile. In addition, normal operation maintains a full storage, unless rate of hot water usage is too high to enable adequate hot water production. However, the heating elements are on as long as the storage level is below maximum. Therefore, the reference storage level is constant at 52 gallons.

Figures 6.4 and 6.5 provide the optimization results for the optimal hot water production (or usage of electricity) and the resulting storage variations, respectively. The optimal production schedule is quite different from the reference or normal production schedule, although for the time horizon, total hot water production equals the total demand at 79.5 gallons. However, the ratio of the optimal total cost to the reference total cost is 5.5 to 3.3.

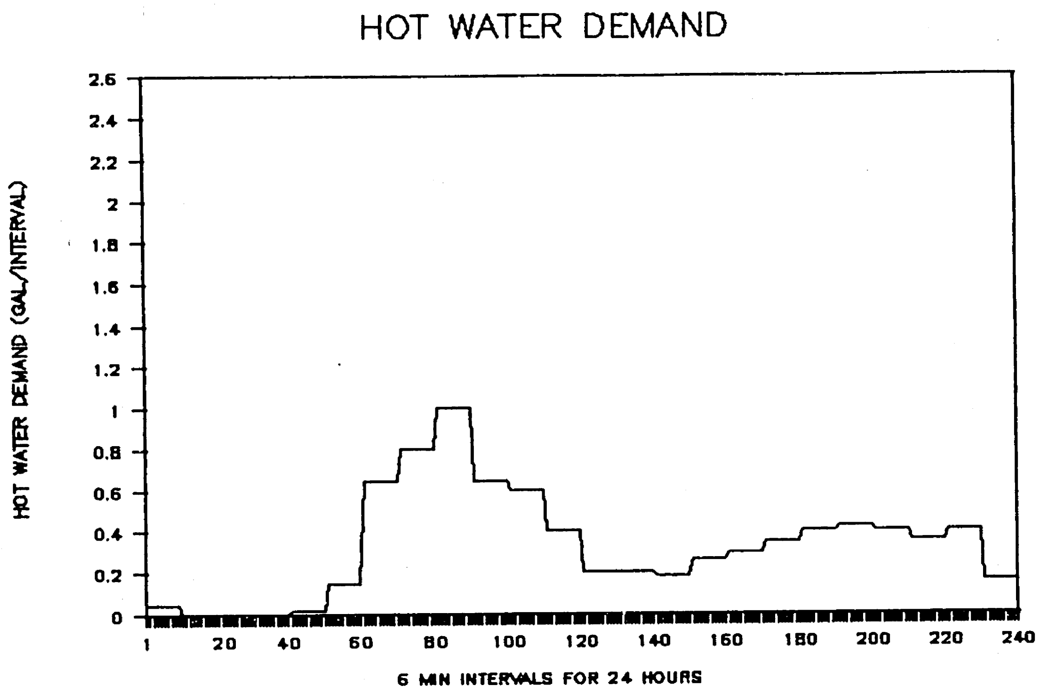

Figure 6.3: Hot Water Demand Profile

It is interesting to note that the optimal production schedule represents a load shift away from the periods of high prices. The reason for increased production at the end of the day is the requirement that the storage level at the end of the time horizon should be equal to initial storage level. This constraint can be relaxed, resulting in a minimum storage level at the end of the time horizon. Another interesting observation is that a full storage is maintained for only a third of the time during the day.

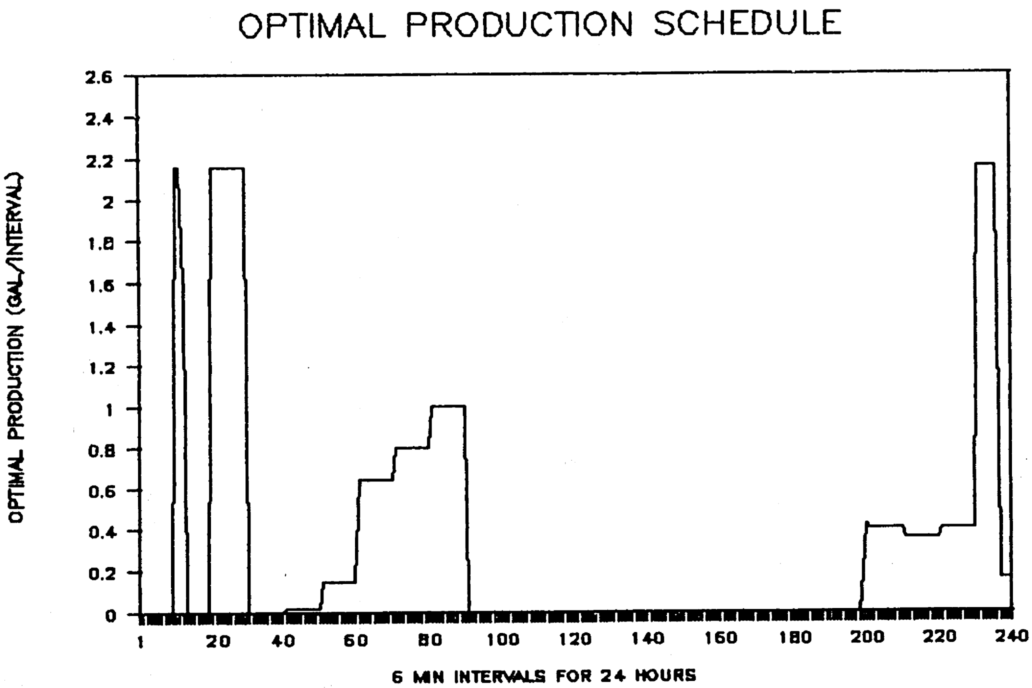

Figure 6.4: Optimal Hot Water Production Profile

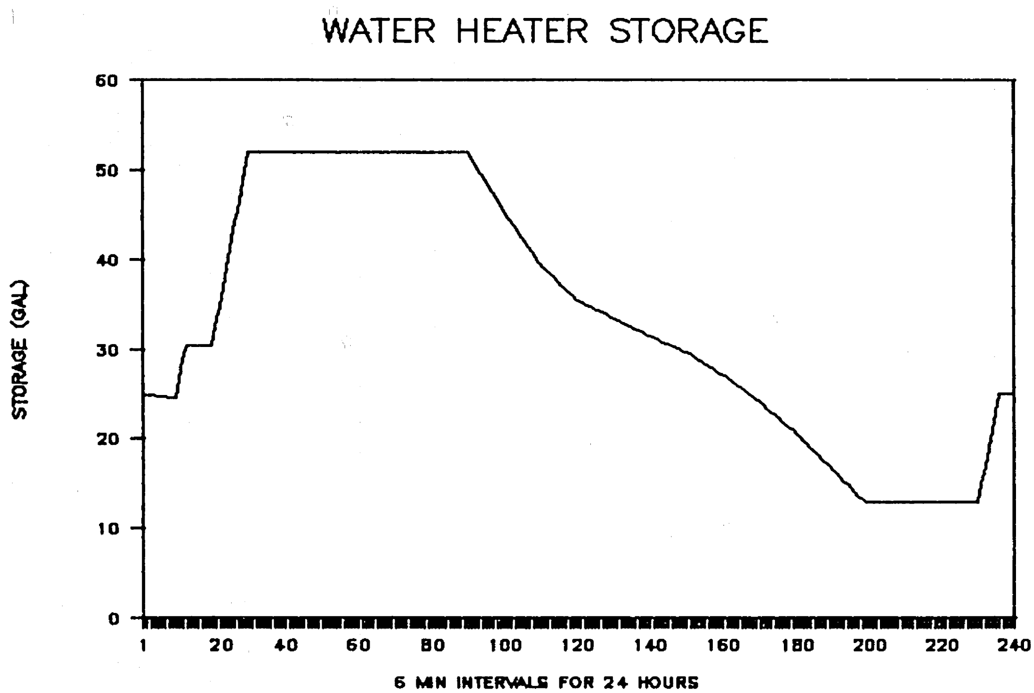

Figure 6.5: Optimal Storage Profile

Chapter 7

SINGLE STORAGE SYSTEM SUBJECT TO STOCHASTIC DEMAND

7.1 Introduction

This chapter presents the results of a dynamic programming approach to the optimal scheduling of electricity consumption under spot pricing and stochastic demand. This study is only a preliminary effort in the formulation of a simple problem, and should be considered more as an appendix to the main body of the thesis. However, it is possible to extend the formulation to more complex systems at the expense of more computation time.

The physical process of interest is the scheduling of input U (electricity usage or hot water) for a single storage unit with no losses (water heater). The storage level may vary between two bounds X_{min} and X_{max}, and the exogenous demand D (for hot water) is stochastic. The input costs for the next 24 hours are available and they are different for each hour. The assumption of no losses is very restrictive, and the results should not be generalized to more complex situations.

Dynamic programming formulation incorporates the physically derived discrete model (state equation) of the heat storage, production, and demand for a single storage unit. The general notations follow those of Bertsekas [B2]. A series of test cases were examined and the preliminary results for an stochastic demand with uniform, but not identical, probability indicate that for a storage unit with no losses:

- The principle of certainty equivalence does not apply to this system.
- Stochasticity of demand does not limit the length of the effective time horizon.
- Compared to ordinary thermostat operation, it is possible to decrease the electricity costs by rescheduling electricity consumption between two periods even if they are separated by a large (infinite) number of periods.
- If there are no storage or production limits, the optimal electricity use level in the first period will not generally remain constant if the length of the time horizon is kept increasing.

The following sections describe the physical model, the dynamic programming formulation, computational implementation, and the results of the case studies.

7.2 The Physical Model

In its simplest form the physical model consists of a storage unit with no losses other than the decrease due to demand. Although demand is stochastic for future periods, it is known for the current period. If at each period demand is high enough to force the storage below its minimum level, the input is automatically increased to meet the demand. Thus storage constraints are never violated.

Starting storage level at any period X_{k+1}, as shown in the following relationship, depends on the scheduled electricity usage (heat production) U_k and demand (for heat) W_k, and the starting storage level X_k of the previous period:

$$X_{k+1} = \min\left[X_{max}, \max\left(0, X_k + U_k - W_k\right)\right] \tag{7.1}$$

$$0 \leq X_k \leq X_{max}$$

$$0 \leq U_k \leq U_{max}$$

where,

X_k	storage at period k
U_k	scheduled electricity use at period k
W_k	stochastic demand at period k
X_{max}	maximum acceptable storage between each period
U_{max}	maximum allowable electricity use at each period

Demand W_k is stochastic with a uniform probability distribution, and it is bounded by W_{min_k} and W_{max_k}, where the bounds may be different for each period. This particular form of demand is chosen to minimize the effort in interpretation of the result. Other probability distributions can be implemented easily. Note that in this section, instead of brackets, subscript k is used to denote time.

The system automatically shuts off operation if storage reaches X_{max}, and as a result, the actual electricity use, $X_{max} - X_k + W_k$, would be less that the scheduled electricity use, U_k in this instance. The system also automatically starts on if the scheduled electricity use is insufficient to keep the storage level above its minimum value (zero level of storage set at X_{min}. In this case the actual electricity use, $0 - X_k + W_k$, is more than the scheduled electricity use U_k.

No additional restrictions were imposed in order to facilitate the dynamic programming formulation and also to gain an insight on the effects of the stochasticity of demand on the optimal scheduling of electricity use without the added complication of more specific assumptions particular to various physical situations. No attempt was made to consider all the particular features of a space heating/cooling unit with storage losses. Although incorporation of storage losses into the dynamic programming formulation is easily possible with no added difficulty, it was not done in order to interpret the results for the simplest case and make the broadest conclusions with regard to the effects of stochasticity.

7.3 Dynamic Programming Formulation

The present problem can be classified as a special case of inventory problems which have been extensively studied by various researchers in mathematical programming and management science. However, there are important differences. The special restrictions imposed on the storage limits and the variability of the price make the analytical approach quite cumbersome for time horizons longer than two periods. The numerical method is based on the discretization of the storage, electricity consumption (hot water production), and demand (for hot water or electricity) levels; and on computation of expected cost associated with each level of electricity utilization for each possible level of storage for each period. The electricity usage levels which result in minimum value of expected cost over the time horizon provide the optimum policy decisions.

In summary, beginning with the last period, dynamic programming formulation finds the optimum level of electricity use for each starting level of storage, together with the associated optimum cost (or cost to go). Then it considers the period before the last one, and finds the level of electricity use at that period which results in total minimum cost to go from that period on for each starting level of the storage. This iterative process is continued backwards in time until all periods within the time horizon are considered.

In this particular problem, future costs after the very last period of the time horizon (period $N+1$) are of no consequence; and thus, the cost to go $J_{N+1}(X_{N+1})$ (cost to go at period $N+1$ as a function of the level of storage at period $N+1$) is zero. At a period k the electricity cost is P_k and for an arbitrarily scheduled electricity usage level U_k and starting storage X_k, the immediate expected cost is:

$$\underset{W_k}{E}\{P_k \ \min[(X_{max} - X_k + W_k), \max(U_k, (0 - X_k + W_k))\ ]\} \qquad (7.2)$$

The expression multiplying the price is the actual electricity consumption (heat production). This formulation implies that if the scheduled U_k is large enough (or W_k small enough) to cause the storage at $k+1$ to go beyond X_{max}, then the system shuts off and the actual electricity usage is limited to $X_{max} - X_k + W_k$. Similarly, if the scheduled U_k is small enough (or W_k is large enough) to cause the storage to drop below zero, then the electricity use is continued beyond U_k with the actual electricity usage being $0 - X_k + W_k$.

The optimal cost to go of the period k, i.e. $J_k(X_k)$, denotes the optimal electricity cost for the periods remaining from period k to the last period, provided that at each period optimal decisions are made. Note that there is one value of J_k associated with each starting level of storage X_k at period k.

Then the optimal value of U_k at period k is the one which minimized the sum of the expected values of the immediate cost at k and the optimal cost to go at $k+1$, or in dynamic programming formulation:

$$\begin{aligned} J_k(X_k) \ &= \ \min_{U_k} \underset{W_k}{E}\{P_k \ \min[(X_{max} - X_k + W_k), \\ &\quad \max(U_k, 0 - X_k + W_k)\ ] + J_{k+1}(X_{k+1})\} \end{aligned} \qquad (7.3)$$

where at the next period k+1 the storage level, as given by the system equation is:

$$X_{k+1} = \min[X_{max}, \max(0, X_k + U_k - W_k)\,] \tag{7.4}$$

Dynamic programming starts with the one to the last period in the time horizon and finds the optimal U_{N-1}, and optimal cost to go $J_{N-1}(X_{N-1})$. Note that in the above formulation the next cost to go, $J_N(X_N)$, is zero for all possible values of S[N] since the costs incurred after the end of the time horizon are immaterial. Therefore, by finding $J_k(X_k)$ for each period and going back, the dynamic programming finds the optimal electricity production and cost to go function for each period.

Actual implementation of dynamic programming algorithm requires some re-working of the main formulation as shown below:

The expression for production at period k,

$$\min\,[(X_{max} - X_k + W_k), \max\,(U_k, (0 - X_k + W_k)\,)] \tag{7.5}$$

can be re-written as:

$$\min[(X_{max} - X_k + W_k), U_k] + \max[U_k, (0 - X_k - U_k + W_k)] - U_k \tag{7.6}$$

which by subtracting and adding $2U_k$ becomes:

$$\min[(X_{max} - X_k - U_k + W_k), 0] + \max[0, (0 - X_k - U_k + W_k)] + U_k \tag{7.7}$$

If we define a new variable Y_k as the sum of storage and scheduled production at period k, then we have:

$$Y_k = X_k + U_k \tag{7.8}$$

Note that at each period the value of Y_k can never be less than the storage at that period. Re-writing in Y_k then the actual production at period k is:

$$\min[(X_{max} - Y_k + W_k), 0] + \max[0, (0 - Y_k + W_k)] = +Y_k - X_k \tag{7.9}$$

Now, the total cost to go at period k which is the expected value of the sum of immediate cost at k and the cost to go at period $k+1$ becomes:

$$\begin{aligned} J_k(X_k) &= \min_{U_k} \mathop{E}_{W_k}\{P_k(\min[(X_{max} - Y_k + W_k), 0] + \max[0, (0 - Y_k + D_k)] \\ &\quad +Y_k - X_k) + J_{k+1}(X_{k+1})\} \end{aligned} \tag{7.10}$$

Now regrouping some of the inner expressions as:

$$\begin{aligned} L(Y_k) &= P_k \mathop{E}_{W_k}\{\min[(X_{max} - Y_k + W_k), 0]\} + \\ &\quad P_k \mathop{E}_{W_k}\{\max[0, (0 - Y_k + W_k)]\} \end{aligned} \tag{7.11}$$

we have the following:

Optimal cost to go function for period k:

$$J_k(X_k) = \min_{Y_k}[P_k\, Y_k + L(Y_k) + \mathop{E}_{X_{k+1}}\{J_{k+1}(X_{k+1})\}] - P_k\, X_k \tag{7.12}$$

for k starting with N and going back to 1, and

$$J_{N+1}(X_{N+1}) = 0 \tag{7.13}$$

for all possible values of X_{N+1}, and, with the equation of state being:

$$X_{k+1} = \min[X_{max}, \max(0, Y_k - W_k)] \tag{7.14}$$

In the program $J_{k+1}(X_{k+1})$ is referred to as JKPLUS, and in the optimal cost to go function the terms in the bracket are referred to as GY, which is the total cost given any Y_k and assuming that X_k is zero. Therefore,

$$GY = P_k\, Y_k + L(Y_k) + E\{J_{k+1}(X_{k+1})\} \tag{7.15}$$

Note that for a given value of storage at period k, X_k, the optimal cost, or the cost or go function $J_k(X_k)$, is attained by the minimization with respect to Y_k of the expressions in the large bracket. We also know that $Y_k = X_k + U_k$, and thus, if the minimization is achieved for some value of Y_k such as Y_{opt} then the optimal decision depends on whether X_k is greater or less that Y_{opt}. In other words:

- If $X_k < Y_{opt}$, then optimal production U_{opt} is $Y_{opt} - X_k$.
- If $X_k > Y_{opt}$, then optimal production U_{opt} is 0.

This is simply due to the fact that Y_k can not be set to less than X_k.

7.4 Computational Implementation

The algorithm was written in *APL*PLUS* and implemented on an *IBM/PC* compatible computer. The program has a modular structure, and therefore, it is quite simple to implement other stochastic forms for demand (driven by the outside temperature) such as a normal distribution. This, however, was not done so as to keep the task of interpretation of results to a minimum.

The program starts by partitioning the range of acceptable storage (0 to X_{max}) and production (0 to U_{max}) and demand (W_{max_k} and W_{min_k} for demand at period

k) into discrete units. It then starts with the one before last period in the time horizon using the our general cost to go equation $J_k(X_k)$ where k is $N-1$. It finds the Y_{opt} by examining all the acceptable values of Y_k and searching for the minimum. For example, it starts with Y_k at zero, and then finds the expectation by examining all the possible discrete values of W_k, which can be between W_{min_k} and W_{max_k}, and multiplying the cost by the probability of demand being at that value (uniform probability in our problem). It then increases Y_k to the next value and finds the expected cost as before, until all the possible values of Y_k have been examined. Finally, it searches for the value of Y_{opt}, which is the value of Y_k which gives the minimum expected cost.

Note that according to our general equation there is one value of cost to go (optimal expected cost) associated which each level of storage at k. Also note that each overall iteration requires the value of cost to go of the next period which, as mentioned before, starts at zero for the period after the end of the time horizon.

In the first iteration the cost to go of the next period, $J_{k+1}(X_{k+1})$, is zero for each level of X_{k+1}. For the first iteration $k+1$ is equal to N. In the intermediate periods at each level of Y_k and W_k the value of X_{k+1} is evaluated from the equation of state, and using this value, the corresponding value of cost to go for period $k+1$, i.e. $J_{k+1}(X_k)$, from the previous iteration is used in the current iteration.

7.5 Results and Conclusions

7.5.1 A hypothetical case

This subsection is not based on our dynamic programming study, but it complements our other results. Compared to ordinary thermostat operation of heating, it is possible to decrease the electricity costs by rescheduling electricity consumption between two periods even if they are separated by a large (infinite) number of periods.

To understand it, imagine a hypothetically infinite number of periods at which storage is allowed to be anywhere between X_1 and X_2. If the system is set at automatic operation with a thermostat setting of X_1, then at each period hot water is produced to compensate for the stochastic demand (for hot water) at that period and keep the storage at X_1. Now suppose that the very last period has a higher electricity cost compared to the first period. Then it is possible to set the thermostats at X_2 for all the periods except the very last and in this way, if there are no storage losses, the additional hot water produced in the first period is carried to the last period in which the storage is allowed to drop to X_1. Since this heat was produced in the first period and not the last, overall cost is lower than the first case. Therefore, when there are no storage losses it is possible, under stochastic demand, to reschedule electricity usage between two infinitely apart periods and still improve the cost.

7.5.2 Selected graphical results

The figures included in this section are meant to provide a graphical presentation of the cost behavior under stochastic demand for the simplest case of the 2 period time horizon. The information shown in Table 7.1 describe the case study considered here. The computation time was 94 seconds.

The total cost to go at the second period, GY_2, is a function of the sum of storage and production at that period. The optimal cost to go value is determined by minimizing this function with respect to the production value U_2. Figure 7.1 shows that this function is a monotonically increasing function of the sum of storage and production, but only up to a point. It remains constant after that point. This is a sensible result, since given a starting storage at the beginning of this period, the total cost will increase with a higher usage rate of electricity, until the storage is saturated. Since the maximum expected demand is 9, then this saturation point occurs when the sum of storage and production is 29 which is the maximum possible.

$N = 2$	number or periods in the time horizon
$P_2 = 4$	price during the first period
$P_1 = 10$	price during the second period
$W_{1_{min}} = 3$	minimum demand during first period
$W_{1_{max}} = 9$	maximum demand during first period
$W_{2_{min}} = 3$	minimum demand during second period
$W_{2_{max}} = 9$	maximum demand during second period
$X_{max} = 20$	maximum possible storage
$U_{max} = 20$	maximum possible electricity consumption
$X_{initial} = 0$	initial storage
$N_{delta} = 20$	discretization number

Table 7.1: The Two-Period Case Study for the Graphical Results

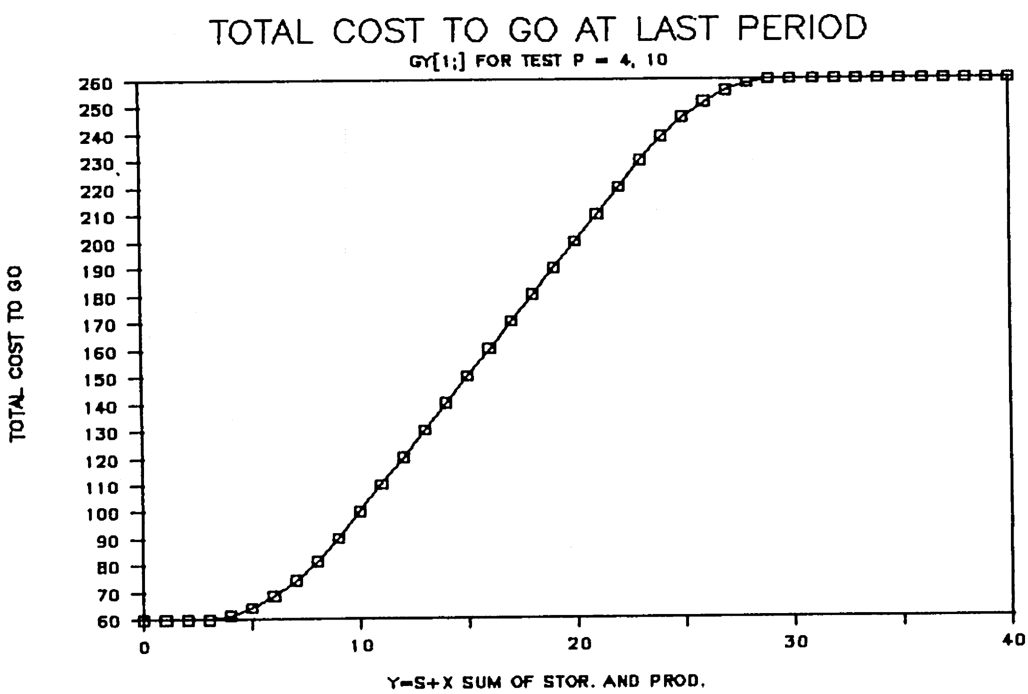

Figure 7.1: Total Cost to Go at the Second Period

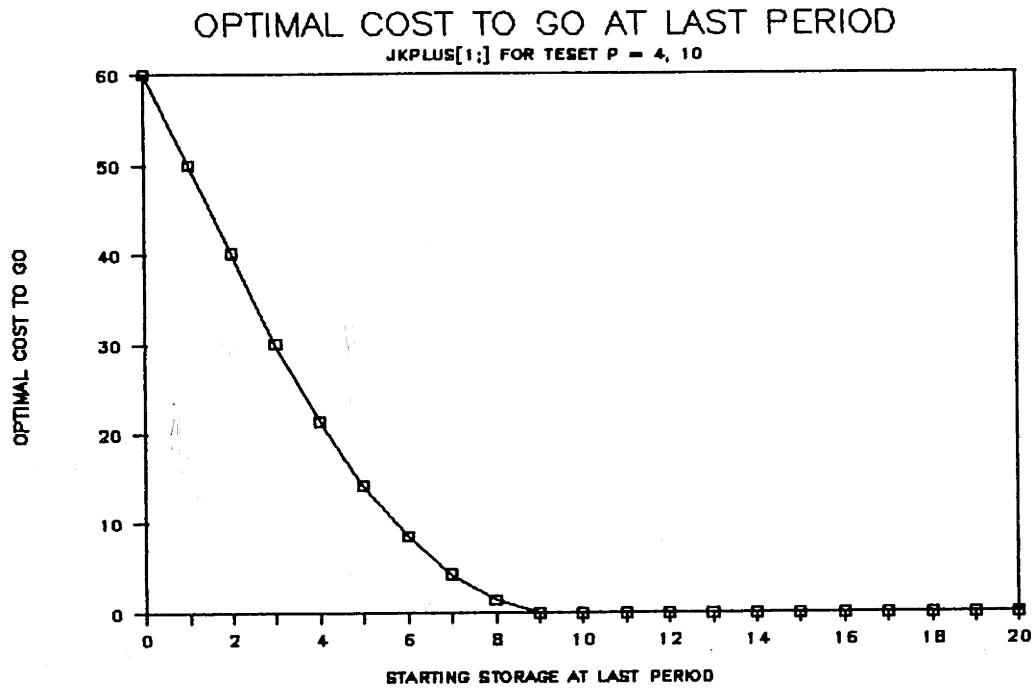

Figure 7.2: Optimal Cost to Go at the Second Period

As seen in figure 7.2 the optimal cost to go at the second period, J_2, is highest when the starting storage in that period is zero. For higher levels of storage at the beginning of the second period, the optimal cost to go decreases monotonically, until it becomes zero when the storage at the start of the second period is 9; which is exactly equal to the value of the maximum expected demand during the second period. Since the total cost to go function of Figure 7.1 is always increasing or constant, its minimum value is obtained when production or U_2 is set to zero.

Figure 7.3 shows the values of the total cost to go at the first period, GY_1, which is at its minimum value when the sum of storage and production is somewhere close to 13. Therefore, if the storage at the start of this period is less than 13, then the production must be set at a value that brings up the sum value to 13. If the storage at the start of the period is more than 13, then production must be set to zero. Note that production can not be negative.

Figure 7.4 depicts the values of the optimal cost to go for the first period, J_1

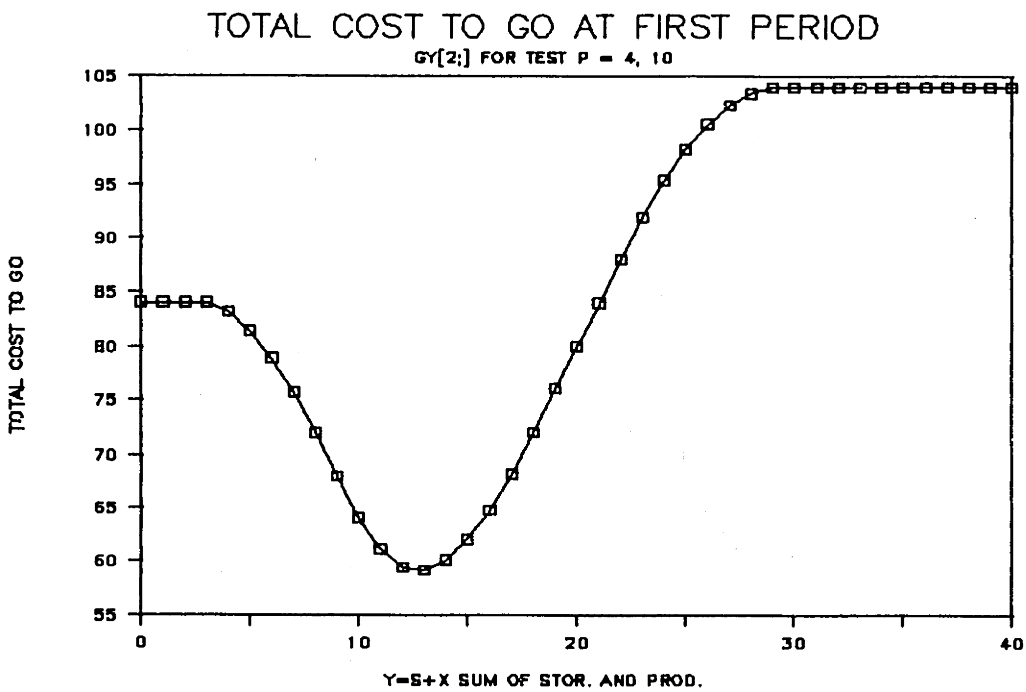

Figure 7.3: Total Cost to Go at the First Period

as a function of the storage at the start of that period. Note that the optimal cost to go is zero if the starting storage is 18, which is the maximum possible total demand.

7.5.3 Numerical results of the two-period case

The optimal value of production at the first period is the value at which the marginal cost of production at the first period (which is an increasing function of production) is equal to the expected marginal savings at the second period (which is a decreasing function of the production at the first period due to the possibility of additional stochastic storage remaining from the previous period).

Note that at Table 7.2 demand at each period is between a minimum of 3 and a maximum of 9 and an expected value of 6. For the two periods combined, the minimum demand is 6, maximum demand is 18, and the expected demand is 12. Since storage level can change only in units of one, then demand at each period

DEMAND=3	3	$X_{initial} = 0$, and $X_{max} = 20$, and $U_{max} = 20$
9	9	$NDELTA = 20$, number of discrete storage levels

IF	THEN
PRICES = 0.1 and 10.0	$U_{opt} = 18$
PRICES = 0.5 and 10.0	$U_{opt} = 17$
PRICES = 1.0 and 10.0	$U_{opt} = 16$
PRICES = 4.0 and 10.0	$U_{opt} = 13$
PRICES = 5.0 and 10.0	$U_{opt} = 12$
PRICES = 7.0 and 10.0	$U_{opt} = 10$
PRICES = 8.0 and 10.0	$U_{opt} = 9$
PRICES = 9.0 and 10.0	$U_{opt} = 8$
PRICES = 9.5 and 10.0	$U_{opt} = 7$
PRICES = 9.8 and 10.0	$U_{opt} = 6$
PRICES = 9.9 and 10,0	$U_{opt} = 6$
PRICES = 20.0 and 10.0	$U_{opt} = 1, 2, 3$

Table 7.2: Numerical Results for The Two-Period Case Study

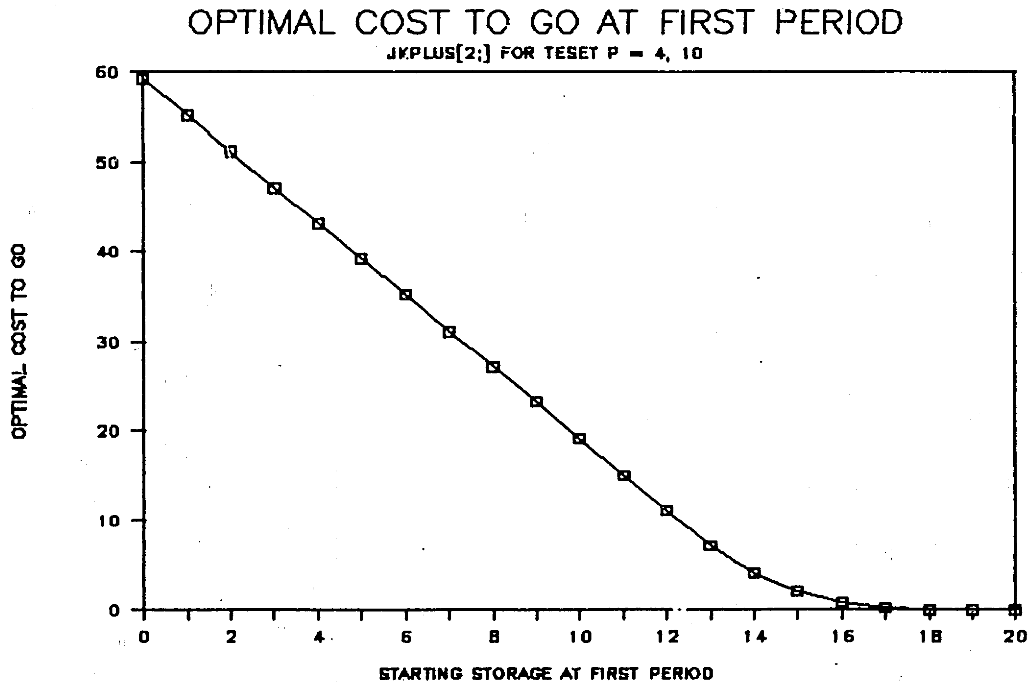

Figure 7.4: Optimal Cost to Go at the First Period

can be either 3 or 4 or 5 or 6 or 7 or 8 or 9, with a probability of 1/7.

Comparing the optimal values of U_{opt} for each period the following observation is made:

- The principle of certainty equivalence does not apply to this system.

7.5.4 Increasing time horizon with limited storage

In this study the time horizon N (number of periods considered) is increased from the earlier period on.

As seen in Table 7.3, if storage and production are unlimited the Y_{opt} would keep increasing with N. For example if X_{max} and U_{max} are each 150 (more than enough) then for $N = 8$, $Y_{opt} = 39$, where total expected demand is 48.

The main result of this case study is that:

- Stochasticity of demand does not limit the length of the effective time hori-

DEMAND = 3 3 3 3 3 3 3 3 3 3 3 $X_{max} = 20$
9 9 9 9 9 9 9 9 9 9 $U_{max} = 20$
$NDELTA = 20$
PRICES = 9 10 10 10 10 10 10 10 10 10

N	$Y_{opt} = X_1 + U_1$(first period)	Total Expected Demand
1	1 or 2 or 3	6
2	8	12
3	13	18
4	19	24
5	27	30
6	29 or more up to 40	36
7	29 or more up to 40	42

Table 7.3: Results for Increasing Time Horizon with Limited Storage

TIME HORIZON= 8
and

DEMAND =	3 3 3 3 3 3 3 3 3 3	$X_{max} = 150$	
	9 9 9 9 9 9 9 9 9 9	$U_{max} = 150$	
		$NDELTA = 50$	
PRICES =	9 10 10 10 10 10 10 10 10 10		

PRICES		Y_{opt}
0.0001	10 10 10 10 10 10 10	72
0.001	10 10 10 10 10 10 10	72
0.1	10 10 10 10 10 10 10	63
9	10 10 10 10 10 10 10	39
9.9	10 10 10 10 10 10 10	33
9.999	10 10 10 10 10 10 10	24

Table 7.4: Effect of Price on Optimal Scheduling

zon.

7.5.5 Effect of prices on optimal consumption

The main results of this study, as shown in Table 7.4, is that everything else remaining the same, not only the relative magnitude order of the prices but also their nominal ratio affects the optimum value of consumption for different periods.

Note that the total expected demand is 48, overall minimum demand is 24, and overall maximum is 72.

Chapter 8

CONCLUSIONS

8.1 Summary and Conclusions

The objective of this thesis was to develop scheduling algorithms for electricity usage under spot prices. The analysis and case studies were restricted to residential electricity consumption, although the results can also be applied to the commercial and industrial sectors.

The first chapter provided the background information on spot pricing and customer types. Chapter 2 presented simple static control schemes for storage and non-storage type customers.

Modeling of thermal systems and their special properties were discussed in Chapter 3. It was shown that thermal systems are asymptotically stable positive dynamic systems, and as a result, behavior of higher order thermal systems has some similarity to simple first-order systems.

Chapter 4 presented an algorithm for the scheduling of electricity usage under spot prices for single-input single-output higher order thermal systems subject to input and output bounds. The algorithm takes advantage of the superposition property of the linear systems, and the monotonically decreasing behavior of the impulse response vector associated with the asymptotically stable positive linear systems. It optimally schedules electricity usage at lower priced periods to compensate for decreased usage at higher priced periods without violating the input and output bounds.

The optimality and efficiency of the algorithms written for the case studies were demonstrated in tests against the simplex method. The relative efficiency of the algorithm increased with the size (time horizon) of the problem.

In Chapter 5 the algorithm was applied to a case study of space heating in a 2R1C (two resistance, one capacitance) building. Savings associated with scheduling under spot prices were defined in terms of electricity usage under equivalent flat rates. In the heating case study, four different price patterns, i.e. impulse, step, single peak, and double peak, were considered. Results reflect the following general conclusions:

- Savings associated with scheduling under spot prices depend on the output bounds, input bounds, amplitude of price variations, and thermal properties of the building.
- The output bounds are defined by the comfort band which specifies the acceptable range of inside air temperature. An increase in the comfort band provides more flexibility in scheduling of the heat input, and therefore, results in higher savings. This temperature band also defines the maximum storage capacity available.
- The input bounds are defined by the power rating of the heating system. A more powerful heating system results in faster heating of the thermal mass of the building. The returned heat from the thermal mass is used later to substitute for the heat input at higher priced hours.
- Usually, depending on the values of the building parameters, only one of the maximum bounds on the input or the output are active. For instance, if at each period, the inside air temperature would reach its maximum level with less heat input than the possible maximum, then a more powerful heating

system will not add to the savings. Thus the algorithm can be used in economic assessment of the value of additional heating capacity or increase in the range of the comfort band.

- Higher savings are associated with more pronounced price variations. For the price patterns studied, and the particular building being considered, a ratio of 10 to 1 between the highest price to lowest price results in a savings rate of about 12% for the impulse response, and 15% or higher for other price patterns. The savings rate increases as the price range is increased.
- Savings depend on the thermal properties of the building. The most interesting behavior is associated with the size of the thermal mass in the building. There appears to exist a unique optimal design value for the thermal mass for each price pattern considered. If the thermal mass of the building is decreased below, or increased above this value, the savings rate is decreased in both cases. The optimal design value of the thermal mass is lowest for the impulse price pattern. This means that under an impulse price pattern, a less massive building results in higher savings than a more massive building.
- Savings depend on the rate of heat transfer between the inside air and the thermal mass of the building. Generally, a higher heat transfer rate results in higher savings. Therefore, increasing this heat transfer rate by various means such as installation of internal fans can increase the savings.
- As common sense dictates, results confirm that savings increase as the rate of heat transfer to the outside is decreased. Therefore, better insulation results in higher savings.
- The last three items indicate that the algorithm, in addition to being a control methodology, can be used as a design tool in determining the optimal values of the thermal properties for buildings under various spot price patterns.

In Chapter 6 A simplified version of the algorithm was applied to the control of a dual element electric water heater. For the hot water demand profile and the price pattern used, a savings of about 40% was observed. The results are based on the unrealistic assumption of perfect predictability for hot water demand. However, hot water demand is confined to certain hours of the day, and a demand profile can be assembled from the observation of data from previous days.

Chapter 7 provides a preliminary application of stochastic dynamic programming methodology to a simple system consisting of a single storage with no losses. Although the results should not be generalized beyond the specificity of the problem addressed, the formulation presented can be used as a starting point for consideration of more complex systems and situations.

8.2 Areas for Further Research

The following list includes possible research studies that either complement or extend the present work. Most are based on issues not explicitly or thoroughly addressed by this thesis. Suggested are:

Studies in optimization:

- Detailed theoretical and computational comparison of the algorithm to alternative techniques such as the simplex method, Lagrange multiplier method, networks with gain, projected gradient method, etc.
- Extension of the algorithm to include more general multi-input multi-output dynamic systems with nonlinear inputs and storage costs.
- Extension of the algorithm to include integer constraints associated with start-up and shut-down operations.
- Extension of the impulse response method to the linear control of more complex dynamic systems

Studies in energy management:

- Application of the algorithm as a control methodology to more complex models of space heating and cooling, specifically, to buildings with greater numbers of lumped thermal masses and more complex heat transfer configurations.
- Application of the algorithm as a design tool for the economic evaluation of the auxiliary storage and insulation in buildings under various spot price patterns.
- A more unified approach to system identification, parameter estimation, and system control for space heating and cooling in buildings.
- Further studies on the demand prediction for dual element electric water heaters based on the observation of the on-off signals from the thermostats and heating elements.
- Design of a complete energy management system under spot prices.
- Actual field experimentation to evaluate the economics and practical implementation of the feed forward and other control schemes.

Studies in related areas:

- Application of the algorithm to other electricity consumer and supplier storage processes.
- Application of the algorithm to other engineering and economic problems with similar structures. Examples are lot sizing and inventory control in management science, and hedging problems in field irrigations.
- Study of the serial storage processes encountered in industrial applications.
- Study of the impact of the customer response on the utility operations and economics.

Stochastic problems:

- A more detailed study of the dynamic response under spot prices for complex systems under stochastic inputs and stochastic prices.

In summary, this thesis has shown the applicability of special purpose optimal control logics for scheduling of residential response to spot prices. The algorithms developed for HVAC, for hot water heaters, and for other more simple electric consuming appliances were shown to be computationally efficient, and to be of practical application in utility load control devices. As with any study of this nature, for every question answered at least as many new avenues for theoretical development and application have been opened.

Appendix A

DETAILS OF THE ALGORITHM FOR A MULTI-STORAGE SINGLE-INPUT SINGLE-OUTPUT THERMAL SYSTEM

A.1 Introduction

In this appendix, the complete algorithm introduced in Chapter 4 is presented in the informal notation of **pidgin pascal**. Following sections are directly based on the various **procedures** of the algorithm. A description of the pidgin pascal is also provided. A cursory look at the procedures prior to the reading of this section will prove helpful in understanding the materials being presented.

Feasibility is always maintained because no constraining limit is violated. However, the intermediate solutions may not necessarily form a *basis*, however, the final solution results in a *spanning tree* configuration which can not be improved further. The program has a modular structure consisting of separate procedures or subroutines. Some procedures are called in within the other procedures. In order to distinguish procedures from variables, the procedure names start with the prescript **P-** before their names.

A.2 Pidgin Pascal

The *pidgin pascal* notation used in this appendix is similar to the *pidgin algol* notation used by Papadimitriou and Steiglitz [P1]. The motivation for using pidgin pascal is the desire to present a general and informal algorithmic code without any strict adherence to specific higher level languages. In addition, the exact forms of the statements and expressions need not be specified. The objective is to simply present the general ideas without being specific. The notations defined in this appendix are very similar to those of **PASCAL** programming language. For clarity, program-specific statements are depicted by bold face letters in the algorithm. A short description of the *pidgin pascal* notation follows:

- **program** *name*: Refers to the main program or routine in the algorithm. Other subprograms or subroutines are called in this program.
- **procedure** *name*: Refers to the subroutines in the algorithm. Each separate **procedure** is identified by its unique name.
- **call** *name*: This statement calls procedure *name* within the main program or the other procedures.
- *Assignment:* has the following form:

  $$variable := expression;$$

  where the *expression* involves any algebraic or logical operation depending on the type of the *variable*. The expression is evaluated and the result is assigned to the variable. The semi-colon marks the end of the expression or statements.

- **if** *condition* **then** *statement*: If the condition is true then the statement is evaluated.
- **else** *statement*: This is an optional addition to the *Conditional Statement*. If the condition is not true then the statement is evaluated.
- **elseif** *statement*: This is similar to **else** , except that more than one condition can be tested by having many **elseif** statements.
- **for** *list* **do** *statement*: The statement is repeated for the parameters identified by the list. This is identical to the *do loop* statements.
- **while** *condition* **do** *statement*: The condition is checked, and as long as it is true, the statement is repeatedly evaluated.
- *Compound Statement:* If there are more than one statement as a part of a conditional or do loop arguments to be evaluated, then these statements are enclosed between a **begin** , and **end** pair.
- **and** , **or** : These are used to construct compound conditions in the **if** and **while** statements.
- **continue** : Is a dummy statement which is ignored in the execution of the program.
- **comment** : Is a statement used for inserting comments anywhere in the program. This statement is also ignored in the execution of the program.
- **input** and **output** : These are the input and output statements in the program.

A.3 Description of the Program and the Procedures

The Master Program

The master program P-Run lays out the main routines of the algorithm. Each particular routine is assigned to a separate procedure. After setting up the inputs and parameters, it calls the P-InitialSolution procedure which finds an initial feasible solution. It then calls the P-FindIVector procedure which finds the *impulse response* vectors $\mathbf{I}^k$ for unit input increase at each period k. Next, it calls the P-FindJVector procedure, which constructs the elementary direction vectors $\mathbf{J_u}^k$ for each period k from which all subsequent feasible direction vectors $\mathbf{F_u}^k$ and also all the output change vectors $\mathbf{F_y}^k$ are constructed. Elements of vector $\mathbf{J_u}^k$ are the amount of decreases in inputs of all and every future period if the input of the first period is increased by one unit, provided that each future input is decreased as much as possible in a chronological order.

Next, the program checks to see if inputs of some periods are already at their minimum level, which, as a result, can not be decreased further. This is carried out in the P-FindFVectors procedure which modifies all the $\mathbf{J}$ vectors into the $\mathbf{F}$ vectors.

Procedure P-MarginalSavings is called to find the savings associated with unit increases and subsequent scheduling in the inputs of each period. The ith element of the $\mathbf{S}$ vector is the marginal saving associated with unit increase in the input of the ith period and the rescheduling of subsequent inputs in the calculated feasible direction.

If some of the marginal savings are positive, then rescheduling of inputs among some periods is economically feasible. At this point the program calls the procedure P-SelectChargePeriod which finds the period with the highest marginal savings. This period is named the *Charge Period* or k_c, and the associated value of its marginal saving is s_{k_c}.

A period at which either the input u or the output y is at the maximum is excluded from this search since any increase in the input of that period will violate one or both of the constraints at that period. Other periods are excluded if they result in an output increase above the maximum in some future periods where no input reductions are possible.

The program terminates if no period has positive marginal savings.

The feasible direction vector $\mathbf{F_u}^{k_c}$ is the change in the vector of inputs $\mathcal{U}_N$ if the rescheduling is performed for one unit increase in the input of the k_c. The output change vector $\mathbf{F_y}^{k_c}$ determines how much vector of outputs $\mathcal{Y}_N$ will change by the same unit increase in the input of k_c and the subsequent rescheduling.

The next task is to determine how much one must go in the feasible direction. In other words how much should the input at k_c be increased until a constraint is met. This value of increase is determined by the procedure P-UIncrease which finds the maximum amount by which the input at k_c can be increased. This factor is named α.

Depending on which constraint is met by moving in the current feasible direction $\mathbf{F_u}^{k_c}$ when $u[k_c]$ is increased by α, the procedure P-AdjustJVector is called to find the new input and/or output constraints which have just become active. If at a certain period the input has just becomes zero, then the input of that period can not be decreased any more. In this case the procedure P-AdjutJVector calls the procedure P-Exclude to modify the $\mathbf{F_u}$ and $\mathbf{F_y}$ vectors. Also, if the output of a certain period reaches the maximum, and, if previously the input of that period had not been decreased because that period had been chosen earlier as a charge candidate, then the P-AdjustJVector calls the P-Include procedure to include that period again in the list of periods at which input an be decreased, and to modify the $\mathbf{F_u}$ and $\mathbf{F_y}$ vectors accordingly.

In addition, the marginal savings vector **S** also changes simply because the $\mathbf{F_u}$ vectors are changed. Updating of the **S** vector is done by the P-Exclude and P-Include procedures.

If the value of α is non-zero, the procedure P-NewSolution is called to find the new block vector of inputs and outputs.

The loop is started again be selecting the next charge candidate. It is quite possible that the previous charge candidate is chosen again.

Each procedure (or subroutine) called in P-Run performs a special task or module. The algorithm is organized in modular form in order to make its structure as clear as possible. The P-Run program controls the main flow the routines and checks for the end of the iterations and the end of the search.

(a) **P-InputData and P-Coefficients :** Procedure P-InputData inputs the system parameters and other data such as the stream of prices $\mathbf{P}[k]$ and the stream of exogenous variables $\mathbf{W}[k]$. In addition, this procedure inputs the initial conditions $\mathbf{X}[1]$, and defines other variables that will be used in the program, and initializes values of some of these variables.

Procedure P-Coefficients uses the parameter data to calculate the system coefficients.

(b) **P-InitialSolution :** Procedure P-InitialSolution uses the state equations to find an initial feasible solution for the inputs $u[k]$ for all k, given the initial conditions $\mathbf{X}[1]$, and the values of the exogenous inputs $w[k]$ for all N periods. This procedure starts with the first period and forces the output to $y[k]_{min}$ at subsequent periods, and finds the N values of u_k for $k = 1, 2, ...N$ without violating the maximum ($u[k]_{max}$) and the minimum ($u[k]_{min}$) input constraints for each period.

At some periods, the input may become maximum before the output reaches the minimum value. This is because at some periods the exogenous input (which is negative in most cases) can be quite large in magnitude. Then, even the maximum input of that period may not be enough to satisfy the demand of the exogenous input. In such a case, one should go back in time and increase the input of the previous periods to satisfy the extra demand of the present period, and if necessary use as many periods as possible.

It is also quite conceivable that no combination of inputs may result in outputs which are within satisfactory bounds. In a sense the bounded output system may not be reachable. Thus, at each period of the recursive calculations, the values of resulting inputs must be checked, and if they fall outside the allowable bounds, they must be adjusted. As a result the output may not be strictly at its lower bound, and possibly, the output may not even be within its permissible bounds.

The state equations are solved in a recursive manner for $k = 1, 2, \ldots N$, i.e. the initial solution is found by recursively finding the values of inputs which result in the minimum values of outputs. Given the initial values $\mathbf{X}[1]$, or alternatively, $\mathbf{Y}[0]$ and $\mathbf{U}[0]$, one solves for the following equations in a loop starting with $k = 0$:

$$u^{init}[k] = \frac{1}{d_u}(y[k]_{min} - \mathbf{CX}[k] - \mathbf{D_w W}[k]) \tag{A.1}$$

It is possible that the solutions for the input found from above may not be feasible. In this case values of the inputs must be adjusted as follows

if $u^{init}[k] \leq u[k]_{min}$
  then $u^{init}[k] = u[k]_{min}$

and

if $u^{init}[k] \geq u[k]_{max}$
  then $u^{init}[k] = u[k]_{max}$

The value of the output must also be re-evaluated :

$$y[k] = \mathbf{CX}[k] + d_u u^{init}[k] + \mathbf{D_w W}[k] \qquad (A.2)$$

Then the values of the state variables for the next period are evaluated and the loop is started again :

$$\mathbf{X}[k+1] = \mathbf{AX}[k] + \mathbf{B_u} u^{init}[k] + \mathbf{B_w} w[k] \qquad (A.3)$$

Thus the procedure P-InitialSolution generates an initial feasible solution, i.e. feasible inputs providing feasible outputs, for all the periods in the time horizon without regard to cost. In fact the initial solution is the feasible solution with the minimum total input usage, but not necessarily the minimum cost.

(c) **P-FindIVectors :** This procedure finds the impulse response vectors **I**. The impulse response is the change in the stream of the outputs as the input of any period is increased by one unit.

The impulse response vector for the input increase at period k is simply the kth column of matrix $\mathcal{D}$.

(d) **P-FindJVectors :** This procedure finds the elementary direction vectors $\mathbf{J_u}$. The elementary direction vector for period k is $\mathbf{J_u}^k$ whose elements are the amount of possible chronological reductions in future inputs if the input of the period k is increased by unity. These vectors are constructed using the information contained in the impulse response vectors.

Also found are the corresponding elementary output change vectors $\mathbf{J_y}$ which represent the movement of the output block due to a change of inputs in the direction of $\mathbf{J_u}$.

This procedure ignores the possibility that a future input may already be at its minimum value. Thus, movements represented by $\mathbf{J_u}$ are not necessarily feasible. This possibility is taken up in the next procedure.

(e) **P-FindFVectors :** This procedure adjusts the elementary direction and output change vectors by considering the future inputs which are already at their minimum. The same routines can also be accomplished by the P-Exclude procedure.

The resulting vectors are the feasible direction vectors $\mathbf{F_u}$ and the feasible output change vectors $\mathbf{F_y}$.

(f) **P-MarginalSavings :** Marginal savings s_k is the net savings due to increase in input at k and the decreases in the inputs of the subsequent periods. In other words it is the savings due to the movement in the feasible direction per unit increase in the value of the input at k.

To find the values, this procedure simply multiplies the prices by the elements of the $\mathbf{F_u}$ vectors and sums up the savings and subtracts the cost of the increasing the input of the associated period.

(g) **SelectChargePeriod :** A charge candidate is selected according to the highest current marginal savings value. One category of exceptions are that periods where either the output or the input are at the maximum. At these periods, the input can not be increased further without violating the constraints, and therefore, can not be considered as charge candidates.

The other group are those which can cause a future output to go above its maximum limit. This may happen if a future input is already at its minimum value.

This task is performed by a simple search for the minimum marginal savings among the permissible set of periods, and the first chronological period among the periods with minimum marginal savings is taken as the charge candidate k_c.

If all the marginal savings are zero or negative, then the program is terminated and the last set of solutions are the final solutions. Otherwise, the program proceeds.

(h) **<u>P-Uincrease :</u>** The procedure P-UINCREASE determines the size of increase in the input at k_c. Let's call the amount to be increased α. To fined α, increase input at k_c until:

- Input at the charging period reaches the maximum level. Then the amount of increase is:

$$\mu_1 = u[k_c]_{max} - u[k_c] \tag{A.4}$$

- Output at the charging period reaches the maximum level. Then the amount of increase is:

$$\mu_2 = \frac{1}{d_u} \cdot (y[k_c]_{max} - y[k_c]) \tag{A.5}$$

- Input at some future period reaches the minimum level. Then the amount of the increase is the *minimum* of the values of

$$\mu_3 = \frac{y[l]_{max} - y[l]}{\delta y_l^{k_c}} \tag{A.6}$$

  for all $l > k_c$ excluding those periods whose inputs are not going to be decreased, i.e. when $\delta y_l^{k_c} = 0$. The first period which results in the minimum is called m_3. This period is used later on to adjust the $\mathbf{F_u}$ and $\mathbf{F_y}$ vectors, as well as the marginal savings $\mathbf{S}$ if necessary.

- Output at some future period reaches the maximum level. Then the amount of the increase is the *minimum* of the values

$$\mu_4 \frac{u[l]}{\delta u_l^{k_c}} \tag{A.7}$$

  for all $l > k_c$ excluding those periods whose inputs are not going to be decreased, i.e. when $\delta u_l^{k_c} = 0$. The first period which results in the minimum is called m_4. Again, this period is used later on to adjust the $\mathbf{F_u}$, $\mathbf{F_y}$ and $\mathbf{S}$ vectors.

The minimum of all these values is the amount by which the input at k_c can be increased without violating any constraint. this is called α.

(i) **P-NewSolution :** Now the new solution is :

$$\mathcal{U} = \mathcal{U} + \alpha \cdot \mathbf{F_u}^{k_c} \tag{A.8}$$

$$\mathcal{Y} = \mathcal{Y} + \alpha \cdot \mathbf{F_y}^{k_c} \tag{A.9}$$

(j) **P-AdjustFVectors :** After finding the new solution set the algorithm must search to see which constraint was encountered first and adjust the values of feasible direction and output change vectors accordingly. This is carried out by the procedure P-AdjustJVector . It first excludes the input of the charge period from future reductions by calling P-Exclude procedure. Then it looks for the period at which a new constraint was encountered. There are four possibilities:

- The new constrained variable is the input at k_c: take no action. In the next round the present charge period will not be a candidate again, since its input is already at the maximum.
- The new constrained variable is the output at k_c: **call** procedure P-INCLUDE (k_c). This is done to reverse the earlier exclusion of k_c which was done by immediately calling procedure P-Exclude earlier.
- The new constrained variable is an output at some future period m_3: **call** procedure P-Include (m_3) if this period was already excluded, and if the input at this period is still above its minimum.
- The new constraint variable is an input at some future period m_4: **call** procedure P-Exclude (m_4). Since input at m_4 is zero, then no further decrease in input of this period is possible, therefore it must be excluded.

The vector Ex keeps a record of the periods which are excluded at any time. For example, if period k is excluded at the moment, the Ex_k is one, otherwise it is zero.

P-Exclude (j) **:** This procedure excludes period j from future input reductions. It updates the $\mathbf{F_u}$, $\mathbf{F_y}$, and $\mathbf{S}$ vectors by taking a compensatory step in the $\mathbf{F_u}^j$ direction. The length of this compensatory movement depends on the feasible direction vector which is being updated. for example, if period j is to be excluded, and $\mathbf{F_u}^i$ is to be updated, then the new vector is:

$$\mathbf{F_u}^i := \mathbf{F_u}^i - \delta u_j^i \mathbf{F_u}^j \tag{A.10}$$

Note that the size of the movement depends both on i and j. Similarly, the update of the output change vector is:

$$\mathbf{F_y}^i := \mathbf{F_y}^i - \delta y_j^i \mathbf{F_y}^j \tag{A.11}$$

This procedure keeps a record of the excluded periods in the vector Ex.

P-Include (j) **:** This procedure does the operations of P-Exclude in reverse. In short, the direction of compensatory movement is the reverse of the one in the P-Exclude procedure. The resulting equations are similar except for a change of sign.

P-MSUpdate : This procedure updates the marginal savings. The change is the savings is:

$$s_i := s_i - \delta u_j^i(-p[j] + \sum_{k=j+1}^{N} p[k]\delta u_k^j) \tag{A.12}$$

The updates have to be performed for all i.

A.4 The General Algorithm in Pidgin Pascal

program P-Run

```
program P-Run;
begin
   call P-Parameters;
   call P-InputData;
   call P-InitialSolution;
   call P-FindIVectors;
   call P-FindJVectors;
   call P-FindFVectors;
   call P-MarginalSaving;
   call P-SelectChargePeriod;
   if s_{k_c} > 0 do
   begin
      call P-UIncrease;
      if α > 0 then
         call P-NewSolution;
      call P-AdjustFVectors;
      ( comment P-AdjustFVectors calls procedures
                  P-Exclude and P-Include );
   end ;
   else
      output results ;
end ;
```

procedure P-Parameters

```
procedure P-Parameters;
begin
   input N
   ( comment number of periods in time horizon );
   input A Bu Bw C Du Dw ;
end ;
```

<u>**procedure P-InputData**</u>

```
procedure P-InputData;
begin
   for k = 1 to N do
   begin
      input P[k], W[k];
      ( comment price and exogenous variables all periods );
   end ;
   input X[1];
   ( comment initial conditions );
   for k = 1 to N do
   begin
      EX_k := 0;
      ( comment to keep track of excluded periods );
   end ;
end ;
```

procedure P-InitialSolution

procedure P-InitialSolution;
begin
  for $k := 1$ to N **do**
  begin
    $u^{init}[k] := \frac{1}{d_u}(y_{min} - \mathbf{CX}[k]) - \mathbf{D_w W}[k]$;
    (**comment** for SI-SO matrix operations result in scalars);
    if $u^{init}[k] \leq u[k]_{min}$
      then $u^{init}[k] := u[k]_{min}$;
    if $u^{init}[k] \geq u[k]_{max}$
      then $u^{init}[k] := u[k]_{max}$;
    $y^{init}[k] := \mathbf{CX}[k] + d_u u^{init}[k] + \mathbf{D_w W}[k]$;
    (**comment** $y[k]$ has to be re-evaluated for readjusted inputs);
    $\mathbf{X}[k+1] := \mathbf{AX}[k] + \mathbf{B_u} u^{init}[k] + \mathbf{B_w W}[k]$;
  end ;
  $\mathcal{U}_N := \mathcal{U}_N^{init}$;
  $\mathcal{Y}_N := \mathcal{Y}_N^{init}$;
  (**comment** *init* : initial);
end ;

procedure P-FindIVectors

```
procedure P-FindIVectors;
begin
   for k := 1 to N do
   begin
```

$$\mathbf{I}^k := \begin{bmatrix} 0_{k-1} \\ \vdots \\ d_u \\ \mathbf{CB_u} \\ \mathbf{CAB_u} \\ \vdots \\ \mathbf{CA}^{N-k-1}\mathbf{B_u} \end{bmatrix};$$

```
   end ;
end ;
```

procedure P-FindJVectors

procedure P-FindJVectors;
begin
  for $k := 1$ to $N - 1$ **do**
  begin
    for $i := 1$ to k **do**
      $\delta u_i^k := 0;$
    $\delta u_1^k := +1;$
    $\mathbf{L} := \mathbf{I}^k;$
    for $i := k + 1$ to N **do**
    begin
      $\delta u_i^k := \frac{\mathbf{L}_i}{i_i^i};$
      $\mathbf{L} := \mathbf{L} + \delta u_i^k \cdot \mathbf{I}^i;$
    end ;

$$\mathbf{J_u}^k := \begin{bmatrix} \delta u_1^k \\ \delta u_2^k \\ \vdots \\ \delta u_N^k \end{bmatrix};$$

    (**comment** subscript u indicates change in inputs);

$$\mathbf{J_y}^k := \mathbf{L} := \begin{bmatrix} \delta y_1^k \\ \delta y_2^k \\ \vdots \\ \delta y_N^k \end{bmatrix};$$

    (**comment** subscript y indicates change in outputs);
  end ;
end ;

procedure P-FindFVectors

procedure P-FindFVectors;
begin
for $k := 1$ to N **do**
begin
$\mathbf{F_u}^k := \mathbf{J_u}^k$;
$\mathbf{F_y}^k := \mathbf{F_y}^k$;
end ;
for $k := 2$ to N **do**
begin
if $u[k] := u[k]_{min}$ **then**
begin
$Ex_k := 1$;
for $i := 1$ to $k - 1$ **do**
begin
$\mathbf{F_u}^i := \mathbf{F_u}^i - \delta u_k^i \mathbf{F_u}^k$;
$\mathbf{F_y}^i := \mathbf{F_y}^i - \delta u_k^i \mathbf{F_y}^i$;
end ;
end ;
end ;
end ;

<u>**procedure P-MarginalSaving**</u>

```
procedure P-MarginalSaving;
begin
   for k := 1 to N do
   begin
      s_k := −p[k] + Σ_{i=k+1}^{N} p[i]δu_i^k;
   end ;
end ;
```

$$s_k := -p[k] + \sum_{i=k+1}^{N} p[i]\delta u_i^k$$

procedure P-Exclude (j)

procedure P-Exclude (j);
begin
  $Ex_j := 1$;
  for $k := 1$ to $j-1$ **do**
  begin
    $s_k := s_k - \delta u_j^k(-p[j] + \sum_{i=j+1}^{N} p[i]\delta u_i^j)$;
    $\mathbf{F_u}^k := \mathbf{F_u}^k - \delta u_j^k \mathbf{F_u}^j$;
    $\mathbf{F_y}^k := \mathbf{F_y}^k - \delta u_j^k \mathbf{F_y}^j$;
  end ;
end ;

procedure P-Include (j)

procedure P-Include (j);
begin
  for $k := 1$ to $j - 1$ **do**
  begin
    $s_k := s_k + \delta u_j^k(-p[j] + \sum_{i=j+1}^{N} p[i]\delta u_i^j)$;
    $\mathbf{F_u}^k := \mathbf{F_u}^k + \delta u_j^k \mathbf{F_u}^j$;
    $\mathbf{F_y}^k := \mathbf{F_y}^k + \delta u_j^k \mathbf{F_y}^j$;
  end ;
end ;

procedure P-SelectChargePeriod

procedure P-SelectChargePeriod;
begin
  $L := \{l \mid u[l] = u[k]_{min}$ **and** $y[l] = y[k]_{max}\}$;
  $M := \{m \mid \delta u_l^m > 0, \forall l \in L\}$;
  $K := \{k \mid u[k] < u[k]_{max}$ **and** $y[k] < y[k]_{max}$ **and** $k \notin M\}$;
  $k_c := \min_{k \in K} (S_k)$;
  if there are more than one such k
    then take the first one as k_c ;
end ;

procedure P-UIncrease;
begin
  $\mu_1 := u_{max} - u[k_c]$;

  $\mu_2 := \frac{1}{d_u[k_c]} \cdot (y_{max} - y[k_c])$;

  $L_3 := \{ l \mid k_c \leq l \leq N$ **and** $\delta y_l^{k_c} \neq 0\}$;
  $\mu_3 := \min\, (\frac{y_{max} - y[l]}{\delta y_l^{k_c}}, l \in L_1)$;
  $m_3 := \{ l \mid \mu_3 = \frac{y_{max} - y[l]}{\delta y_l^{k_c}}\}$;
  if there are more than one such l
    then take the first l as m_3 ;

  $L_4 := \{ l \mid k_c \leq l \leq N$ **and** $\delta u_l^{k_c} \neq 0\}$;
  $\mu_4 := \min\, (\frac{u_l}{\delta_l^{k_c}}, l \in L_2)$;
  $m_4 := \{ l \mid \mu_4 = \frac{u_l}{\delta u_l^{k_c}}\}$;
  if if there are more than one such l
    then take the first l as m_4 ;

  $\alpha := \min\, (\mu_1, \mu_2, \mu_3, \mu_4)$;
end ;

procedure P-NewSolution

procedure P-NewSolution;
begin
  $\mathcal{U}_N = \mathcal{U}_N + \alpha \cdot \mathbf{F_u}^{k_c}$;
  $\mathcal{Y}_N = \mathcal{Y}_N + \alpha \cdot \mathbf{F_y}^{k_c}$;
end ;

procedure P-AdjustFVectors

```
procedure P-AdjustFVectors;
begin
   call P-Exclude (k_c);
   ( comment k_l is the last period where a constraint is met );
   if α = μ_1 then
   begin
      k_l := k_c;
      continue ;
   end ;
   if α = μ_2 then
   begin
      k_l := k_c;
      if u[k_l] > 0 and EX_{k_l} = 1 then
      begin
         Ex_{k_l} := 0;
         call P-Include (k_l);
      end ;
   end ;
   if α = μ_3 then
   begin
      k_l := m_3;
      if u[k_l] > 0 and Ex_{k_l} = 1 then
      begin
         EX_{k_l} := 0;
         call P-Include (k_l);
      end ;
   end ;
   if α = μ_4 then
   begin
      k_l := m_4;
      call P-Exclude (k_l);
   end ;
end ;
```

procedure P-MSUpdate

procedure P-MSUpdate;
begin
  for $i := 1$ to N **do**
  begin
    $s_i := s_i - \delta u^i_j(-p[j] + \sum_{k=j+1}^{N} p[k]\delta u^j_k)$;
  end ;
end ;

Appendix B

ALGORITHM FOR THE 2R1C MODEL

B.1 Introduction

The 2R1C (two resistances and one capacitance) model applies to the heating and cooling of simple buildings. The algorithm presented in this appendix is written in APL, and uses a slightly different notation, but it is based on the general algorithm of the previous appendix. The version included here has been extensively debugged.

For this model, many of the calculations become very simple. The state equations are:

$$X[k+1] = aX[k] + b_u U[k] + b_w W[k] \tag{B.1}$$

$$Y[k] = cX[k] + d_u U[k] + d_w W[k] \tag{B.2}$$

The impulse response vector for the first period is:

$$\mathbf{I}^1 = \begin{bmatrix} d_u \\ cb_u \\ cab_u \\ \vdots \\ ca^{N-2}b_u \end{bmatrix} \tag{B.3}$$

and the elemental direction vector has a very simple form. If β is defined as:

$$\beta = \frac{cb_u}{d_u} \tag{B.4}$$

Then, the elemental direction vector becomes:

$$\mathbf{J}^1 = \begin{bmatrix} 1 \\ -\beta \\ -\beta(a-\beta) \\ \vdots \\ -\beta(a-\beta)^{N-2} \end{bmatrix} \tag{B.5}$$

In fact, feasible direction vector of any period can be constructed from the feasible direction vector of the first period using a proportionality factor. Thus, both the calculations and the storage requirements for the algorithm reduce considerably.

The names of the main procedures start with **H-**.

B.2 APL Code for the 2R1C Model

```
      ∇H∆RUN∆2R1C[⎕]
    H∆RUN∆2R1C
    ⍝.....
    ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
    ⍝.....MASTER FUNCTION (2R1C MODEL)
    ⍝.....
     Z∆TIMESTART
    ⍝.....
     H∆INPUT∆2R1C ⋄ H∆COEF∆2R1C ⋄ H∆SETVARS∆2R1C
     H∆PRICES
     H∆INITSOLN∆2R1C
     H∆IVECTORS∆2R1C ⋄ H∆JVECTORS∆2R1C ⋄ H∆FVECTORS∆2R1C
     H∆MARGSAVE
    NEXT: ⍝.....CONTINUE
     UHISTORY ⍝.....TYPE OUT HISTORY
     H∆SELECT
     →FINISH×⍳(CHECKVALUE=0)
     →JUMP2×⍳(HISTORY[CPERIOD]≠1)
     KK←CPERIOD ⋄ H∆EXCLUDE KK
     H∆VT3 ⋄ H∆MSUPDATE KK
    JUMP2: ⍝.....CONTINUE
     H∆UINCREASE
     →NEXT×⍳(NEXTC=1)
     H∆NEWSOLN ⋄ H∆ADJUSTF ⋄ H∆VT3 ⋄ H∆MSUPDATE KK
     →NEXT
    FINISH: ⍝.....CONTINUE
     H∆COMPUTE ⋄ H∆RESULTS
     '...H∆RUN∇2R1C FINISHED...'
    ⍝.....
     Z∆TIMEDONE
    ∇
      ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆INPUT∆2R1C[⎕]
[0]  H∆INPUT∆2R1C
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN, 12/16/87, LAST 2/1/89
[3]  ⍝.....SETS OR ASKS FOR THE 2R1C MODEL PARAMETERS
[4]  ⍝.....
[5]   ⎕TCFF
[6]   'HOW MANY HOURS IN THE TIME HORIZON (24.0 HRS) ?' ⋄ HOURS←⎕
[7]   'HOW MANY MINUTES PER PERIOD (60.0 MINS) ?' ⋄ MPP←⎕
[8]   N←⌊(HOURS×(60÷MPP)) ⋄ DELT←MPP÷60 ⍝.....ASSUMES N INTEGER.....
[9]   'WHAT IS HAE (0.300 KW/C) ?' ⋄ HEA←HAE←⎕
[10]  'WHAT IS HAI (0.500 KW/C) ?' ⋄ HIA←HAI←⎕
[11]  'WHAT IS CI (2.00 KWH/C) ?' ⋄ CI←⎕
[12] ⍝.....
[13]  'WHAT IS UKW (MAXIMUM POWER RATING, EX. 6.0 KW ) ?' ⋄ UKW←⎕
[14]  UMAX←UKW×DELT ⍝....UMAX IN UNITS OF KWH
[15]  'WHAT IS MIN. ACCEPTABLE TEMP. YMIN (18.0 C) ?' ⋄ YMIN←⎕
[16]  'WHAT IS MAX. ACCEPTABLE TEMP. YMAX (22.0 C) ?' ⋄ YMAX←⎕
[17] ⍝.....
[18]  'WHAT IS INITIAL WALL TEMP. XO (18.0 C) ?' ⋄ XO←⎕
[19]  ORDER←1
[20]  'IF W VARIABLE, YOU MUST INPUT IT DIFFERENTLY...'
[21]  'WHAT IS CONSTANT EXTERNAL TEMP. WO (12.0 C) ? ' ⋄ WO←⎕
[22] ⍝.....
[23]  'WHAT IS NUMERICAL TOLERANCE EPSILON (1E¯15) ?' ⋄ EPSILON←⎕
[24] ∇
      ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆COEF∆2R1C[⎕]
 H∆COEF∆2R1C
 ⍝.....
 ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
 ⍝.....CALCULATES THE COEFFICIENTS
 ⍝.....
  DELT←MPP÷60 ⋄ N←⌊(HOURS÷DELT)
  RAE←REA←1÷HAE ⋄ RAI←RIA←1÷HAI
  A11←1-((DELT×HAI×HAE)÷(CI×(HAI+HAE)))
  BU11←HAI÷(CI×(HAI+HAE))
  BW11←(DELT×HAI×HAE)÷(CI×(HAI+HAE))
  C11←HAI÷(HAI+HAE)
  DU11←1÷(DELT×(HAI+HAE))
  DW11←HAE÷(HAI+HAE)
 ⍝.....
  MU1←(CI÷HAE)÷DELT ⍝.....ALSO MU1←DU11÷BU11=DW11÷BW11
  T2←CI÷HAI ⋄ T3←CI÷HAE
  T1←T2+T3 ⍝.....ALSO TAU←(1÷BW11)÷DELT
  TTEST←(DELT<T2)∧(DELT<T3) ⋄ →NEXT×⍳(TTEST=1)
  '...DELT > TIME CONSTANTS...TERMINATED...' ⋄ →
 ⍝.....
 NEXT:UMAX←UKW×DELT
  G1←A11,BU11,BW11,C11,DU11,DW11,T1,T2,T3
  G1← 1 9 ⍴G1
  A← 1 1 ⍴A11 ⋄ BU← 1 1 ⍴BU11 ⋄ BW← 1 1 ⍴BW11
  C← 1 1 ⍴C11 ⋄ DU← 1 1 ⍴DU11 ⋄ DW← 1 1 ⍴DW11
 ∇
  ⎕TCFF
```

APL Code for the 2R1C Model

```
      ∇H∆SETVARS∆2R1C[⎕]
[0]   H∆SETVARS∆2R1C
[1]   ⍝.....
[2]   ⍝.....BY BAHMAN DARYANIAN, 12/16/87, LAST 2/1/89
[3]   ⍝.....SETS AND INITIALIZES THE PRIMARY VARIABLES
[4]   ⍝.....
[5]    PERIODS←⍳N
[6]    X←X0,N⍴0
[7]    U←N⍴0
[8]    Y←N⍴0
[9]    W←N⍴W0
[10]  ⍝.....
[11]   HISTORY←UHISTORY←YHISTORY←N⍴0
[12]   MS←MSOLD←MSNEW←N⍴0
[13]  ⍝.....
[14]   DELU0←DELY0←DELU←DELY←N⍴0
[15]  ⍝.....
[16]   '...THIS FUNCTION GOOD ONLY FOR CONSTANT W...'
[17]   '...AND IF INITIAL STATE VARIABLE IS GIVEN...'
[18]   ⎕TCLF
[19]  ∇
       ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆PRICES[⎕]
 H∆PRICES
 ⍝.....
 ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
 ⍝.....PRICES FOR HEATING/COOLING CASE STUDIES
 ⍝.....
 ⍝.....CASE A: IMPULSE PRICE PATTERN
  PA←24⍴1 ◇ PA[12]←10
 ⍝.....
 ⍝.....CASE B: STEP PRICE PATTERN
  PB←(12⍴1),(12⍴10)
 ⍝.....
 ⍝.....CASE C: SINGLE PEAK PRICE PATTERN
  PC← 1 1 1 1 1 3 3 5 5 7 7 9 9 10 9 9 7 7 5 5 3 3 2 1
 ⍝.....
 ⍝.....CASE D: DOUBLE PEAK PRICE PATTERN
  PD← 1 1 1 1 1 3 5 7 9 10 9 7 5 3 3 5 7 10 9 7 5 3 3 1
 ⍝.....
  '...FOR N≠24, PRICES AND PERIODS MUST BE ADJUSTED...'
  '...SELECT PRICES: A, B, C, OR D ? ' ◇ PSELECT←⍞
  →JUMP1×⍳(PSELECT≠'A')
  P←PA ◇ →OUT
 JUMP1:→JUMP2×⍳(PSELECT≠'B')
  P←PB ◇ →OUT
 JUMP2:→JUMP3×⍳(PSELECT≠'C')
  P←PC ◇ →OUT
 JUMP3:→JUMP4×⍳(PSELECT≠'D')
  P←PD ◇ →OUT
 JUMP4:'..NO CORRECT PRICES.., PROGRAM TERMINDATED..'
  P←24⍴1 ◇ →
 OUT:'...PRICES SELECTED...' ◇ PSELECT
 ∇
     ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆INITSOLN∆2R1C[⎕]
[0]   H∆INITSOLN∆2R1C;K
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]   ⍝.....INITIAL SOLUTION (2R1C MODEL)
[4]   ⍝.....
[5]   ⍝.....Z∆TIMESTART
[6]    X[1]←X0
[7]    K←1
[8]   BEGIN: ⍝.....CONTINUE
[9]    U[K]←(1÷DU11)×(YMIN-((C11×X[K])+(DW11×W[K])))
[10]   U[K]←U[K]×(U[K]≥0)
[11]   U[K]←U[K]-((U[K]-UMAX)×(U[K]≥UMAX))
[12]   X[K+1]←(A11×X[K])+(BU11×U[K])+(BW11×W[K])
[13]   Y[K]←(C11×X[K])+(DU11×U[K])+(DW11×W[K])
[14]   →NEXT×⍳(Y[K]≥YMIN)
[15]   '...TERMINATED, UMAX BELOW REQUIREMENTS...' ◊ →
[16]  NEXT: ⍝.....CONTINUE
[17]   K←K+1
[18]   →BEGIN×⍳(K≤N)
[19]   UINITIAL←U ◊ YINITIAL←Y
[20]   '... END H∆INITSOLN∆2R1C ...'
[21]  ⍝.....Z∆TIMEDONE
[22]  ∇
```

```
     ∇H∆EQUATION∆2R1C[⎕]
[0]   H∆EQUATION∆2R1C KK;K
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]   ⍝.....FINDS ALL X GIVEN U AND W (2R1C MODEL)
[4]   ⍝.....
[5]   ⍝.....Z∆TIMESTART
[6]    X[1]←X0
[7]    K←KK
[8]   BEGIN: ⍝.....CONTINUE
[9]    X[K+1]←(A11×X[K])+(BU11×U[K])+(BW11×W[K])
[10]   Y[K]←(C11×X[K])+(DU11×U[K])+(DW11×W[K])
[11]  ⍝.....Y[K]←X[K]+MU1×(X[K+1]-X[K])
[12]   K←K+1
[13]   →BEGIN×⍳(K≤N)
[14]   '... END H∆EQUATION∆2R1C ...'
[15]  ⍝.....Z∆TIMEDONE
[16]  ∇
      ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆IVECTORS∆2R1C[⎕]
[0]  H∆IVECTORS∆2R1C
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]  ⍝.....IMPULSE FUNCTION VECTOR (COEFF. METHOD FOR 2R1C MODEL)
[4]  ⍝.....
[5]   YIMP←(C11×BU11),(N-2)⍴A11
[6]   YIMP←×\YIMP
[7]   YIMP←DU11,YIMP
[8]  ⍝.....ALSO CAN FIND IMPFUNC RECURSIVELY FROM EQUATIONS
[9]  ∇
```

```
     ∇H∆JVECTORS∆2R1C[⎕]
[0]  H∆JVECTORS∆2R1C;SKIP;NN;LL;K
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]  ⍝.....FINDS JVECTORS
[4]  ⍝.....COEFFICIENT METHOD (FOR 2R1C MODEL ONLY)
[5]  ⍝.....
[6]   BETA←(C11×BU11)÷DU11
[7]   DELUO←(-BETA),(N-2)⍴(A-BETA)
[8]   DELUO←×\DELUO ⋄ DELUO←1,DELUO
[9]   DELYO←DU11,(N-1)⍴0
[10]  DELU←DELUO ⋄ DELY←DELYO
[11] ∇
     ⎕TCFF
```

APL Code for the 2R1C Model

```
      ∇H∆FVECTORS∆2R1C[⎕]
[0]   H∆FVECTORS∆2R1C;SKIP;NN;LL;K
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]   ⍝.....FINDS F VECTORS (2R1C MODEL)
[4]   ⍝.....
[5]    DELU←DELU0 ⋄ DELY←DELY0
[6]    YHISTORY←1×(Y≥YMAX)
[7]    HISTORY←(YHISTORY=0)∧(UHISTORY=1)
[8]    SKIP←((U≤EPSILON)∨(HISTORY=1))/PERIODS
[9]    NN←⍴SKIP
[10]   →OUT×⍳(NN=0)
[11]   K←1
[12]  AGAIN:LL←SKIP[K]
[13]   DELU←DELU+((LL⍴0),(-DELU[LL])×((N-LL)↑(1↓DELU0)))
[14]   DELY[LL]←DELY[LL]+(-DELU[LL]×DELU0[1]) ⋄ DELU[LL]←0
[15]   K←K+1
[16]   →AGAIN×⍳(K≤NN)
[17]  OUT: ⍝.....CONTINUE
[18]  ∇
      ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆MARGSAVE[⎕]
  H∆MARGSAVE;UT1;UT2;UT3;K;UREMAIN;PREMAIN
  ⍝.....
  ⍝.....BHAMAN DARYNIAN 12/16/87, LAST 2/1/89
  ⍝.....FINDS MARGINAL SAVINGS FOR EACH PERIOD
  ⍝.....
   MS←N⍴0
   K←1
  AGAIN: ⍝.....CONTINUE
   UREMAIN←K↓DELU
   UT1←DELY[K+1]+DU11×(-DELU[K+1])
   UT2←(-DELUO[2])×DU11
   UT3←UT2÷UT1
   PREMAIN←K↓P
   MS[K]←(UT3×(+/(-UREMAIN)×PREMAIN))-P[K]
   K←K+1
   →AGAIN×⍳(K<N)
   MS[N]←-(P[N])
   MSNEW←MS
  ∇
```

```
     ∇H∆SELECT[⎕]
  H∆SELECT
  ⍝.....
  ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
  ⍝.....PICKS THE PERIOD WITH THE HIGHEST MS
  ⍝.....
   TAKE←(Y<YMAX)∧(U<UMAX)
   MSFEILD←TAKE/MSNEW
   PERIODFEILD←TAKE/PERIODS
   PICK←(⌈/MSFEILD)=MSFEILD
   CPERIODS←PICK/PERIODFEILD
   →NEXT×⍳((⍴CPERIODS)>0)
   '...TERMINATED, CPERIODS EMPTY...' ◇ →
  NEXT:CPERIOD←1↑CPERIODS
   CVALUE←MSNEW[CPERIOD]
   CHECKVALUE←CVALUE>0
   UHISTORY[CPERIOD]←1
   YHISTORY[CPERIOD]←1×(Y[CPERIOD]≥YMAX)
   HISTORY[CPERIOD]←(YHISTORY[CPERIOD]=0)∧(UHISTORY[CPERIOD]=1)
  ∇
   ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆EXCLUDE[⎕]
[0]  H∆EXCLUDE KK;UT1;UT2;UT3;OLDDELU;OLDDELY;NEWDELU;NEWDELY
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]  ⍝.....EXCLUDES A PERIOD FROM FVECOTRS
[4]  ⍝.....
[5]   →OUT×⍳(KK=1)
[6]   OLDDELU←DELU ⋄ OLDDELY←DELY
[7]   →JUMP×⍳(KK=N)
[8]   UT1←DELY[KK+1]+(DU11×(-DELU[KK+1]))
[9]   UT2←(-DELU[KK])×(-DELUO[2])×DU11
[10]  UT3←UT2÷UT1
[11]  DELU←DELU+((KKρ0),(UT3×((KK)↓DELU)))
[12]  DELY←DELY+((KKρ0),(UT3×((KK)↓DELY)))
[13] JUMP:DELY[KK]←DELY[KK]+(-DELU[KK]×DELYO[1])
[14]  DELU[KK]←0
[15]  NEWDELU←DELU ⋄ NEWDELY←DELY
[16]  DELCHANGE←NEWDELU-OLDDELU
[17] OUT: ⍝.....CONTINUE
[18] ∇
```

```
     ∇H∆INCLUDE[⎕]
[0]  H∆INCLUDE KK;UT1;UT2;UT3;OLDDELU;OLDDELY;NEWDELU;NEWDELY
[1]  ⍝.....
[2]  ⍝.....BHAMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]  ⍝.....INCLUDES A PERIOD FOR FVECTORS
[4]  ⍝.....
[5]   →OUT×⍳(KK=1)
[6]   OLDDELU←DELU ⋄ OLDDELY←DELY
[7]   UT1←DELY[KK+1]+(DU11×(-DELU[KK+1]))
[8]   UT2←(-DELY[KK]÷DU11)×(-DELUO[2])×DU11
[9]   UT3←UT2÷UT1
[10]  DELU←DELU-((KKρ0),(UT3×((KK)↓(-DELU))))
[11]  DELY←DELY-((KKρ0),(UT3×((KK)↓DELY)))
[12]  DELU[KK]←(-DELY[KK])÷DU11
[13]  DELY[KK]←0
[14]  NEWDELU←DELU ⋄ NEWDELY←DELY
[15]  DELCHANGE←NEWDELU-OLDDELU
[16] OUT: ⍝.....CONTINUE
[17] ∇
     ⎕TCFF
```

APL Code for the 2R1C Model

```
     ∇H∆UINCREASE[⎕]
[0]  H∆UINCREASE
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
[3]  ⍝.....CHECKS TO SEE HOW MUCH U CAN BE INCREASED
[4]  ⍝.....
[5]   NEXTC←0
[6]   →OUT×ι(CPERIOD=N)
[7]   FEILD←CPERIOD↓PERIODS
[8]   UPLUS←(U[FEILD]>EPSILON)∧(HISTORY[FEILD]=0)
[9]   UZERO←(U[FEILD]≤EPSILON)∨(HISTORY[FEILD]=1)
[10]  →OUT×ι((+/UPLUS)=0)
[11]  DELTAU←CPERIOD↓DELU ⋄ DELTAY←CPERIOD↓DELY
[12]  H∆FCPERIOD
[13]  DEL1←DEL2←DEL3←DEL4←UMAX ⍝...INITIALIZES ALL THE DELS...
[14]  DEL1←UMAX-U[CPERIOD]
[15]  →OUT×ι(DEL1≤EPSILON)
[16]  DEL2←(1÷DU11)×(YMAX-Y[CPERIOD])
[17]  →OUT×ι(DEL2≤EPSILON)
[18]  →JUMP×ι((+/UZERO)=0)
[19]  DEL3S←((YMAX-Y[UZERO/FEILD])÷DELTAY[UZERO/ι(N-CPERIOD)])
[20]  DEL3←⌊/DEL3S
[21]  →OUT×ι(DEL3≤EPSILON)
[22] JUMP:DEL4S←U[UPLUS/FEILD]÷(-DELTAU[UPLUS/ι(N-CPERIOD)])
[23]  DEL4←⌊/DEL4S
[24]  →OUT×ι(DEL4≤EPSILON)
[25]  UINC←⌊/(DEL1,DEL2,DEL3,DEL4)
[26]  H∆ECONTEST
[27]  →OUT×ι(ETEST=0)
[28]  →NEXT
[29] OUT:NEXTC←1 ⋄ ETEST
[30] NEXT: ⍝.....CONTINUE
[31] ∇
     ⎕TCFF
```

APL Code for the 2R1C Model

```
      ∇H∆FCPERIOD[⎕]
    H∆FCPERIOD;UT1;UT2;UT3
    ⍝.....
    ⍝.....BAHMAN DARYANIAN  12/16/87, LAST 2/1/89
    ⍝.....UPDATES FVECTOR FOR CPERIOD
    ⍝.....
     UT1←DELTAY[1]+(DU11×(-DELTAU[1]))
     UT2←(-DELUO[2])×DU11
     UT3←UT2÷UT1
     DELTAU←UT3×DELTAU
     DELTAY←UT3×DELTAY
    ⍝.....
    ∇
```

```
      ∇H∆ECONTEST[⎕]
    H∆ECONTEST
    ⍝.....
    ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
    ⍝.....TEST OF ECONOMIC FEASIBILITY
    ⍝.....
     ETEST←(P[CPERIOD]<(+/P[UPLUS/FEILD]×(-DELTAU[UPLUS/⍳(N-CPERIOD)])))
    ⍝.....
    ∇

      ⎕TCFF
```

APL Code for the 2R1C Model

```
∇H∆NEWSOLN[⎕]
H∆NEWSOLN
⍝.....
⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
⍝.....FINDS NEW U AND X GIVEN UINC
⍝.....
⍝..... UBEFORE←U
 U←U+(((CPERIOD-1)⍴0),(UINC×((1),DELTAU)))
 U[(U<0)/PERIODS]←0
 Y←Y+(((CPERIOD-1)⍴0),(UINC×(DU11,DELTAY)))
 YHISTORY←1×(Y≥YMAX)
 HISTORY←(YHISTORY=0)∧(UHISTORY=1) ⍝..FOR ALL PERIODS
 '... END : H∆NEWSLON ...'
∇


∇H∆VT3[⎕]
H∆VT3;VT1;VT2
⍝.....
⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
⍝.....FINDS ARRAY OF VT3, PROPORTINALITY OF F VECTORS
⍝.....
 VT1←DELY+DU11×(-DELU) ⋄ VT1[1]←DELYO[1]
 VT2←(-DELUO[2])×DU11
 VT3←VT2÷VT1
⍝.....
∇
 ⎕TCFF
```

```
      ∇H∆ADJUSTF[⎕]
    H∆ADJUSTF
    ⍝.....
    ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
    ⍝.....ADJUSTS DELU AND DELY DEPENDING ON THE CONSTRAINTS MET
    ⍝.....
     KK←1
     →NEXT1×⍳(UINC≠DEL2)
     KK←CPERIOD
     H∆INCLUDE KK
    NEXT1: ⍝.....CONTINUE
     →NEXT2×⍳(UINC≠DEL3)
     KK←1↑((DEL3=DEL3S)/(UZERO/FEILD))
     →NEXT2×⍳(U[KK]<EPSILON)
     H∆INCLUDE KK
    NEXT2: ⍝.....CONTINUE
     →NEXT3×⍳(UINC≠DEL4)
     KK←1↑((DEL4=DEL4S)/(UPLUS/FEILD))
     H∆EXCLUDE KK
    NEXT3: ⍝.....CONTINUE
    ∇
```

```
      ∇H∆MSUPDATE[⎕]
    H∆MSUPDATE KK;MSOLD;DELS1;DELS2;MSCHANGE
    ⍝.....
    ⍝.....BAHMAN DARYANIAN 12/16/87, LAST 2/1/89
    ⍝.....UPDATES MARGINAL SAVINGS AS DELU CHANGES
    ⍝.....
     →OUT×⍳(KK=1)
     MSOLD←MSNEW
     DELS1←P×DELCHANGE ⋄ DELS1←(-DELS1)
     DELS2←⌽(+\⌽DELS1)
     MSCHANGE←(1↓VT3)×(¯1↓DELS2)
     MSCHANGE←MSCHANGE,0
     MSCHANGE←(KK-1)↑MSCHANGE
     MSCHANGE←MSCHANGE,(N-(KK-1))⍴0
     MSNEW←MSOLD+MSCHANGE
    OUT: ⍝.....CONTINUE
    ∇
      ⎕TCFF
```

APL Code for the 2R1C Model

```
      ∇H∆COMPUTE[⎕]
[0]   H∆COMPUTE
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN, 12/16/87, LAST 2/1/89
[3]   ⍝.....COMPUTES SAVINGS AND LOAD SHIFT
[4]   ⍝.....
[5]    UTOTALOLD←+/UINITIAL
[6]    UTOTALNEW←+/U
[7]    USHIFT←U-UINITIAL
[8]    USHIFTNET←+/USHIFT
[9]    USHIFTRATIO←USHIFTNET÷UTOTALOLD
[10]   COSTOLD←+/P×UINITIAL
[11]   COSTNEW←+/P×U
[12]   SAVENET←COSTOLD-COSTNEW
[13]   SAVERATIO←SAVENET÷COSTOLD
[14]  ∇


      ∇H∆RESULTS[⎕]
[0]   H∆RESULTS
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN, 12/16/87, LAST 2/1/89
[3]   ⍝.....SHOWS RESULTS: SAVINGS AND LOAD SHIFT
[4]   ⍝.....
[5]    ' UTOTALOLD:       ',⍕UTOTALOLD
[6]    ' UTOTALNEW:       ',⍕UTOTALNEW
[7]    ' USHIFTNET:       ',⍕USHIFTNET
[8]    ' USHIFTRATIO :    ',⍕USHIFTRATIO
[9]    ' COSTOLD   :      ',⍕COSTOLD
[10]   ' COSTNEW   :      ',⍕COSTNEW
[11]   ' SAVENET   :      ',⍕SAVENET
[12]   ' SAVERATIO    :   ',⍕SAVERATIO
[13]  ∇
      ⎕TCFF
```

APL Code for the 2R1C Model

```
      ∇Z∆TIMESTART[⎕]
[0]   Z∆TIMESTART
[1]   ⍝.....
[2]   ⍝.....WRITTEN BY ROGER E. BOHN  10/2/85
[3]   ⍝.....REVISED BY BAHMAN DARYANIAN 10/26/87
[4]   ⍝.....
[5]   ⍝.....SEE FUNCTION Z∆TIMEDONE.....
[6]   ⍝.....
[7]    TSTART←⎕AI[2]
[8]   ∇
```

```
      ∇Z∆TIMEDONE[⎕]
[0]   Z∆TIMEDONE
[1]   ⍝.....
[2]   ⍝.....WRITTEN BY ROGER E. BOHN  10/2/85
[3]   ⍝.....REVISED BY BAHMAN DARYANIAN 10/26/87
[4]   ⍝.....
[5]    TDONE←⎕AI[2]
[6]    TRUN←TDONE-TSTART
[7]   ∇
```

```
      ∇Z∆IF[⎕]
[0]   R←A Z∆IF B
[1]   ⍝.....
[2]   ⍝.....FROM WEPLAMDA.AWS BY ROGER E. BOHN 7/10/84
[3]   ⍝.....THIS PROGRAM PERFORMS AN IF ROUTINE
[4]   ⍝.....B IS A ZERO-ONE ARRAY, AND A IS A BRANCH NAME ARRAY
[5]   ⍝.....
[6]    R←B/A
[7]   ∇
      ⎕TCFF
```

Appendix C

ALGORITHM FOR THE 1R1C AND 0R1C MODEL

C.1 Introduction

The 1R1C (storage with losses) and 0R1C (storage without losses) models apply to the water heater. The algorithm presented in this appendix is also written in APL, and again it uses a different notation than the ones used already. Although It is based on the general algorithm presented in this thesis, it has been structurally simplified to bypass many of the steps which are no longer necessary. Again, the reason for the inclusion of this version of the algorithm is that it has been extensively debugged.

The important feature of the 1R1C and 0R1C models is that the output variable and the state variable are identical. If the input at some period is increased, the storage level is increased also. For the 1R1C model, the storage level decays back to its previous value as time goes by. However, it remains constant for the 0R1C model. In both cases, when the input at some future period is decreased, the storage level, is brought back to its previous level. This can be seen from the form of the elemental direction vectors. The state equation is:

$$X[k+1] = aX[k] + b_u U[k] + b_w W[k] \tag{C.1}$$

$$Y[k] = X[k] \tag{C.2}$$

If the input at period 1 is increased by δ_1, then the output at period j is increased by $a^{j-1}b_u\delta_1$. The impulse response vector for the first period takes the following form:

$$\mathbf{I}^1 = \begin{bmatrix} b_u \\ ab_u \\ \vdots \\ a^{N-1}b_u \end{bmatrix} \tag{C.3}$$

The input at period j can be decreased by:

$$\delta_j = \frac{a^{j-i}b_u\delta_i}{b_u} = a^{j-i}\delta_i \tag{C.4}$$

However, the storage levels for periods greater than j return to their previous values. Therefore, the elemental direction vector takes the following simple form for the first period:

$$\mathbf{J}^1 = \begin{bmatrix} 1 \\ a \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{C.5}$$

As a result, all feasible direction vectors are similarly sparse, and the overall program code becomes very simple. However, the number of iterations can be reduced by combining a number of adjacent schedulings together. In other words, instead of scheduling of inputs for adjacent periods, as dictated by the form of the elemental direction vector, we can do the scheduling for any pair of periods within the time horizon.

For a 0R1C model, the algorithm takes the following simple form:

- Find a feasible initial solution at each k for $U[k]$ given $X[0]$ and $W[k]$ by assuming that prices remain constant, and by keeping storage levels as low as possible. Initial $U[k]$ can be found by recursive application of the state equation.

- Find the period with the highest electricity price. If there are more than one, take any one and call it the *discharge* period k_D.
- Find the period with the lowest electricity price. If there are more than one, take the one nearest to k_D and call it the *charge* period k_C.
- Decrease the electricity usage at k_D and at the same time increase it at k_C. Continue until one of the following happens:
  - $U[k_D]$ reaches the U_{min}.
  - $U[k_C]$ reaches the U_{max}.
  - Storage of a period in between reaches one of the bounds Y_{min} or Y_{max}.
- If $U[k_D]$ is still not equal to U_{min}, then search for another period with the next lowest price among the remaining periods, i.e. search for another k_C.
- If no more periods can be found to substitute for the electricity usage at k_D, or if $U[k_D]$ reaches U_{min}, then search for the period with the next highest price, and go through the same steps.
- Algorithm ends after a finite number of steps when all the periods have been considered as a possible k_D period.

For the 1R1C model, the algorithm is similar to the one given above, except that the storage loss discounts the value of charging for the long intervals between k_C and k_D. For example, if the input at period i is increased by one unit, it can be used to decrease the input at period j by a^{j-i}. Thus in evaluating the economics of scheduling for the two periods, instead of comparing the prices $P[i]$ and $P[j]$, we must compare $P[i]$ and $a^{j-i}P[j]$.

The 0R1C algorithm was developed by the author and the proof of the optimality is given in the Reference [D1]. For an industrial application of the algorithm see References [D1] and [D2]. The 1R1C algorithm, as mentioned before, is identical to that of Tsitsiklis [T2].

The APL code presented here is for the 1R1C model. For the 0R1C algorithm the coefficient a is simply unity, and the methodology does not require some of the additional calculations and procedures used in the 1R1C algorithm. However, since it is only a special case of the 1R1C algorithm, it is not included here.

In the following program, the names of procedures start with **W-**.

C.2 APL Code for the 1R1C Model

```
     ∇W∆RUN[⎕]
[0]  W∆RUN
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN  7/4/85, 5/10/86, LAST 2/1/89
[3]  ⍝.....MASTER FUNCTION FOR 1R1C MODEL
[4]  ⍝.....
[5]   W∆PRICES
[6]   W∆WHDATA
[7]   W∆SETVARS
[8]   W∆INITSOLN ◇ W∆IVECTOR
[9]   NNN←1 ◇ NEWCHARGE←0
[10] TEST:HISTORYTEST←(0=HISTORY) ⍝...TEST OF THE END OF RUN
[11]  FINISHED←(0=(+/HISTORYTEST))
[12]  →TERMINATE IF FINISHED
[13] DPOINT:W∆DISCHARGE ⍝...CALLS FUNCTION  DISCHARGE
[14] LPOINT:W∆LEFTFIELD ⍝...CALLS FUNCTION  LEFTFIELD
[15] RPOINT:W∆RIGHTFIELD ⍝...CALLS FUNCTION RIGHTFIELD
[16]  DECISION1←W∆COMBINEFIELDS ⍝...DECISIONS BY COMBINEFIELDS
[17]  →(TEST,LPOINT,CPOINT)IF DECISION1
[18] CPOINT:W∆PADJUST
[19]  DECISION2←W∆CHARGE ⍝...RESULTS FROM SETVALUES, SUBFUNC. OF CHARGE
[20]  →(TEST,LPOINT,CPOINT)IF DECISION2
[21] TERMINATE:'...END OF THE RUN...'
[22] ∇
```

```
     ∇IF[⎕]
[0]  R←A IF B
[1]  ⍝ THIS PROGRAM PERFORMS AN IF ROUTINE......FROM WEPLAMDA.AWS BY REB 7/84
[2]  ⍝ B IS A ZERO-ONE ARRAY, AND A IS A BRANCH NAME ARRAY
[3]   R←B/A
[4]  ∇
     ⎕TCFF
```

APL Code for the 1R1C Model

```
      ∇W∆PRICES[⎕]
[0]  W∆PRICES
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN  1/1/87, LAST 2/1/89
[3]  ⍝.....BASE CASE SPOTPRICES
[4]  ⍝.....
[5]   N←24
[6]   P1← 69 66 83 33 53 46 29 13 92 49 23 96
[7]   P2← 47 22 90 19 25 17 43 33 23 81 51 42
[8]   P←P1,P2
[9]  ∇


      ∇W∆WHDATA[⎕]
[0]  W∆WHDATA
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 1/1/87, LAST 2/1/89
[3]  ⍝.....DATA FOR WATER HEATER
[4]  ⍝.....
[5]   UMAX←15
[6]   YMAX←52
[7]   YINITIAL←YMIN←13
[8]   W←N⍴12
[9]   W1←W
[10]  A11←0.98
[11]  BU11←1
[12]  BW11←¯1
[13]  EPSILON←1E¯15
[14] ∇


      ∇W∆SETVARS[⎕]
[0]  W∆SETVARS
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 1/1/87, LAST 2/1/89
[3]  ⍝.....INITIALIZES THE VARIABLES
[4]  ⍝.....
[5]   PERIODS←⍳N
[6]   U←N⍴0
[7]   Y←YINITIAL,N⍴0
[8]   HISTORY←N⍴0
[9]  ∇
      ⎕TCFF
```

APL Code for the 1R1C Model

```
      ∇W∆INITSOLN[⎕]
[0]   W∆INITSOLN
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN  1/1/87, LAST 2/1/89
[3]   ⍝.....INITIAL U GIVEN Y[1], YMIN, AND W
[4]   ⍝.....
[5]    K←1
[6]   BEGIN:U[K]←(1÷BU11)×(YMIN-((A11×Y[K])+(BW11×W[K])))
[7]    U[K]←U[K]×(U[K]≥0)
[8]    U[K]←U[K]-((U[K]-UMAX)×(U[K]≥UMAX))
[9]    Y[K+1]←(A11×Y[K])+(BU11×U[K])+(BW11×W[K])
[10]   K←K+1
[11]   →BEGIN×ι(K≤N)
[12]   K←1
[13]   INITIALU←U ◊ INITIALY←Y
[14] ∇
```

```
      ∇W∆EQUATION[⎕]
[0]   W∆EQUATION KK
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN 1/1/87, LAST 2/1/89
[3]   ⍝.....
[4]    K←KK
[5]   BEGIN: ⍝.....CONTINUE
[6]    Y[K+1]←(A11×Y[K])+(BU11×U[K])+(BW11×W[K])
[7]    K←K+1
[8]    →BEGIN×ι(K≤N)
[9]    K←1
[10] ⍝.....'END W∆EQUATION'
[11] ∇
```

```
      ∇W∆IVECTOR[⎕]
[0]   W∆IVECTOR
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN    LAST 2/1/89
[3]   ⍝.....IMPULSE RESPONSE VECTOR : ALPHA METHOD (1R1C MODEL)
[4]   ⍝.....
[5]    YIMP←BU11,(N-1)⍴A11
[6]    YIMP←×\YIMP
[7]    YIMP←0,YIMP
[8]   ∇
      ⎕TCFF
```

APL Code for the 1R1C Model

```
     ∇W∆DISCHARGE[⎕]
[0]  W∆DISCHARGE;DHISTORY;DATA;DATAHOURS;DBOOLEAN
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 7/4/85, 5/10/86, LAST 2/1/89
[3]  ⍝.....SELECTS DISCHARGE CANDIDATES STARTING WITH REMAINING
[4]  ⍝.....HIGHEST PRICES. MAY HAVE MULTIPLE CANDIDATES.
[5]  ⍝.....ZERO-ONE VARIABLE HISTORY KEEPS TRACK OF PREVIOUS DISCHARGE AND
[6]  ⍝.....CHARGE. DHISTORY SELLECTS THE REMAINING HOURS. DBOLEAN IS A
[7]  ⍝.....ZERO-ONE ARRAY. ONES STAND FOR NEW DISCHARGE CANDIDATES.
[8]  ⍝.....DVALUES AND DHOURS:PRICES AND HOURS OF SELECTED CANDIDATES.
[9]  ⍝.....
[10]  J←1 ⍝...KEEPS TRACK OF DVALUES AND DHOURS.
[11]  DHISTORY←(0=HISTORY) ⍝...ZERO-ONE ARRAY TO PICK WHAT'S LEFT
[12]  DATA←(DHISTORY/P) ⍝...REMAINING PRICES
[13]  DATAPERIODS←(DHISTORY/PERIODS) ⍝...REMAINING PERIODS
[14]  DBOOLEAN←(⌈/DATA)=DATA ⍝...SEE ABOVE
[15]  DVALUES←DBOOLEAN/DATA ⍝...PICKS THE HIGHEST PRICES
[16]  DPERIODS←DBOOLEAN/DATAPERIODS ⍝...AND THEIR HOURS
[17]  DNUMBER←⍴DVALUES ⍝...HOW MANY PICKED AT SAME PRICE
[18] ∇
     ⎕TCFF
```

```
     ∇W∆LEFTFIELD[⎕]
[0]  W∆LEFTFIELD;LSTORAGE;LHISTORY;LVALUES;LPERIODS;LEFTLIMIT
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 7/4/85, 5/10/86, LAST 2/1/89
[3]  ⍝.....FINDS THE FIELD TO THE LEFT (PAST HOURS) OF CURRENT
[4]  ⍝.....DISCHARGE FOR SELECTION OF CHARGEING CANDIDATES.
[5]  ⍝.....LEFT FIELD EXTENDS TO THE HOUR OF FULL STORAGE.
[6]  ⍝.....
[7]   LHISTORY←(DPERIODS[J]-1)↑HISTORY ⍝...HISTORY TO THE LEFT
[8]   LSTORAGE←(DPERIODS[J]-1)↑(1↓Y) ⍝...STORAGE TO THE LEFT
[9]   LVALUES←(DPERIODS[J]-1)↑P ⍝...VALUES TO THE LEFT
[10]  LCHARGE←(DPERIODS[J]-1)↑U
[11]  LPERIODS←(DPERIODS[J]-1)↑PERIODS ⍝...PERIODS TO THE LEFT
[12]  MAXTEST←(YMAX=LSTORAGE) ⍝...TEST TO SEE IF STORAGE FULL
[13]  MAXPERIODS←MAXTEST/LPERIODS ⍝...HOURS WHEN STORAGE IS FULL
[14]  LEFTLIMIT←¯1↑MAXPERIODS ⍝...NO CHARGE BEFORE LEFTLIMIT
[15]  LVALUES←LEFTLIMIT↓LVALUES ⍝...NEW LIMITED LEFTFIELD
[16]  LCHARGE←LEFTLIMIT↓LCHARGE
[17]  LPERIODS←LEFTLIMIT↓LPERIODS ⍝...AND ITS HOURS
[18]  LHISTORY←LEFTLIMIT↓LHISTORY ⍝...RELEVANT HISTORY ARRAY
[19]  LHISTORY←((0=LHISTORY)∧(LCHARGE<UMAX)) ⍝...WHAT IS REMAINING
[20]  LEFTVALUES←LHISTORY/LVALUES ⍝...PICKS AVAILABLE CHARGING
[21]  LEFTPERIODS←LHISTORY/LPERIODS ⍝,,,FIELD TO THE LEFT
[22] ∇
      ⎕TCFF
```

APL Code for the 1R1C Model

```
     ∇W∆RIGHTFIELD[⎕]
[0]  W∆RIGHTFIELD
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 7/4/85, 5/10/86, LAST 2/1/89
[3]  ⍝.....FINDS THE FIELD TO THE RIGHT (FUTURE HOURS) OF
[4]  ⍝.....CURRENT DISCHARGE FOR SELECTION OF CHARGING
[5]  ⍝.....CANDIDATES. RIGHT FIELD EXTENDS TO THE LEFT OF THE
[6]  ⍝.....HOUR AT WHICH STORAGE IS ZERO.
[7]  ⍝.....
[8]   RIGHTVALUES←RIGHTPERIODS←0⍴0
[9]   →OMIT×⍳(YMIN=Y[DPERIODS[J]+1])
[10]  RHISTORY←(DPERIODS[J])↓HISTORY ◊ RSTORAGE←(DPERIODS[J])↓Y
[11]  RVALUES←(DPERIODS[J])↓P ◊ RPERIODS←((DPERIODS[J])↓PERIODS),(N+1)
[12]  RCHARGE←(DPERIODS[J])↓U
[13]  MINTEST←(YMIN=RSTORAGE) ⍝...TEST TO FIND WHEN STOR IS MIN
[14]  MINPERIODS←MINTEST/RPERIODS ⍝...HOURS AT WHICH STOR IS EMPTY
[15]  RIGHTLIMIT←(1↑MINPERIODS) ⍝...CHARGE ONLY BEFORE RIGHTLIMIT
[16]  →NOADJUST×⍳(RIGHTLIMIT≠0) ⍝...ADJUST RIGHTLIMIT IF ZERO
[17]  RIGHTLIMIT←(N+1) ⍝...MAKES RIGTHFIELD GO ALL WAY
[18] NOADJUST:TAKE←(RIGHTLIMIT-1)-DPERIODS[J] ⍝...LENGTH OF THE FIELD
[19]  RVALUES←(TAKE)↑RVALUES ⍝...NEW LIMITED RIGTHFIELD
[20]  RCHARGE←(TAKE)↑RCHARGE
[21]  RPERIODS←(TAKE)↑RPERIODS ⍝...AND ITS HOURS
[22]  RHISTORY←(TAKE)↑RHISTORY ⍝...RELEVANT HISTORY ARRAY
[23]  RHISTORY←((0=RHISTORY)∧(RCHARGE<UMAX)) ⍝...WHAT REMAINS TO CHARGE
[24]  RIGHTVALUES←RHISTORY/RVALUES ⍝...PICKS AVAILABLE CHARGING,,,
[25]  RIGHTPERIODS←RHISTORY/RPERIODS ⍝,,,FIELD TO THE RIGHT
[26] OMIT:→0
[27] ∇
     ⎕TCFF
```

```
     ∇W∆COMBINEFIELDS[⎕]
[0]   R←W∆COMBINEFIELDS;SF;DECISION11;DECISION12
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN 7/4/85, 5/10/86, LAST 2/1/89
[3]   ⍝.....COMBINES LEFT AND RIGHT FIELDS FOR THE CURRENT DISCHARGE.
[4]   ⍝.....CHECKS TO SEE IF ANY CHARGING IN THE COMBINED FIELD IS
[5]   ⍝.....PHYSICALLY POSSIBLE. THEN GOES TO THE CHARGE FUNCTION TO PICK
[6]   ⍝.....CHARGES. CURRENT DISCHARGE CANDIDATE IS IGNORED IF BOTH
[7]   ⍝.....RIGHT FIELD AND LEFT FIELD ARE NULL, OR IF RIGHT FIELD IS
[8]   ⍝.....NULL AND STORAGE BEFORE THE DISCHARGE CANDIDATE IS ZERO.
[9]   ⍝.....
[10]   R←3⍴0 ⍝...NOTE: J (BELOW) IS A GLOBAL VARIABLE, SET EARLIER
[11]   SF←Y[DPERIODS[J]] ⍝...STORAGE PRIOR TO DISCHARGE
[12]   DECISION11←(0=⍴LEFTPERIODS)∧(0=⍴RIGHTPERIODS)
[13]   DECISION12←(0=⍴LEFTPERIODS)∧(YMIN=SF) ⍝...THE TWO DECISION VARIABLES
[14]   →NODISCHARGE×⍳(DECISION11∨DECISION12) ⍝...IGNORE DISCHARGE DECISION
[15]   FVALUES←LEFTVALUES,RIGHTVALUES ⍝...COMBINE VALUES OF FIELDS
[16]   FPERIODS←LEFTPERIODS,RIGHTPERIODS ⍝...COMBINE HOURS OF FIELDS
[17]   R[3]←1 ◇ →EXIT ⍝...GO TO FUNCTION CHARGE
[18]  NODISCHARGE:HISTORY[DPERIODS[J]]←1 ⍝...KEEP TRACK OF NODISCHARGE
[19]   J←J+1 ⍝...NEXT DISCHARGE (SAME LEVEL)
[20]   →NEWDISCHARGELEVEL×⍳(J>DNUMBER) ⍝...NEXT DISCHARGE (NEW  LEVEL)
[21]   R[2]←1 ◇ →EXIT ⍝...GO TO FUNCTION LEFTFIELD
[22]  NEWDISCHARGELEVEL:J←1 ◇ R[1]←1 ⍝...GO TO FUNCTION DISCHARGE
[23]  EXIT:→0
[24]  ∇
     ⎕TCFF
```

```
     ∇W∆PADJUST[⎕]
[0]  W∆PADJUST;TEST1;TEST2;FIELD1;FIELD2
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 1/1/87, LAST 2/1/89
[3]  ⍝.....ADJUSTS PRICES BECAUSE OF LOSSES
[4]  ⍝.....
[5]   XFVALUES←P[FPERIODS]×YIMP[2+FPERIODS-(1↑FPERIODS)]
[6]   XDVALUES←P[DPERIODS]×YIMP[2+DPERIODS-(1↑DPERIODS)]
[7]  ⍝.....
[8]   ALPHA←A11÷BU11
[9]   TEST1←DPERIODS[J]>FPERIODS
[10]  TEST2←DPERIODS[J]<FPERIODS
[11]  FIELD1←TEST1/FPERIODS
[12]  FIELD2←TEST2/FPERIODS
[13]  FVALUES1←P[FIELD1]÷(ALPHA×YIMP[1+DPERIODS[J]-FIELD1])
[14]  FVALUES2←P[FIELD2]×(ALPHA×YIMP[1+FIELD2-DPERIODS[J]])
[15]  FVALUES←FVALUES1,FVALUES2
[16]  DVALUE←P[DPERIODS[J]]
[17] ∇
      ⎕TCFF
```

APL Code for the 1R1C Model

```
     ∇W∆CHARGE[⎕]
[0]  R←W∆CHARGE
[1]  ⍝ ;CCHOICE;DROPPERIOD;HIS
[2]  ⍝.....
[3]  ⍝.....BAHMAN DARYANIAN 7/4/85, 5/10/86, LAST 2/1/89
[4]  ⍝.....PICKS CHARGE PERIODS ONE BY ONE FROM COMBINEFIELDS
[5]  ⍝.....
[6]   R←3⍴0
[7]  PICK:CCHOICE←FVALUES⍳(⌊/FVALUES) ⍝...PICK LOWEST VALUED PERIOD
[8]   CVALUE←FVALUES[CCHOICE] ⋄ CPERIOD←FPERIODS[CCHOICE]
[9]   →CONTINUE×⍳(DVALUE>CVALUE) ⍝...TEST ECONOMIC FEASIBILITY
[10]  HISTORY[DPERIODS[J]]←1 ⋄ →NEXTDISCHARGE ⍝...NOT TO BE  SELECTED AGAIN
[11] CONTINUE: ⍝...GO TO  SETVALUE FUNCTIONS
[12]  →COND2×⍳(DPERIODS[J]<CPERIOD)
[13] COND1:W∆SETVALUES1
[14]  →JUMP
[15] COND2:W∆SETVALUES2
[16] JUMP: ⍝.....CONTINUEW
[17] ⍝...PICK AGAIN..BUT FIRST CHECK TO SEE IF ANY MORE CHARGE CANDIDATE
[18] ⍝...IS LEFT ,DROP THE PICKED CHARGE HOUR FROM THE FIELD. ANY MORE LEFT ?
[19]  DROPPERIOD←FPERIODS⍳CPERIOD ⋄ HIS←(⍴(FPERIODS))⍴1 ⋄ HIS[DROPPERIOD]←0
[20]  FVALUES←HIS/FVALUES ⋄ FPERIODS←HIS/FPERIODS
[21]  →PICK×⍳((⍴FPERIODS)>0)∧(HISTORY[DPERIODS[J]]≠1) ⍝...ANY DISCHARGE LEFT?
[22] NEXTDISCHARGE:J←J+1 ⍝...NEXT DISCHARGE,SAME PRICE
[23]  →NEWDISCHARGELEVEL×⍳(J>DNUMBER) ⋄ R[2]←1 ⋄ →EXIT ⍝...GO TO LEFTFIELD
[24] NEWDISCHARGELEVEL:J←1 ⋄ R[1]←1 ⍝...OR TO NEW DISCHARGE LEVEL
[25] EXIT:→0
[26] ∇
      ⎕TCFF
```

<u>**APL Code for the 1R1C Model**</u>

```
     ∇W∆SETVALUES1[⎕]
[0]   W∆SETVALUES1
[1]   ⍝.....
[2]   ⍝.....BAHMAN DARYANIAN 1/1/87, LAST 2/1/89
[3]   ⍝.....
[4]    ALPHA←A11÷BU11
[5]    CDPERIODS←DPERIODS[J]-CPERIOD
[6]    CHECKPERIODS←CPERIOD+⍳CDPERIODS ⋄ CHECKSTOR←Y[CHECKPERIODS]
[7]    DEL1←UMAX-U[CPERIOD]
[8]    DEL2←U[DPERIODS[J]]÷(ALPHA×YIMP[CDPERIODS+1])
[9]    DEL3←⌊/((YMAX-CHECKSTOR)÷YIMP[1+⍳CDPERIODS]) ⍝...Y MAX ?
[10]   NEWCHARGE←⌊/(DEL1,DEL2,DEL3) ⍝...MINIMUM OF DEL1,DEL2,DEL3
[11]   U[CPERIOD]←U[CPERIOD]+NEWCHARGE
[12]   U[DPERIODS[J]]←U[DPERIODS[J]]-NEWCHARGE×(ALPHA×YIMP[CDPERIODS+1])
[13]   →NEXT×⍳(U[DPERIODS[J]]>EPSILON)
[14]   HISTORY[DPERIODS[J]]←1
[15]  NEXT:CHANGE←NEWCHARGE×YIMP[1+⍳CDPERIODS] ⍝...CHANGE IN STORAGE
[16]   Y[CHECKPERIODS]←Y[CHECKPERIODS]+CHANGE ⍝...UPDATE STORAGE
[17]  ∇
```

APL Code for the 1R1C Model

```
     ∇W∆SETVALUES2[⎕]
[0]  W∆SETVALUES2
[1]  ⍝.....
[2]  ⍝.....BAHMAN DARYANIAN 1/1/87, LAST 2/1/89
[3]  ⍝.....
[4]   ALPHA←A11÷BU11
[5]   CDPERIODS←CPERIOD-DPERIODS[J]
[6]   CHECKPERIODS←DPERIODS[J]+⍳CDPERIODS ◇ CHECKSTOR←Y[CHECKPERIODS]
[7]   DEL1←UMAX-U[CPERIOD]
[8]   DEL2←U[DPERIODS[J]]×ALPHA×YIMP[CDPERIODS+1]
[9]   DEL3←(⌊/(CHECKSTOR-YMIN)÷YIMP[1+⍳CDPERIODS])×ALPHA×YIMP[CDPERIODS+1]
[10]  NEWCHARGE←⌊/(DEL1,DEL2,DEL3) ⍝...MINIMUM OF DEL1,DEL2,DEL3
[11]  U[CPERIOD]←U[CPERIOD]+NEWCHARGE
[12]  U[DPERIODS[J]]←U[DPERIODS[J]]-(NEWCHARGE÷(ALPHA×YIMP[CDPERIODS+1]))
[13]  →NEXT×⍳(U[DPERIODS[J]]>EPSILON)
[14]  HISTORY[DPERIODS[J]]←1
[15] NEXT:CHANGE←(NEWCHARGE÷(ALPHA×YIMP[CDPERIODS+1]))×YIMP[1+⍳CDPERIODS]
[16]  Y[CHECKPERIODS]←Y[CHECKPERIODS]-CHANGE ⍝...UPDATE STORAGE
[17] ∇
      ⎕TCFF
```

Bibliography

[A1] Agnew, C. E., "The Dynamic Control of Congestion-Prone Systems Through Pricing", Ph.D. dissertation, Stanford University, 1973.

[B1] Bazaraa, M. S., and Jarvis, J. J., *Linear Programming and Network Flows*, John Wiley & Sons, New York, NY, 1977.

[B2] Bertsekas, D. P., *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1987.

[B3] Bischke, R. F., and Sella, R. A., "Design and Controlled Use of Water Heater Load Management", *IEEE Transactions on Power Apparatus & Systems*, June 1985, pages 1290-1293.

[B4] Bohn, R. E., "Industrial Response to Spot Electricity Prices: Some Empirical Evidence", MIT Energy Laboratory Working Paper, MIT-EL 080-016 WP, 1980.

[B5] Bohn, R. E., "A Theoretical Analysis of Customer Response to Rapidly Changing Electricity Prices", MIT Energy Laboratory Working Paper, MIT-EL 081-001 WP, 1981.

[B6] Bohn, R. E., Caramanis, M. C., and Schweppe, F. C., "Optimal Spot Pricing of Electricity: Theory", MIT Energy Laboratory Technical Report, MIT-EL 81-026, 1981.

[B7] Bohn, R. E.,"Spot Pricing of Public Utility Services", Ph.D. dissertation, MIT. Also available as MIT Energy Laboratory Technical Report, MIT-EL 82-031, 1982.

[B8] Bohn, R. E., Caramanis, M. C., and Schweppe, F. C., "Optimal Pricing in Electrical Networks Over Space and Time," *The Rand Journal of Economics*, Vol. 15, No. 3, Autumn 1984, pages 360-376.

[B9] Bohn, R. E., Schweppe, F. C., and Caramanis, M. C., "Using Spot Pricing to Coordinate Deregulated Generators, Customers and Utilities", in J. Plummer (ed.), *Electric Power Strategic Issues*, Arlington, VA : Public Utility Reports, 1983.

[B10] Bohn, R. E., Golub, B. W., Tabors, R. D., and Schweppe, F. C., "Deregulating the Generation of Electricity Through the Creation of Spot Markets for Bulk Power", *The Energy Journal*, Vol. 5, No. 2, pages 71-91, 1984.

[B11] Bradley, S. P., Hax, A. C., and Magnanti, T. L., *Introduction to Mathematical Programming*, Addison- Wesley Publishing Co., Reading, MA., 1977.

[C1] Caramanis, M. C., "Investment Decisions and Long-Term Planning Under Electricity Spot Pricing", *IEEE Transactions on Power Apparatus and Systems*, Vol. 101, No. 12, 1982.

[C2] Caramanis. M. C., Bohn, R. E., and Schweppe, F. C., "Optimal Spot Pricing: Practice and Theory", *IEEE Transactions on Power Apparatus & Systems*, Vol. 101, No. 9, 1982.

[C3] Caramanis, M. C., Tabors, R. D., and Stevenson, R., "Utility Spot Pricing: Wisconsin", MIT Energy Laboratory, MIT-EL 82-025, June 1982.

[C4] Chang, C. S., "A Fast Interactive Optimal Load Management Algorithm Under Time-Varying Tariffs", *Electric Power & Energy Systems*, Vol. 10, No. 4, October 1988.

[C5] Constantopoulos, P., Larson, R. C., and Schweppe, F. C., "Electric Load Management by Consumers Facing a Variable Price of Electricity", 1982 Joint National Meeting, San Diego, California, October 26, 1982.

[C6] Constantopoulos, P., "Consumer Decision Making with a Variable Price of Electricity", Ph.D. Thesis Dissertation, MIT, 1983.

[D1] Daryanian, B., "The Definition and Application of an Optimal Response Algorithm for Electricity Consumers and Small Power Producers Subject to Spot Prices", Master's Thesis, MIT, May 1986.

[D2] Daryanian, B., Bohn, R. E., and Tabors, R. D., "Optimal Demand-Side Response to Electricity Spot Prices for Storage-Type Customers", Paper presented at the IEEE/PES 1989 Winter Meeting, New York, N.Y., January 29-February 3, 1989.

[D3] David, A. K., "Optimal Consumer Response for Electricity Spot Pricing", *IEE Proceedings*, Vol. 135, Pt. C, No. 5, UK, September 1988.

[D4] Delgado, R. M. "Demand-Side Management Alternatives", *Proceedings of the IEEE*, Vol. 73, No. 10, Oct. 1985, pages 1471-1488.

[E1] *Demand-Side Planning Program: Projects and Products, 1974-1987*, EPRI Special Report, EM-5917-SR, July 1988.

[E2] *Electric Power Research Institute, Electric Water Heating for Single-family Residences: Group Load Research and Analysis*, EPRI EA-4006, Project 1101-1, May 1985.

[E3] *Development of a Residential Load Control Emulator*, EPRI Project RP 2830, Project Manager: Laurence Carmichael.

[F1] Flory, J. E., and Parker, O. L., "Communication and Metering Equipment for Electricity Spot Pricing", paper presented at the IEEE/PAS (IEEE Power Engineering Society) 1984 Summer Meeting, Seattle, Washington, July 15-20, 1984, 84 SM 552-6.

[G1] Gellings, C. W., "The Concept of Demand-Side Management", *Proceedings of the IEEE*, Vol. 73, No. 10, Oct. 1985, pages 1468-1470.

[G2] Gellings, C. W., Rabl, V. A., Matthews, A. J., and Delgado, R., "Status of Demand Side Management in the USA", *IEEE Power Engineering Review*, August 1987, Vol. PER-7, No. 8, pages 3-8.

[G3] Glover, J. D., Schweppe, F. C., "Control of Linear Dynamic Systems with Set Constrained Disturbances", *IEEE Transactions on Automatic Control*, Vol. AC-16, No. 5, October 1971, pages 411-423.

[H1] Hastings, B. F., "Ten Years of Operating Experience with a Remote Controlled Water Heater Load Management Systems at Detroit Edison", T-PAS, July/August 1980, pages 1437-1441.

[I1] Ihara, S., and Schweppe, Fred, "Physically Based Modeling of Cold Load Pickup", *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-100, No. 9, September 1981.

[K1] Karnopp, D. C., and Rosenberg, R. C., *System Dynamics: A Unified Approach*, John Wiley & Sons, Inc., New York, NY, 1975.

[L1] Lee, S. J., and Wilkins, C. L., "A Practical Approach to Appliance Load Control Analysis: A Water Heater Case Study", *IEEE Transactions on Power Apparatus & Systems*, Vol. PAS-102, No. 4, April 1983, pages 1007-1013.

[L2] Ljung, L., *Systems Identification: Theory for the User*, Printice-Hall, Inc., Englewood Cliffs, New Jersey, 1987.

[L3] Luenberger, D. G., *Introduction to Dynamic Systems*, John Wiley & Sons, Inc., New York, NY, 1979.

[L4] Luenberger, D. G., *Linear and Nonlinear Programming*, 2nd edition, Addison-Wesley Publishing Company, Reading, MA, 1984.

[M1] Manichaikul, Y., and Schweppe, Fred C., "Physically Based Load Modeling", 1978 IEEE Summer Power Meeting, paper No. F78 518-3.

[M2] Manichaikul, Y., and Schweppe, Fred C., "Physically Based Industrial Electric Load", *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-99, No. 2, March/April 1980.

[M3] Manichaikul, Y., and Schweppe, Fred C., "Physical/Economic Analysis of Industrial Demand", *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-99, No. 2, March/April 1980.

[M4] Morgan, M. G., and Talukdar, S. N., "Electric Power Load Management: Some Technical, Economic, Regulatory and Social Issues", *Proceedings of the IEEE*, Vol. 67, No. 2, February 1979, pp. 241-313.

[N1] – "New Era for Electric Utilities: Residential Rates Might Rise", *New York Times*, Tuesday, August 11, 1987, page A1.

[N2] – "Non-Utility Electricity Rising", *New York Times*, Wednesday, August 12, 1987, page D1.

[N3] – "Shopping Around for Electric Power", *New York Times*, Thursday, August 13, 1987, page D1.

[O1] O'Rourke, F. P., "A Physically Based Model of the Space Conditioning Load Under Spot Pricing", MIT Master's Thesis, January 1982.

[O2] O'Rourke, P., and Schweppe, F. C., "Space Conditioning Load Under Spot or Time of Day Pricing", *IEEE Transactions on Power Apparatus and Systems*, Vol. 102, May 1983, pages 1294-1301.

[P1] Papadimitriou, C. H., and Steiglitz, K., *Combinatorial Optimization: Algorithms and Complexity*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1982.

[P2] Peddie, R., "Customer/Utility Relations: A new approach to Credit and Load Management Systems", *IEEE Transaction on Power Apparatus and Systems*, Vol. PAS-103, No. 2, August 1983.

[R1] Rabl, A., "Parameter Estimation in Buildings: Methods for Dynamic Analysis of Measured Energy Use", Center of Energy and Environment Studies, Princeton University, Princeton, NJ 08544, June 1987, revised November 1987.

[R2] Ross, S. M., *Introduction to Stochastic Dynamic Programming*, John Wiley & Sons, New York, 1983.

[R3] Ruane, M., " Physically Based Load Modeling", MIT Ph.D. Thesis, 1980.

[S1] Schweppe, F. C., Caramanis, M. C., Tabors, R. D., and Flory J., "Utility Spot Pricing: California", MIT Energy Laboratory Technical Report, MIT-EL 82-044, Dec. 1982.

[S2] Schweppe, F. C., Tabors, R. D., and Kirtley, J. L., et all "Homeostatic Control for Electric Power Usage", *IEEE Spectrum*, 1982.

[S3] Schweppe, F., Constantopoulos, P., and Larson, R., *Studies in Management Science and Systems; Energy Models and Studies*, edited by Burton Dean and Benjamin Leb, North Holland Publishing Company, New York, Volume 9, 1983, pages 273-292.

[S4] Schweppe, F. C., Caramanis, M. C., Tabors, R. D., and Bohn, R. E., *Spot Pricing of Electricity*, Kluwer Academic Publishers, Boston, 1988.

[S5] Schweppe, F. C., Daryanian, B., Tabors, R. D., "Algorithms for a Spot Price Responding Residential Load Controller", presented at the IEEE/PES Summer Meeting, Portland, Oregon, July 24-29, 1988.

[S6] Shearer, J. L., Murphy, A. T., and Richardson, H. H., *Introduction to System Dynamics*, Addison-Wesley Publishing Company, Reading, MA, 1971.

[S7] Sonderegger, R. C. "Dynamic Models of House Heating Based on Equivalent Thermal Parameters", Princeton University, Phd Thesis, 1977.

[S8] Sonderegger, R. C. " Diagnostic Tests Determining the Thermal Response of a House", *ASHRAE Transactions*, Vol. 84, 1978, pt. 1, pp. 691-702.

[T1] TranstexT, TranstexT is a trademark of Integrated Communication Systems of Atlanta Georgia. It is an experimental/commercial system for residential customer communications, a major component of which is spot price based.

[T2] Tsitsiklis, J. N., "Linear Optimization Problems with Dynamic Structure", Bachelor's Thesis, MIT, December 1979.

[V1] Vickrey, W., "Responsive Pricing of Public Utility Services", *Bell Journal of Economics and Management Science*, Vol. 2, No. 1, 1971.

[W1] Wang, D., "Stochastic Modelling and Parameter Estimation of Residential Electric Loads", Ph.D. Thesis Dissertation, MIT, 1984.

[W2] Wilber, J., "A Micro-Processor Based Energy Usage Rescheduler for Electric Utility Load Leveling", Bachelor's Thesis, MIT, 1981.

[W3] Wilson, N. W., Wagner, B. S., and Colborne, W. G., "Equivalent Thermal Parameters for an Occupied Gas-Heated House", *ASHRAE Transactions*, Vol. 91, pt. 2, 1985, pp. 1875-1884.

[W4] Woodard, James, *Electric Load Modeling*, Garland Press , N.Y., 1979; also MIT Ph.D Thesis, 1974.