NYSERDA--95-1

AUTOMATIC CONTROL OF ELECTRIC THERMAL STORAGE (HEAT) UNDER REAL-TIME PRICING

New York State Energy Research and Development Authority

The New York State Energy Research and Development Authority (NYSERDA) is a public benefit corporation created in 1975 by the New York State Legislature. A 13-member board governs NYSERDA.

NYSERDA's primary mission is to carry out a broad program of energy research, development and demonstration projects designed to develop and apply efficient technologies to help ensure that New York has secure and economical future supplies of energy, while protecting environmental values and promoting economic growth.

NYSERDA derives its basic research revenues from an assessment levied on the intrastate sales of New York State's investor-owned electric and gas utilities. Additional research dollars come from limited corporate funds and a voluntary annual contribution by the New York Power Authority.

In its research program, NYSERDA stresses consultation and collaboration with other organizations, including utilities, universities, industries, private engineering and scientific research firms, local governments, and State and Federal agencies. These efforts stretch NYSERDA's limited research funds and ensure the involvement of those who can use the results of the research.

NYSERDA also has responsibility for:

- Managing the 3,300-acre Western New York Nuclear Service Center at West Valley 35 miles south of Buffalo, the site of a former commercial nuclear fuel reprocessing plant and a low-level radioactive waste disposal area. These responsibilities include:
  - Participating in the West Valley Demonstration Project, a joint Federal/State effort to solidify the high-level radioactive wastes left over from the reprocessing operation and to clean up the facilities used.
  - Maintaining the portion of the site not being used in the Demonstration Project, including the shut-down low-level radioactive waste disposal area.
- Issuing tax-exempt bonds to finance facilities for electric and gas utilities and energy projects for private companies.
- Constructing and operating facilities for disposal of low-level radioactive wastes produced in New York State, once the disposal method and site decisions have been made by the State Low-Level Radioactive Waste Siting Commission and approvals have been issued by State regulatory authorities.
- Managing a 365-acre portion of a Superfund clean-up site in Malta, 20 miles north of Albany. Part of the site was once owned by the Federal government. Portions of it have been used by the Federal government and its contractors since the 1940s for activities that have included rocket engine and fuel testing, weapons testing, and space research.

For more information, contact the Technical Communications Unit, NYSERDA, 2 Empire State Plaza, Suite 1901, Albany, New York 12223-1253, (518) 465-6251, ext. 272.

DISCLAIMER

DISCLAIMER

New York State
Energy Research and Development Authority
Two Rockefeller Plaza • Albany, New York 12223-9998
(518) 465-6251 FAX: (518) 432-4630

An Energy Authority Report in Brief

Report: **Automatic Control of Thermal Electric Storage (Heat) Under Real-Time Pricing, Report, 95-1**

Project Manager: **Davetta Montgomery**

Contractor: **Tabors Caramanis & Associates, Inc.**

Cosponsors: **Consolidated Edison Company of New York, Inc; New York State Electric & Gas Corporation; Electric Power Research Institute; and the Empire State Electric Energy Research Corporation**

Background: Real-Time Pricing (RTP), an innovative rate structure expected to encourage energy conservation and demand reduction, is a demand-side management (DSM) initiative.

Objectives: This project's objective was to design, demonstrate, and assess combining RTP with electric thermal storage (ETS) in several commercial buildings in New York State.

R & D Results: RTP/ETS systems were field-tested at three sites in NYSEG's service territory. Project work included optimizing an energy storage control algorithm. The project contractor worked closely with participating utilities to analyze, develop, and prescribe a procedure to implement a real-time rate structure. A simulation model was applied to each selected building to further test control strategies.

Specific control, communication, and RTP calculation hardware was assembled and installed at each site; systems initially used simple forecasting and response algorithms. Systems were tested during the 1989-90 and 1990-91 heating seasons. Several versions of control software were compared.

Compared with time-of-use-based control, under RTP-based control, the utility's actual cost of service for two water-based heat storage ETS units was reduced eight percent. Savings varied by month, depending on the utility's price pattern, the storage size, and the building's heating load patterns. Simulation study results demonstrated that greater savings are possible for thermal storage systems with larger storage and charging capacities. Savings are expected to be higher for mature RTP-based control systems.

Copies Available: A limited number of copies of the full report are available from The New York State Energy Research and Development Authority, 2 Empire State Plaza, Suite 1901, Albany, New York 12223-1253.

AUTOMATIC CONTROL OF ELECTRIC THERMAL STORAGE (HEAT) UNDER REAL-TIME PRICING

Final Report

Prepared for

THE NEW YORK STATE
ENERGY RESEARCH AND DEVELOPMENT AUTHORITY
Project Manager
T. Davetta Montgomery

NEW YORK STATE ELECTRIC & GAS CORPORATION
Project Engineer
Ronald E. Foster II

ELECTRIC POWER RESEARCH INSTITUTE
Senior Project Manager
Laurence A. Carmichael

EMPIRE STATE ELECTRIC ENERGY RESEARCH CORPORATION
Progam Manager
Frank E. Porretto

CONSOLIDATED EDISON COMPANY OF NEW YORK, INC.
Rates Manager
John D. Flood

Prepared by

TABORS CARAMANIS & ASSOCIATES, INC.
Bahman Daryanian
Richard D. Tabors
Roger E. Bohn

1314-EEED-BEA-90

NYSERDA
Report 95-1

January 1995

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

NOTICE

First printing: January 1995

ABSTRACT

Real-time pricing (RTP) can be used by electric utilities as a control signal for responsive demand-side management (DSM) programs. Electric thermal storage (ETS) systems in buildings provide the inherent flexibility needed to take advantage of variations in prices. Under RTP, optimal performance for ETS operations is achieved under market conditions where reductions in customers' costs coincide with the lowering of the cost of service for electric utilities. The RTP signal conveys the time-varying actual marginal cost of the electric service to customers. The RTP rate is a combination of various cost components, including marginal generation fuel and maintenance costs, marginal costs of transmission and distribution losses, and marginal quality of supply and transmission costs.

This report describes the results of an experiment in automatic control of heat storage systems under RTP which were carried out under the sponsorship of the New York State Energy Research and Development Authority, New York State Electric & Gas Corporation (NYSEG), Electric Power Research Institute, Empire State Electric Energy Research Corporation, and Consolidated Edison Company of New York, Inc. The project was performed under an experimental hour-ahead RTP rate on three buildings in New York State Electric & Gas Corporation territory during the winter seasons of 1989-90 and 1990-91.

The report covers the technical and economic issues involved, including a description of the hardware and software setup and the results of the experiment. The results demonstrate the technical and economic feasibility of automatic RTP-based control of ETS systems. Compared to time-of-use- (TOU-) based control, it was shown that under RTP-based control, the utility's actual cost of service for two water-based ETS systems was reduced by 8%. The savings to the utility of RTP-based control of the ETS over TOU-based control constituted more than one-third of the total savings compared to the no storage case (NSC). The savings varied by month, depending on the price pattern and heating load. Results of a simulation study indicate that greater savings are possible for ETS systems with larger storage and charging capacities. It is expected that savings will be higher for a mature RTP-based control system.

Key Words: Real-time pricing (RTP), electric thermal storage (ETS), demand-side management (DSM), energy management system (EMS), customer response, intelligent buildings, automatic control, optimization.

ACKNOWLEDGEMENTS

The experiment was sponsored under New York State Energy Research and Development Authority (Energy Authority) agreement No. 1314-EEED-BEA-90. The project was funded jointly by the Energy Authority (Project Manager, T. Davetta Montgomery), New York State Electric & Gas Corporation (Project Engineer, Ronald E. Foster II), Electric Power Research Institute (Senior Project Manager, Laurence Carmichael), and Empire State Electric Energy Research Corporation (ESEERCO; under ESEERCO agreement No. EP 89-24, "Real-Time Pricing of Electric Thermal Storage," ESEERCO Program Manager, Frank E. Porretto, ESEERCO Project Manager, John D. Flood of Consolidated Edison Company of New York).

We would also like to acknowledge the cooperation and support of Joseph Rizzuto of the Energy Authority; A. Kumar David of Hong Kong Polytechnic; Greg Kern of Ascension Technologies; Steve Carlson of the Fleming Group; Hank Host-Steen, Jeff Crandall, and Joseph Rosenberger of New York State Electric & Gas Corporation; Dick McDonald of Precision Parts Corporation; Mat Milliorn of Hydrokinetix; and Les Norford, Michael Caramanis, Greg Byrd, David Zakon, Brian Chin, and Deanne Samuels of Tabors Caramanis & Associates, Inc.

TABLE OF CONTENTS

FIGURES

TABLES

SUMMARY

OBJECTIVE AND OVERVIEW

The objectives of the real-time pricing (RTP) experiment performed by Tabors Caramanis & Associates, Inc. (TCA), in the winters of 1989-90 and 1990-91 in the New York State Electric & Gas Corporation (NYSEG) service territory were to:

- Develop experimental rates based on NYSEG's hourly generating costs or opportunity costs of generation;

- Develop the mathematical response algorithms for use in RTP-based control of electric thermal storage (ETS) systems in selected commercial buildings in the NYSEG service territory;

- Implement a fully automated hardware and software system for:

    - Communicating hourly prices from NYSEG, weather data from a commercial weather service, and building energy use data from the sites to TCA's central computer in Cambridge, Massachusetts;

    - Applying the RTP response algorithms in the central computer for control of ETS in the individual buildings;

    - Sending charging schedules to on-site controllers in the buildings;

- Calculate utility service costs under RTP, time-of-use (TOU), and no storage case (NSC) behavior; and

- Determine the relationship between storage size and operational savings.

The experiment was successful from four perspectives:

- Real-time prices were generated based on actual hourly data used in the utility control room;

- Fully automated algorithms for RTP-based control were developed and implemented;
- RTP-based control of ETS provided positive savings to the participating utility; and
- Customer comfort was either maintained or improved under RTP-based control.

STRUCTURE OF THE EXPERIMENT

The RTP-based control system developed for the experiment consisted of a central computer in Cambridge, Massachusetts that communicated via modem with on-site integrated data loggers/controllers. The on-site equipment monitored the local thermal conditions and status of the thermal storage system and transmitted the information to the central computer. If there was a power or telephone failure, the on-site equipment returned the thermal storage system to standard time clock operation.

The experiment was performed at three commercial building sites in NYSEG service territory.

- Brewster office building: A 27,000 square feet (sq ft), one-story building in Brewster, New York. One-half of the building was conventional office space, while the other half contained repair facilities, including some high bay areas, for the local NYSEG district. The entire facility was heated and cooled by a sophisticated, multi-function heating, ventilation, and air conditioning (HVAC) system with a dedicated Andover control system. The heat storage was provided by a 12,614-gallon, pressurized, Megatherm water tank. Water heating in the tank was controlled by a dedicated Megatherm control unit using standard logic (outdoor temperature reset and time clock control). Energy withdrawals from the tank to heat the building were controlled by the Andover system.

- Brewster storefront: A 1,100 sq ft shoe store located in a strip mall in Brewster, New York. The store had its own non-pressurized, 400-gallon, hot water tank manufactured by Hydrokinetix. Store heating was controlled by a simple thermostat that was adjusted by the store owner based on the store schedule. The Hydrokinetix tank was controlled by its own simple outdoor reset plus time clock. This site was the smallest of the three in both square footage and energy use.

- Plattsburgh: A 4,000 sq ft office of a small manufacturing company in Plattsburgh, New York. This building used an earth storage system in which electrical resistance heating mats were embedded in sand beneath a concrete slab. The design made control more difficult, since once heat was stored in the sand its release into the building space was entirely uncontrolled. Before the experiment, the Plattsburgh office site was controlled by a time clock and a very simple thermostat embedded in the concrete slab.

The three sites represent two different technologies and two different climates. Each site was monitored for indoor and outdoor temperature and for charging, energy flow, and temperature of the heat storage system. In the Brewster office building, extensive monitoring had been started under a different project. TCA captured this information from the in-place monitoring system and transmitted it to the computer in Cambridge. The Brewster office monitoring system was transformed into an integrated data logger/controller by adding controller software and connections. Integrated data loggers/controllers were installed by TCA at the Brewster storefront and the Plattsburgh site.

The central computer in Cambridge collected data automatically, including contacting the weather service for the next day's forecasts and contacting the individual sites for current temperature and storage system data. Some of this information was an estimate of the total amount of heat the building would need during the next several hours and the following day. Because of the small size of the experiment, and in order to maintain flexibility, minimize cost, and avoid overburdening the utility, it was decided that all the pricing information would be sent to TCA's office once a day, although this information could have been generated and sent to the RTP sites hourly. As a result, the experiment used the previous week's price data as if it had been received in real-time. The use of week-old price data for a current day's price calculation maintained the correct daily profile of costs, and facilitated the automatic retrieval of prices. Use of week-old prices assumes price and weather (heating load) independence. The study of storage size versus savings which was performed later, used same-week prices and heating load data in a simulation, and the order of magnitude of the results was the same as that of the actual experiment, which used the one-week-old prices. Thus, the results indicated lack of correlation between price and weather. It should also be noted that scheduling of usage under RTP depends on the relative level of prices and not their absolute values.

The scheduling algorithms were run in the central computer and the calculated charging schedules were downloaded to the individual sites. The experimental RTP-based control system was developed to be fully automated. However, there would be differences between the control system described in this report and non-experimental, commercial systems. The NYSEG price data, for example, was generated in the control room in real-time but was sent to TCA once daily in order to simplify communications. In a non-experimental, commercial system, the price information would be sent either to each customer's computer controller or to an electronic mailbox. Consolidated Edison Company of New York, Inc. (Con Edison) and Niagara Mohawk Power Corporation (Niagara Mohawk) currently have similar systems for communication of price information. In a non-experimental system, most of the functions performed by the central computer in this experiment would be performed by on-site equipment, including data gathering, calculating the charging schedule, and controlling ETS charging.

RESULTS

The experiment showed that utility savings under RTP-based control at both Brewster sites compared with those under TOU-based operations. Systems under TOU-based control are forced to follow a consistent pattern. They begin to charge at the start of the off-peak TOU period and turn off either when the storage is full or the off-peak period ends. Since TOU rates only roughly approximate utility system costs, lower, off-peak TOU rates reflect the fact that system costs are normally lower at night. RTP more accurately reflects system costs by allowing prices to vary every hour. The total savings for the two Brewster sites under RTP-based control compared to TOU-based control constituted more than one-third of the total savings under RTP-based control compared to NSC. Monthly results vary and depend on the storage size, heating load, and the price pattern. For the two Brewster sites, the RTP-based control reduced the utility's actual cost of service under TOU-based control by 8%. The costs under each mode of operation were evaluated by multiplying the electricity use (in kilowatt hours [KWH]) and the actual utility price at the hour of use. The absolute monetary savings for the Brewster office building were much larger than those for the Brewster storefront, since its absolute size and pre-experiment energy consumption were more than 10 times larger. The savings varied from week to week, partly because the RTP-based control algorithms were improved over the course of the experiment, but also because of variations in prices and heating loads. Utility hourly costs were determined after the fact and reflected the true expense of serving the load. However, the actual control was based on day-ahead and hour-ahead price forecasts. Therefore, a better forecast would result in even higher savings. Using RTP-based control, the two Brewster sites maintained sufficient charge in the storage tanks to guarantee that heat could be delivered when requested; thus, customer comfort was maintained while reducing total operating costs.

At the Plattsburgh site, the slow dynamics of the storage system and the inability to control the withdrawals of thermal energy required that the charging schedule be based on total hours of daily charging, as well as on prices. There was little cost savings, but customer comfort increased significantly, based largely on using temperature and building energy load forecasts to develop a charging schedule that brought the desired level of thermal energy to the surface of the floor's concrete slab.

The objective of the experiment was to decrease the economic cost of operation without changing customer comfort or level of service. Consequently, no attempt was made to reduce energy consumption, but only to reschedule it to less expensive times. However, if it could be shown that there is a direct relationship between a utility's real-time prices and the heat-rate of its primary fuel for generation, then the thermal storage system operated under RTP could also be considered an energy conservation device, because it would use electricity produced at maximum efficiency, provided the increase in generation efficiency was not more than off-set by storage losses. However, establishment of such a relationship was beyond the scope of this study.

The RTP ETS experiment operated successfully from January to April 1990, and from January to May 1991. The experiment established the technical feasibility of automated control of thermal storage under RTP. Furthermore, the experiment established that RTP-based control significantly reduces the cost of supplying electricity to ETS units, thereby increasing economic efficiency.

Using the results of the experiment as a base case, operations under different storage sizes were simulated to study the relationship between storage size and operational savings. Results of the simulations indicate that, under RTP-based control there is a clear relationship between savings and storage size. Savings are defined as reductions in cost to the utility of supplying the service, where cost is computed by multiplying the electricity use by the utility's actual hourly price (This definition, however, does not include the cost of increased storage which must be considered in optimizing storage capacity.) Under TOU-based control, increasing storage size does not necessarily increase savings. Under RTP-based control, however, increasing either of the two storage size attribute--thermal capacity or charging rate--increases savings, since larger storage capacity provides opportunities for storing across days, and higher charging rate results in fewer hours of charging needed to fill up the storage. As long as the storage system is capable of meeting the heating load, the RTP-based control gives a more economical performance for both undersized and oversized systems compared to the operations of the same storage system under TOU-based control. Under RTP-based control, storage systems half the size of the systems at the Brewster office building and the Brewster storefront performed more economically compared to the actual storage systems under TOU-based control. At both Brewster sites, increasing the storage size attributes increased the savings.

TERMINOLOGY

Prices

Actual Hourly Prices The actual hourly price of electricity determined after the hour.

Day-Ahead Prices	A set of 24 prices forecasted to be in effect for each hour of the next day, determined in the afternoon of the previous day.
Hour-Ahead Prices	A price forecasted to be in effect at the next hour, determined 15 minutes before the start of the hour.

<u>Optimization</u>

Deterministic Algorithm	The primary optimization algorithm that determines, once every few hours, a storage charging schedule for the given time horizon using the day-ahead prices and the current hour-ahead price.
Adjusting Controller	The secondary algorithm which adjusts, once every hour, the solution of the deterministic algorithm as hour-ahead prices become available.
Time Horizon	The number of hours in the future for which the deterministic algorithm determines a charging schedule. Depending on the time of the day, the extent of the time horizon is between 48 to 72 hours.

<u>Types of Control</u>

RTP-Based Control	Control logic under real-time pricing (RTP) rates, resulting in electricity usage hours which change from day to day.
TOU-Based Control	Time-scheduled control logic which requires a fixed daily start time for the charging of the storage. This is the conventual control logic under time of use (TOU) pricing rates.
NSC-Based Control	No storage case (NSC) designates control for buildings similar to ones in the experiment but without any storage. NSC-based control means that electricity is used directly and concurrently to meet the building's thermal load.

Section 1
INTRODUCTION

Electricity is a unique form of energy in that it is simultaneously produced and consumed and cannot easily be stored. The cost of production at any time depends on generating plants, and especially on the last plant loaded, i.e., the marginal generator. A utility's costs, therefore, vary in real time, as load and unit availability change. Logically, then, prices for inter-utility sales have varied by hour to reflect actual load and generation conditions. In contrast, customers' prices traditionally have been frozen in advance, either at a flat rate or on a fixed and predetermined time-of-day pattern that reflects average costs of generation.

Real-time pricing (RTP) conveys information about the time-varying costs of electric service directly to the customer[1]. When the cost of generation is low, the price charged to the customer is low and usually less than both the average and off-peak prices under time-of-use (TOU) rates. When the marginal cost of generation is high, the prices are high. The result is that customers on a real-time rate are integrated physically and economically into the utility system, paying higher costs at times of expensive generation and reaping the benefits when prices are low. In principle, the price could be determined and communicated to the customer in close to real time--every five minutes, for example--but hourly price blocks are more manageable and correspond to current utility operating practices. Hourly prices can be determined in two ways:

- Estimated and set a day ahead (day-ahead pricing), or
- Determined a few minutes ahead of the hour and then communicated to the customer for the next hour (hour-ahead pricing).

One key to realizing net benefits in using RTP is the nature of the customer load. A customer with a large and price-responsive load (i.e., a load which can be turned on and off for several hours if the cost incentive is sufficiently large) would find RTP highly attractive. A smaller and less flexible customer might find that it is not cost-effective to manage its electric load against RTP, since some fixed costs are needed to set up the RTP system at the customer's site. The characteristics of the load are also important to the

[1] The concept goes back to Vickrey, W., "Responsive Pricing of Public Utility Services," Bell Journal of Economics and Management Science, Vol. 2, No. 1, Spring 1971, pp. 337-346. The most extensive discussion of how to calculate real-time prices is in Schweppe, F. C., Caramanis, M. C., Tabors, R. D., and Bohn, R. E., Spot Pricing of Electricity, Kluwer Academic Publishers, Boston, MA, 1988. Several European countries have had limited versions of real-time prices, usually at two or three levels, for years.

customer's ability to respond. Manufacturing customers with the ability to store either the intermediate or final product would find RTP advantageous because it would allow them to produce when prices are low and perform other tasks when prices are high. Commercial or residential customers with electric thermal storage (ETS, for heating) or thermal energy storage (TES, for cooling) would also find RTP attractive. These customers could develop automated control systems to charge and discharge storage systems that respond to RTP. Customers would receive the same final services (heating or cooling), but would benefit through lower electricity bills.

This report presents the results of an experiment by Tabors Caramanis & Associates, Inc. (TCA), to evaluate the benefits of RTP in ETS applications in the commercial sector. The experiment was performed in the winter heating seasons of 1989-90 and 1990-91 for three commercial heat storage customers in New York State Electric & Gas (NYSEG) service territory. Three types of storage systems were evaluated in two climatic regions. Pressurized and non-pressurized water systems were controlled in Brewster, in the southeastern region of New York State, and an earth storage system was controlled in Plattsburgh, in the colder northern region of the state.

The experiment showed significant benefits to the utility. The RTP-based control decreased the utility's cost of service by allowing it to provide the service at times when its actual hourly costs were lowest. Heating costs were evaluated using the utility's actual hourly marginal costs, the base case being similar buildings with no storage (no storage case [NSC]). Savings for the water-based systems under RTP-based control were shown to increase significantly over those for the same systems operating under TOU-based control. First-season savings were 80% for the larger system and 27% for the smaller system. Second-season savings were 60% for the larger system and 45% for the smaller system. This was accomplished with the building occupants experiencing no change in comfort. Control of the earth storage system resulted in no significant savings but did give increased customer comfort and satisfaction with the heating system compared with that under the TOU-based control system.

The remainder of this section provides an overview of the design and implementation of the experiment. Section 2 describes the real-time price-setting system developed for this experiment. Section 3 describes the control system design and implementation, including a discussion of the control algorithms developed, the software systems implemented, and the hardware used in the experiment. Sections 4 and 5 discuss detailed results of the experiment. Conclusions and overall results are given in Section 6.

STRUCTURE OF THE EXPERIMENT

The RTP experiment involved four principal activities:

- Development of the logic for and implementation of an RTP-based rate for this experiment;
- Development and refinement of algorithms and input data for RTP-based control of ETS systems;
- Implementation and operation of the experiment, with real-time communication and control; and
- Analysis and reporting of results.

These activities were conducted using the following steps:

1. Developing real-time prices and rates for the experiment. This included defining the basis for the rates (e.g., the use of short-run incremental costs, defined in a particular way); the choice of time interval for the rate; and treatment of revenue reconciliation. For this project, the decision was made to use hourly real-time prices, based on the opportunity cost to NYSEG of economy energy transactions through the New York Power Pool (NYPP).
2. Developing prices for the experiment. For this project, utility personnel forecasted prices for each hour. The forecasted hourly prices were then transferred for experimental use with a time lag. The prices used in the experiment were slipped by one week, but treated as though they were received each hour. This method is discussed in sections 2 and 3.
3. Selecting customers (sites). Three sites were chosen based on the type of storage system in use, the willingness of the customer to participate, and location.
4. Determining procedures and algorithms to respond to real-time prices. An optimal algorithm was implemented to maximize the impact and benefits of RTP.
5. Communicating with the sites in real time. The real-time prices and other information such as weather forecasts were communicated on a regular basis from the utility and other off-site locations. A computer at TCA offices in Cambridge, Massachusetts was the central point for communications to and from the sites.
6. Controlling charging of the storage systems. After calculating an optimal charging schedule for the storage systems, the charging schedule was sent by modem to an on-site, unattended data logger/entry management unit at each site.

7. Calculating the bills. RTP requires hour-by-hour metering of customers' loads, with once-monthly bill calculations based on these hourly loads. For this experiment, all billing calculations were simulated. The three customers were actually billed under conventional rates.
8. Evaluating the results. The results of the experiment were evaluated as the experiment proceeded to allow updating and refinement of the control algorithms. Utility costs under RTP-based control were compared with those under TOU-based control and NSC.
9. Determining the relationship between storage size and operational savings. Using the experiment as a base case, storage charging under different storage sizes was simulated and compared to base case results.

PRICE METHODOLOGY AND RESULTING PRICES

Prices for this experiment were designed to reflect NYSEG's opportunity cost of serving the customer hourly. Since NYSEG is heavily involved in economy sales within the NYPP, the price at which it could sell into the pool was taken as the opportunity cost. This price is half-way between NYSEG's own incremental cost for the hour, and the cost to whichever pool member is buying from NYSEG. This price was forecasted ahead of the hour by NYSEG and then cleared at the end of each hour by NYPP. As a result, it was possible for NYSEG to generate hour-ahead forecasts with an acceptable level of additional effort. This hour-ahead price was the basis for the price used in the experiment. In addition to the hour-ahead forecasts and actual prices, each weekday at noon, NYSEG forecasted the hourly prices for the following day. This day-ahead information provided a price forecast for the experiment.

Two adjustments were made to these hour-ahead pool prices. First, since all customers were low-voltage (660 volts or lower) electricity users, an hourly adjustment was made for distribution system losses. Following the theory of spot-pricing of electricity, this was calculated as a percentage of the high voltage price, rather than as an absolute number. The percentage usually varied from 9% to 16%, depending on how heavily loaded NYSEG's distribution system was at that time. Second, an approximate adjustment was made for revenue reconciliation, which consisted of 10 mills, or one cent, added to the price per kilowatt hour (KWH).

It is believed that real-time prices calculated in this way are fundamentally sound and representative of the pricing patterns which a fully approved RTP rate would give. However, for a rate to be approved by the New York State Public Service Commission, various refinements and more precise calculations would be required at several points. The goal in this experiment, however, was to test the operational feasibility of RTP price calculations and response, rather than to develop an approved RTP rate.

The resulting hourly real-time prices showed considerable variation by hour, day, and week. Note that the greater the variation in RTP prices, and the greater the discrepancies between RTP and TOU prices, the greater the benefits of RTP-based control over conventional TOU-based control of the sites. For all of the sites it was feasible to charge the storage during times of low prices, almost entirely avoiding charging during high-priced periods. The lowest prices were generally at night, but the exact hours of lowest prices varied by day. Some nights, for example, had only a few hours of low prices.

For operational reasons, while the prices and price forecasts were recorded by NYSEG control room operators each hour as the data came in, they were transmitted to the experiment only once each weekday, which reduced the burden on the operator's time. Since prices were not received in real time, to maintain proper daily price patterns, all prices were lagged by one week (168 hours) in the experiment, then internally communicated to the optimization software one hour at a time. Thus, to the experiment, they were apparently arriving in real time. In a fully operational system, either the operators would input the prices hourly and they would be transmitted automatically to the sites or the process would be automated through the control system computer to calculate and communicate the hourly prices.

SITES AND ETS SYSTEMS

Although alternating current (AC) electric power is expensive to store as electricity, it is often quite inexpensive to store after conversion into a product such as hot or cold water. ETS uses a physical storage medium, such as water, to store electrically generated thermal energy which will be used for space conditioning. For storage heating, electrical energy is used at night, or whenever the RTP price is low, to heat water in a dedicated tank. This water is later withdrawn as needed to heat the building. Thus, the time at which the utility generates the electricity is divorced from the time the customer uses the heating services.

ETS is well established in the United States for cooling and is also used by some utilities for space heating. Considerable research and commercial development have been conducted on alternative physical technologies for storage. For this experiment two technologies were used: hot water and earth storage.

In contrast, the ETS control mechanisms have received much less attention. With the notable exception of one Electric Power Research Institute- (EPRI-) sponsored project by Honeywell[2], existing control

[2] "Cool Storage Supervisory Controller," EPRI report, 1989. Although this report deals with cooling, many of the principles expounded can be applied to heating.

systems use simple methods for load forecasts and response optimization. For any control system a forecast of tomorrow's building load is needed to determine how much to charge the storage system (i.e., how hot to heat it). This traditionally is accomplished using outdoor temperature reset, which means using the contemporaneous outdoor temperature to forecast the next day's heating load. The decision of when to fill the storage is generally based on simple time clock control, synchronized to the lowest price period of a utility's TOU electrical rate. For example, for the NYSEG customers in this experiment, the storage system can charge from 10 P.M. to 7 A.M. weekdays, and at any time on weekends. This is suboptimal for the utility, since NYSEG's lowest-cost hours vary from day to day. RTP-based control of the storage seeks the actual lowest-cost times whenever they may occur, and charges during these times.

Three heating sites were selected: two single-story office buildings, and one small store in a strip mall. The first site (Brewster office building) was a one-story building in Brewster, New York. One-half of it was conventional office space, and the other half contained repair facilities for the local NYSEG district, including some high bay areas. The entire facility was heated and cooled by a sophisticated multi-function heating, ventilation, and cooling (HVAC) system with a dedicated Andover control system. The heat storage was provided by a pressurized Megatherm water tank. Water-heating in the tank was controlled by a dedicated Megatherm control unit which used standard logic (outdoor temperature reset and time clock control). Withdrawals from the tank to heat the building were controlled by the Andover system. A major reason for selecting this building was that it already had a dedicated data logging/monitoring system, specifically to evaluate the effectiveness of the ETS system. For this experiment, the Andover control system of the tank was bypassed, and the tank was instead controlled by RTP-based control signals transmitted via the data logger.

The second site (Brewster storefront) was a shoe store, one of a row of stores in a strip mall. The store had its own non-pressurized 400-gallon hot water tank by Hydrokinetix. Store heating was controlled by a simple thermostat, adjusted by the store owner based on the store schedule. The Hydrokinetix tank was controlled by its own simple outdoor reset plus time clock. For the experiment, a Trinet data logger/energy management system (EMS) was installed at the site, and RTP-based control signals were transmitted via this device. This site was the smallest of the three, in terms of both square footage and energy usage.

The third site was located in Plattsburgh, New York, which is a colder region. This building was the office of a small manufacturing company. It used an earth storage system in which electrical resistance heating mats were embedded in sand below a concrete slab. During the night or at times of low prices, the mats were turned on to heat the sand. Heat retained in the sand gradually moved up through the slab and into the building. The building also had supplemental electric resistance heaters in the office spaces. The thermal characteristics of this system made control more difficult, since once heat was stored in the sand, its release into the building space was entirely uncontrolled. Conversely, the water-based storage systems had separate control loops for putting energy into storage via the electrical heating elements and removing it via the hot water loop to the building.

Prior to the experiment, the Plattsburgh office site was controlled by a very simple thermostat embedded in the concrete slab, in series with a time signal. Under TOU-based control, the system tended to run all night and most of each weekend. For the experiment, a Trinet data logger/EMS was installed at the site, and RTP-based control signals were transmitted via this device to the electrical mats. The supplemental electric resistance heaters were not controlled. Table 1-1 provides a summary of features of the three sites.

Table 1-1. Summary of the Sites.

Specification	Brewster Office Building	Brewster Storefront	Plattsburgh Office
Size (sq ft)	27,000	1,100	4,000
Activities	Office; repairs	Retail store	Office
Storage type	Pressurized water	Non-pressurized water	Earth storage
Storage size	12,500 gallons	400 gallons	----
Storage capacity	10,000,000 Btu 2,900 KWH	300,000 Btu 88 KWH	Undefined
EMS (pre-RTP)	Andover controls and time clock	Simple thermostat and time clock	Simple thermostat
Storage control logic (pre-RTP)	Outdoor reset + timer	Outdoor reset + timer	Timer only

The type and the size of the original storage systems were influenced by the rebates and incentives from the utility and the TOU rates under which the systems were designed to operate.

SYSTEM DESIGN

The system architecture is shown in Figure 1-1. The system was divided into a number of sections and components. The following operations were repeated every hour or every few hours in rolling fashion:

- Forecasting the building load, based on historical data and a weather forecast;
- Forecasting the real-time price, based primarily on a forecast conducted by NYSEG operators;
- Determining a tentative schedule of tank charging for the following several days, using the above forecasts as input to a deterministic optimization program;
- Adjusting the tentative schedule on an hour-by-hour basis, in response to hour-ahead prices and building conditions. This was accomplished by adjusting the controller algorithm. This hourly schedule was communicated to the energy monitoring system/site controller at the site;
- Turning the heating elements on and off in the storage tanks. This was done by closed loop control, using a temperature sensor inside the tanks (or in the concrete, for the Plattsburgh office) by changing the maximum storage temperature setting on an hourly basis; and
- Sensing the space-conditioning temperature profile and demands made by the occupants of the system without altering them in any way. The existing conventional HVAC control system decided when to use stored energy to heat the building, in response to conventional thermostatic control or time-clock control of a night setback thermostat. Thus, customer comfort is not adversely affected by the RTP system. This separation was not possible in the Plattsburgh earth storage system, due to the inherent linkage between storage and use.

The two-step procedure used to calculate a storage-charging schedule (desired storage temperature trajectory) is important to successful RTP-based control. First, a deterministic linear program is solved to calculate an optimal schedule, assuming the price and load forecasts are completely accurate. The non-simplex algorithm takes advantage of special features of the problem and is much faster than a standard

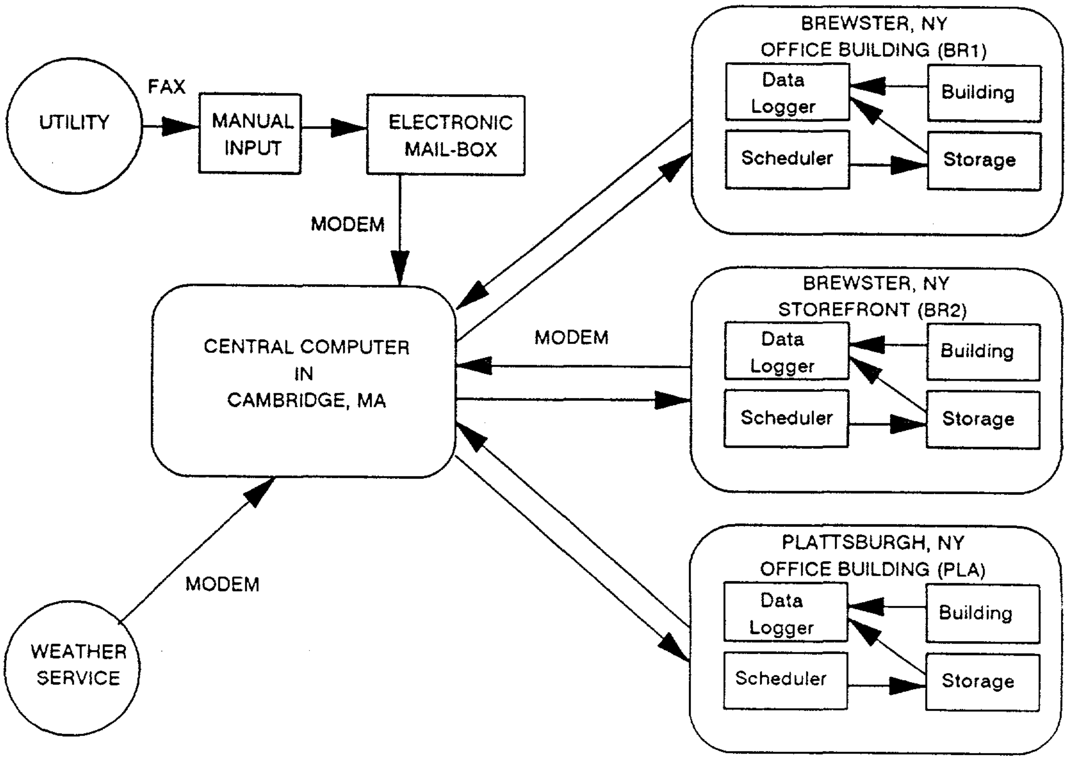

Figure 1-1. Overall System Architecture of the RTP-Based Control System.

linear program[3]. However, as we know, load and price forecasts can never be completely accurate. Therefore, a second-stage adjusting controller algorithm was used to determine, hour-by-hour, whether the charging schedule (or calculated storage temperature trajectory) required readjustment during that hour. If the forecasts were accurate, the algorithm simply gave the same results as the deterministic optimization. More often, the adjusting controller reacted to unexpected changes (either increases or decreases) in the prices and loads by calculating a threshold price and comparing it with the actual hourly price.

Hardware for the project was divided into several components. As discussed, each site required its own data logger, including sensors for tank (or concrete) temperature, water flow rates, and actual KWH consumed. The data loggers were also used to change the maximum storage temperature setting, and thus acted as controllers. All three sites communicated by modem with a single central computer in Cambridge, Massachusetts, which executed the periodic calculations for all three sites. The central computer was an IBM-compatible 386-based PC. Another IBM-compatible 286-based PC was used as a substitute at various times. Prices were received from the NYSEG control room by a facsimile machine, then entered by hand into a computer database.

The hardware architecture described here was well-suited for this small experiment, but different architecture would probably be used in a large-scale RTP project. The computer was centralized because of the need for frequent software changes and testing. The data loggers provided reliable off-the-shelf mechanisms for measurement and control. In a full RTP project, each site could have a dedicated small computer, with an analog-to-digital board for sensing and control. All of this could be embedded in future top-of-the-line EMS's. This would make the site completely autonomous except for the price and weather data, which would be received by modem, radio, pager, or a hybrid system. A single computer at the utility would handle price transmission to each site.

The system was designed to be fail-safe. If telephone communication with the building were lost, the building had preprogrammed instructions covering the following 18 hours. Thereafter it reverted to conventional or TOU-based control. If price and weather forecasts were unavailable, crude forecasts were done based on average historical information.

---

[3] Daryanian, B., Bohn, R. E., and Tabors, R. D., "Optimal Demand-Side Response to Electricity Spot Prices for Storage-Type Customers," IEEE Transactions on Power Systems, Vol. 4, No. 3, August 1989, pp. 897-903.

Section 2

RATE DESIGN AND PRICE FORECASTS

Rate design and price forecasts are critical elements in the design and implementation of RTP-based control of ETS. The basic concept of RTP is that the hour-by-hour cost of electricity reflects the true opportunity cost to society (the marginal cost of generation to the utility) for that hour. The design of the RTP rate used in this experiment to control the ETS units consisted of determining the appropriate marginal cost definition and suggesting an experimental rate. Price forecasts were needed as an input to the control algorithms used to schedule the charging of the ETS systems.

THE THEORY OF RTP

The theory of spot-pricing, the basis of RTP, was developed to communicate the actual short-run marginal costs of electricity generation and transmission from the utility to the customer. The marginal cost of generation varies in real time over a day, a week, and a year, based on a number of factors, including daily demand, weather patterns, unexpected forced outages, and maintenance outages. The short-run marginal cost of the next KWH of electricity generated (and consumed) depends on what generating unit is on the margin, its fuel type and efficiency, and transmission losses. This marginal cost is the first component of the ideal RTP rate.

The ideal real-time price equals the short-term marginal cost of delivering a unit of electric energy at time t. Therefore, the real-time price is defined as[4]:

RTP at time t = marginal generation fuel cost at time t
+ marginal generation and network maintenance cost at time t
+ marginal cost of network losses at time t
+ generation and network quality of supply cost at time t

The first three components--fuel costs, maintenance costs, and costs of losses--are relatively straightforward. The quality of supply can be thought of as that price increase needed to maintain the balance of supply and demand during a shortage. In practice, the quality of supply for generation is estimated as the cost of additional capacity, such as the annualized cost of a gas turbine, distributed across

---

[4] Schweppe, F. C., Caramanis, M. C., Tabors, R. D., and Bohn, R. E., Spot Pricing of Electricity, Kluwer Academic Publishers, Boston, MA, 1988.

only those hours during which there is a probability of loss of load. The quality of supply for the transmission and distribution (network) system is directly analogous. There is usually an additional component of the real-time price imposed by the requirements of the current embedded-cost-based regulatory system. This revenue reconciliation component is designed to account for differences in revenue recovered using an RTP rate and using an embedded-cost rate.

REVENUE RECONCILIATION

Marginal costs show more accented fluctuations than average costs. Under marginal cost pricing, the utility captures more revenue than its average costs, thus enabling it to recover some of its fixed costs. Marginal costs become highest when utility generation resources are strained and when the most expensive units are operating. Therefore, customers who do not respond to price signals and consume electricity at times of high marginal cost contribute a greater amount of revenue to capital recovery. They have to contribute more because their energy consumption during peak periods results in extra capital requirements for the utility in order to meet the peak consumption requirements. Marginal cost pricing acts as a load management signal. Load leveling as a result of marginal cost pricing has a long-term impact on the reduction of reserve margin requirements and the corresponding investment on high-end (more costly) and none-base (less costly) generation.

There are two essential principles or criteria which underlie the pricing decisions of a regulated monopoly:

1. Prices should reflect marginal cost.
2. Prices should recover average total cost, including costs of capital that reflect an acceptable rate of return, averaged over the number of electricity units sold.

The first principle or criterion is a policy prescription meant to reflect conditions in a competitive market, which is clearly most important to the regulator. The second is a financial prescription that ensures the financial viability of the company, which is the condition most vital for the utility. The former determines equitable burden sharing among the utility's customers, while the latter determines the utility's bottom line. Thus, in developing marginal cost-based pricing structures, the second criterion is always taken as given.

If, at the marginal unit of demand, marginal costs are less than average costs (as is almost always the case in a natural monopoly), setting price equal to marginal cost would result in revenue deficit. This corresponds to the case of a positive residual charge (revenue reconciliation adder) between average cost-based recovery and marginal cost-based recovery. If, on the other hand, marginal costs are above average costs, setting price equal to marginal costs would result in a surplus for the utility. Only coincidentally are average costs and marginal costs the same for any period of time.

In one year, marginal costs may recover the full average cost-based revenue requirements, while two years later, there could be a significant shortfall. In a regulated monopoly there is no well-defined correspondence between these two cost concepts.

The residual or revenue reconciliation adder serves two purposes which are absolutely vital to utility profitability:

1. It supplements marginal cost-based revenue to recover total costs.
2. It absorbs the risks associated with the utility's marginal cost based revenue stream.

Once the principle of a residual is solidly established, it is possible to entertain a debate on the magnitude of marginal cost-based recovery versus residual recovery, and the most appropriate method by which the marginal cost component and the residual component may be charged to customers. However, no substantive debate need be entertained on the fundamental principle that the utility needs to recover its average costs (and not its marginal costs) in order to operate as a viable business.

The TOU rate in place at the commencement of the RTP experiment can be used as a basis for determining a revenue reconciliation adder to the RTP rate if it can be assumed that this TOU rate recovered NYSEG's average costs. The adder is the amount added to the RTP rate which results in the same bill as the bill under the TOU rate for a customer who does not change its TOU-based behavior.

In the beginning of the experiment, the revenue reconciliation adder was estimated at 10 mills (1 cent). This value was added to the RTP rates used in day-to-day control of the systems. Analysis of the results after the experiment indicated that the actual revenue reconciliation residual varied from week to week. However, it must be emphasized that the adder has no impact on customer response to prices, because under an RTP-based control algorithm, price differentials, rather than the absolute value of prices, dictate the customer's consumption pattern.

To the extent that distribution costs can be associated with an individual customer (dedicated costs) rather than the system, these costs can be charged to the customer as a constant amount billed every month independent of actual consumption.

RATE DESIGN

The objective in this experiment was to design an RTP rate to meet four conditions:

- The RTPs must elicit the same customer behavior as pure marginal cost rates;
- The rate design must be acceptable to the regulators and be simple and easy to implement and to explain to participating customers;
- Customer-specific revenue neutrality must exist if the customer does not respond to RTP, i.e., does not alter his or her consumption from its pattern under TOU rates [(L@TOU(t))]. More specifically, revenue neutrality requires that:

  $\Sigma_{(\text{over all hours } t)}$ { RTP(t) $\times$ L@TOU(t) } = Customer bill under TOU rates

  where:

  L@TOU(t) = Customer load for heating during hour t under a TOU-based behavior

- If the customer responds to RTP's, the bill must change by an amount that does not exceed the savings (positive or negative) to NYSEG from the customer's behavior under RTP's compared to the behavior and consequent billing costs under TOU rates, i.e.,

  $\Sigma_{(\text{over all hours } t)}$ { RTP(t) $\times$ L@RTP(t) } $\leq$ (Customer bill under TOU rates) - (NYSEG Savings)

  where:

  NYSEG Savings = $\Sigma_{(\text{over all hours } t)}$ { MC(t) $\times$ [L@TOU(t) - L@RTP(t)] }

  L@RTP(t) = Customer load for heating during hour t under real-time prices and TCA storage control logic

  MC(t) = NYSEG marginal cost during hour t

The final RTP rate consisted of hourly marginal costs adjusted by the marginal cost of line losses, plus a constant revenue reconciliation adder[5]. Specifically:

$$RTP(t) = MC(t)_{adj} + (TOU - AMC)$$

where:

$$AMC = \text{Average Marginal Cost}$$

$$MC(t)_{adj} = MC(t) \times (\text{Marginal Losses at Time } t).$$

The rate was expected to satisfy the following objectives:

- Assuming that TOU rate-based charging takes place uniformly during the off-peak hours, the objective of customer-revenue neutrality in the absence of response to the RTP rate would be met;

- The additive adjustment would elicit the same behavior for storage dispatch as pure marginal costs since decisions depend on the cost differential between hours rather than the absolute value of costs; and

- The higher $MC(t)$ values during peak hours assure that no charging would occur during system capacity shortage hours (the RTP during these hours could be as high as 1,000 mills/KWH). Hence a capacity charge would not be needed, nor is a different MC adjustment factor needed for peak hours.

MARGINAL COSTS

Because NYSEG is part of the NYPP, the ideal starting point for calculation of an RTP would be the pool system lambda or resource marginal energy costs, corrected for area reliability constraints. In the absence of a pool lambda, the NYPP-based NYSEG lambda was the best available proxy. This proxy differed from true lambda primarily in the following three ways:[6]

[5] The revenue reconciliation adder was estimated at 10 mills/KWH.

[6] This information was provided by the NYSEG Operations Department.

- NYSEG primarily sells to NYPP during hour t:[7]

$$MC(t) = \quad NYSEG_{F\&OM} + (\text{Average NYPP Buyers}_{F\&OM} - NYSEG_{F\&OM}) \div 2$$

- NYSEG buys primarily from NYPP but is not deficient during hour t:

$$MC(t) = \quad NYSEG_{F\&OM} - (NYSEG_{F\&OM} - \text{Average NYPP Seller}_{F\&OM}) \div 2$$

- NYSEG buys from NYPP and is deficient during hour t:

$$MC(t) = \quad (\text{Average Marginal NYPP Seller}_{F\&OM} \times 1.1) + (3.0 \text{ mills/KWH})$$

DISTRIBUTION LOSSES

Marginal losses were estimated as a linear function of the NYSEG system load ($2 \times B \times$ system load) by assuming that total losses were adequately modeled as a quadratic function of load. The value of the coefficient B used was 0.000037 per megawatt (MW), which corresponds to total losses of 3.7% when the NYSEG system load is 1,000 MW.

FORECASTS

Forecasts of RTP were needed to optimize charging of the ETS systems, since the RTP forecast was based on a forecast of NYSEG's NYPP lambda. While several different forecasting methods were considered, including using a unit commitment model and an historical/statistical model, it was decided that the most accurate forecasts were generated by NYSEG system operators and control room personnel. They were already forecasting prices on a day-ahead basis under contract obligations, using a multi-hour (three or four hours) time period. They agreed to provide forecasts of hourly prices on a day-ahead basis for the RTP experiment. Because of other contract obligations, the NYSEG operators were at times forecasting the costs for the next hour 45 minutes ahead of the hour. This forecast, or quote, was verbal and non-binding and only communicated if the other contracting party requested it. For the experiment, the operators simply recorded this hour-ahead price forecast. An entire day of forecasted hour-ahead prices as well as

---

[7] The marginal costs are based on fuel and operation and maintenance (F&OM) costs of NYSEG, the average seller, or the average buyer.

the actual prices were recorded on a form and sent by facsimile to TCA along with the day-ahead price forecast and the day-ahead hourly load forecast sent to NYPP.

It was not feasible for this experiment to send the forecasts in real time to the central computer because of the additional burden it would have placed on the utility's operations room personnel. (Under full implementation of RTP rates, hourly price can be generated and sent to the RTP sites on an hourly basis.) Therefore, price information was sent to TCA once a day instead of once an hour. In order to preserve price patterns for each day of the week, the experiment was conducted using the previous week's data. The data was entered into a database, and the storage control algorithm simply accessed the previous week's prices for the corresponding day of the week, and simulated real-time price communication such that day-ahead prices were received the previous day, and hour-ahead prices were received in the previous hour.

The use of week-old data for a current day's price determination:

- Maintained the correct daily profile of costs;

- Removed the burden of hourly data communication from the operators, while allowing them to practice one hour-ahead marginal cost forecasting; and

- Did not allow the ETS results to show the true correlation between weather (which influences building load) and prices. This matter was further investigated using simulation, and no significant effect on percent savings was observed, thereby indicating a lack of correlation between weather and prices.

RTP AND FORECAST RESULTS

Actual real-time prices exhibited a great deal of variation throughout the experiment. During the 1989-1990 season, the highest price recorded during the experiment was $99.4/megawatt hour (MWH) ($1/MWH = 1 mill/KWH); the lowest was $15/MWH. The average price was $33.8/MWH. Although the highest prices occurred during the TOU on-peak period (7 A.M. to 10 P.M.), low prices (ranging from $25 to $30/MWH) occurred during this period as well. During the TOU off-peak hours (10 P.M. to 7 A.M., plus Saturday and Sunday), the highest price that occurred was around $75/MWH. Prices on Saturday and Sunday varied almost as much as on weekdays, with weekend prices often as high as their weekday counterparts. The maximum price increase from one hour to the next was $59.7/MWH and the maximum price decrease from one hour to the next was $49.1/MWH. Appendix A contains weekly graphs of hourly real-time prices recorded during the experiment. The graphs show that prices can vary substantially by hour, day, and week.

The large short-term variation in prices was due partly to short-term purchases/sales within NYPP and between NYSEG or NYPP and its neighboring utilities which take place on this time scale. The sales were hour-by-hour deals negotiated by utility system operators to meet loads at the lowest possible costs. One system would offer to sell power to another system at an estimated but non-binding price. The price represented the operator's best guess of the cost of producing electricity by the seller's system. The price actually paid was the actual price calculated by the pool at the end of the hour in which the transaction took place. Other types of short-term sales that affected the hourly price included sales to other pools and transmission transactions.

Analysis of prices indicates that NYSEG's native load was generally not a good predictor of the real-time price. Figure 2-1 shows that price could vary a great deal for a given load. The higher loads displayed larger variations. Note that prices were adjusted to account for losses. This had the effect of increasing prices at every load level, with prices associated with higher loads being adjusted upward slightly more than prices at lower loads.

The monthly price-duration curves, shown in Figure 2-2, indicate that although December 1989 had generally higher prices than January, February, or March of 1990, approximately 50% of January 1990 prices were higher than those of December 1989. Price-duration curves, analogous to load-duration curves, show how often different prices were achieved. December was expected to have higher prices due to extremely cold weather. The low prices on the December curve are partly due to the Christmas holidays.

In the experiment, the hour-ahead forecasts were assumed to be the prices used for billing. The rationale behind hour-ahead forecasts was that customers should know the price of electricity when it is being used. This is the price used for control, although it is only an estimate of the actual cost to the utility. The more accurate the forecast, the more accurately it reflects the actual cost of producing the electricity, and the more economical is the optimal scheduling of ETS charging.

The precision of the forecasts varied. Figure 2-3 shows forecasts versus actual prices for a week. On days when actual prices did not vary widely, both day-ahead and hour-ahead forecasts were close to actual, while on days when prices varied widely, day-ahead forecasts were generally inaccurate. Hour-ahead forecasts intermittently captured the price variation of such days but were less successful for extreme

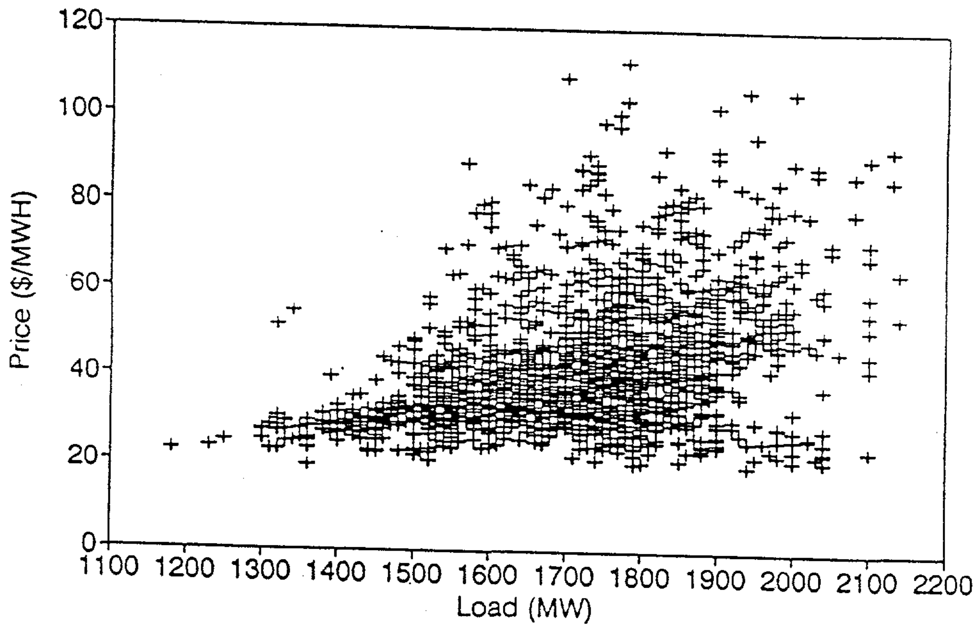

Figure 2-1. NYSEG Load Versus Loss-Adjusted Prices for December 1989 to March 1990.

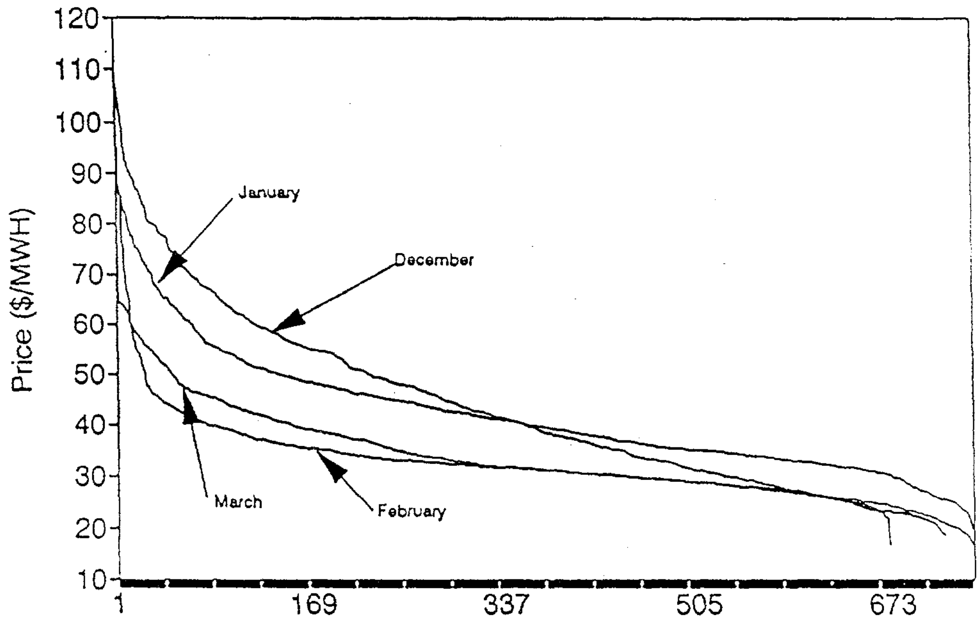

Total Number of Hours at or above a Given Price

Figure 2-2. Monthly Price Durations for NYSEG Loss-Adjusted Prices for December 1989 to March 1990.

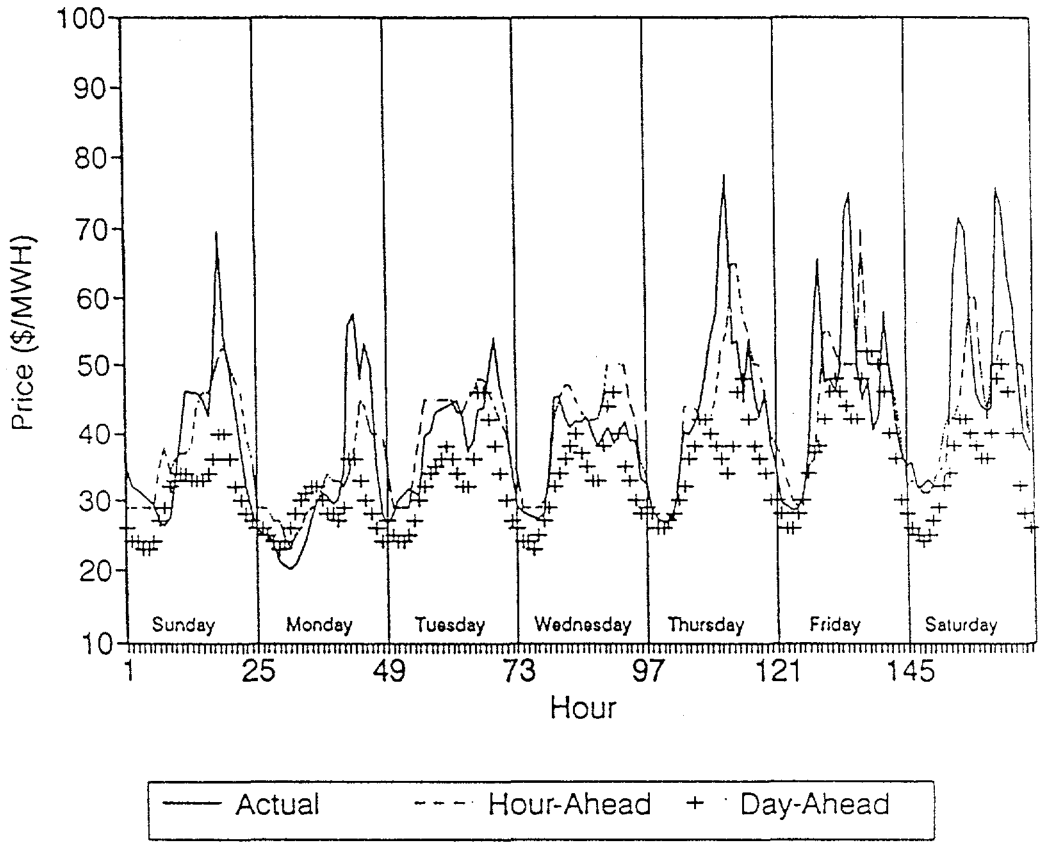

Figure 2-3. Forecast Precision -- Forecasts Versus Actual Prices for a Typical Week.

variations. The difference between hour-ahead forecasts and actual prices during off-peak hours was significant. Since charging was done primarily during off-peak hours, the inaccuracy of hour-ahead forecasts may have produced suboptimal charging. The accuracy of the forecasts is shown in Appendix B in plots of actual prices versus hour-ahead and day-ahead forecasts. The 45-degree line represents a perfect forecast; deviations above and below indicate forecasting errors.

To improve the accuracy of the day-ahead forecasts used to set day-by-day charging plans, more contemporaneous information contained in hour-ahead forecasts was used. The averages of both the current day's hour-ahead and day-ahead prices from midnight to the current hour were computed every four hours. The difference between the two averages was then added to the day-ahead forecast for each hour. Thus, if the day-ahead forecast were lower than the hour-ahead, the entire day-ahead forecast would be shifted up by the average difference, a change that improved the performance of the optimization algorithm.

Calculating the hourly RTP rate was done routinely throughout the experiment. While the experiment used forecasts of prices that were "hidden" for a week, it is clear that the forecasting procedure in the NYSEG control room could be automated to communicate the prices to customers hourly. The professional skills and attitude of the staff at NYSEG's control room contributed to the experiment's success, and proved that implementing an RTP rate based on one-hour-ahead forecasts is possible.

In summary, both day-ahead and hour-ahead prices were used in the optimization algorithm of the RTP-based control system. Day-ahead prices were adjusted upward or downward based on the latest hour-ahead price information. Furthermore, they were extended for the following two days using the most recent price data as a predictor of future price trends. The day-ahead prices in adjusted and extended form were used as input in a *deterministic algorithm* to determine a charging schedule for the next two to three days. The hour-ahead prices were then used in an *adjusting controller algorithm* to adjust the schedule based on the latest price information. The algorithms are described in Section 3.

FORECAST BIAS

In future utility RTP programs, if prices are based on utility forecasts and consistent bias (underestimation or overestimation) is observed in forecasts, a feedback mechanism should be used to eliminate this bias.

The forecast bias can also be viewed as an over- or under-recovery in the customer's bill. A correction can be reflected in the final billing prices as an adder to the actual billing prices. As noted before, customer response depends on the pattern and the relative value of prices, and modification of prices by a positive or negative adder will not change storage customer behavior.

Continuous monitoring of forecast bias and proper correction either during real-time operations or during billing will eliminate the possibility of exaggeration of the RTP benefits for the customer. In this experiment, day-ahead and hour-ahead prices were used to determine control schedules for the storage customers. However, economic performance was evaluated using after-the-fact actual hourly prices. Consequently, for this experiment, a better price forecast could only improve the results, since it would bring the customer response closer to the ideal response to actual prices.

Section 3

CONTROL SYSTEM DESIGN AND OPERATION

This section discusses the architecture of the RTP-based control system designed and operated for this experiment, beginning with a brief description of its general features. It continues with an overview of the operation, which was an interplay of data logging at the three test sites, with the control schedule determined by and transmitted from the central computer at the TCA office. The control system elements implemented in the experiment are described, including the optimization algorithms which determined an economical schedule of storage charging for each building, and the price, load, and weather forecasts used as inputs. The hardware used in the experiment is described at the end of this section.

CONTROL SYSTEM DESIGN

Existing control systems for ETS under TOU rates are based on two simple rules. The first rule allows storage charging only during off-peak periods. This rule is optimal for the building owner, but not for the utility, since TOU rate structures do not reflect true day-to-day and hour-to-hour variations in the cost of electricity. The second rule is to use the current day's outdoor temperature to predict the next day's heating/cooling load. Present outdoor temperature reset mechanisms use a linear relationship to adjust the storage limits for nightly charging according to the outdoor temperature--the lower the temperature, the higher the storage limit is set. Since outdoor temperature provides only a crude forecast of the following day's load, a more economical use of storage requires a more sophisticated load-forecasting method.

To respond optimally to real-time prices, the HVAC control system for a building with thermal storage must have the capability to sense its current status and predict its future heating (or cooling) loads. Based on future prices, the control system then decides on a least-cost schedule of electricity use to satisfy its thermal demands. Optimal response requires daily variations in storage charging periods to account for hourly changes in loads and prices. The utility cannot control the storage activities of numerous customers because daily customer requirements vary depending on occupancy, scheduling of holidays, physical performance of the storage system, and other building-specific factors, some of which cannot be predetermined.

RTP-based control system design relies on locally distributed intelligence, whereby the utility calculates real-time prices and makes these available to the customers, whose computers then sense and optimize the response schedule for their individual buildings. All the data gathering and computational capability resides in the building being controlled. However, for efficiency, the RTP-based controls for this experiment for the three sites in New York State were implemented on a single PC located in Cambridge, Massachusetts. This was useful in an experiment for purposes of monitoring and debugging the software but is feasible where only a few sites are being controlled.

The RTP-based control system was designed to provide supervisory control over the buildings' existing control systems. Price and weather information was gathered electronically, building status was sensed, and a near-optimal storage charging schedule for the building was calculated and activated--all under computer control. The additional computer technology required for RTP-based control beyond the existing building control system were an IBM-compatible 286- or 386-based personal computer (PC), modems for communications, a data logger/controller unit, and several additional sensors to measure the status and performance of the HVAC system in more detail. The elements of the RTP-based control system are shown in Figure 3-1. The control algorithms were run in the central computer at TCA's office in Cambridge. At each site, data were collected continuously, stored by the data logger, and downloaded to the central computer when called. The central computer also received price forecasts from the utility and outdoor temperature forecasts from a commercial weather service. The temperature forecast information was combined with the site data to calculate a building load forecast for the following three days.

The load and price forecasts were used in a two-stage procedure to determine a near-optimal storage charging schedule. The first stage assumed that the forecasts were accurate and determined an optimal storage charging schedule under deterministic conditions. A second stage considered the uncertainty in the future loads and prices, and modified the charging schedule accordingly.

The system was designed to work unattended. During routine operation, no human operator intervened in the continuous automatic operation of the RTP-based control system, except for a once-daily compilation into a data file of the utility price forecasts which were received via facsimile in a non-machine readable form.

The system was also designed to be fail-safe in several respects. For example, if telephone communication with the building was lost, the building had preprogrammed instructions covering the following 18 hours; thereafter it reverted to conventional TOU-based control. If price and weather forecasts were not available, crude forecasts were calculated based on average historical data.

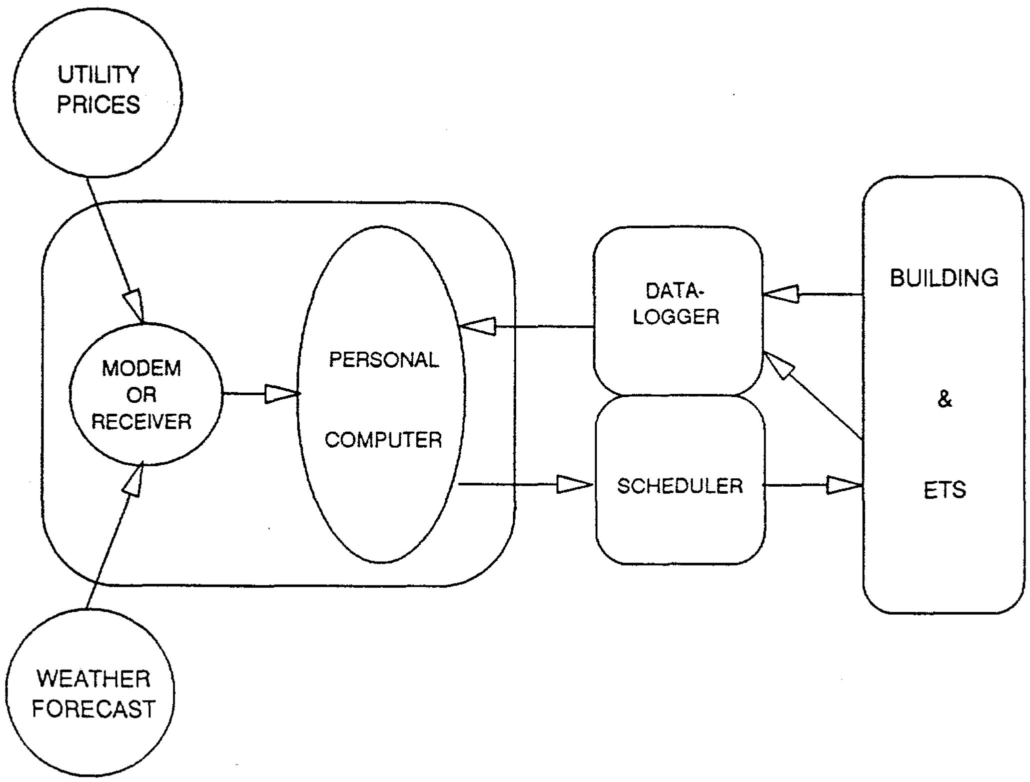

Figure 3-1. Elements of the RTP-Based Control System.

CONTROL SYSTEM IMPLEMENTATION

This control system was implemented for the heat storage systems in three buildings during the winter seasons of 1989-90 and 1990-91. The two water storage systems in Brewster had independent storage and heating loops, so while energy could be put into the tanks under RTP-based control, the hot water was drawn under the normal customer control of the building's heating system. The systems were analogous to cool storage for air conditioning. The physical characteristics of the buildings and their storage systems were not altered for this experiment.

The hourly least-cost storage charging schedule was downloaded into the controller unit of the hardware at the sites. At specified times, the controller reset the storage temperature band limits, thereby activating the storage heating elements, which controlled the timing and amount of storage charging. The space-conditioning temperature profile and heating demands of the occupants were sensed by the system, but were left unaltered. Each building's conventional HVAC thermostat control system determined independently when to draw heat from the storage system in response to conventional thermostatic control (Brewster storefront) or the Andover EMS (Brewster office building). Therefore, the size and pattern of the actual daily heating load were unaltered by the RTP-based control system, and remained under the control of the buildings' occupants. As shown in Figure 3-2, the sole function of the RTP-based control system was timing of storage charging decisions and heating of the tanks.

The elements of RTP-based control were in place by late 1989, and all three sites were under RTP-based control by early January 1990. Before each site was placed under RTP-based control, information about the site was gathered by the data logger. Next, the site was controlled by the central computer under the TOU strategy previously used at the site. This testing allowed the contractor to understand the dynamics of each site and to test the field equipment. The sites remained under RTP-based control, with the control system being continuously improved, through March 30, 1990, when they were returned to TOU-based control. Data were continuously logged from each site until May 5, 1990. Each day, beginning November 1, 1989, and continuing through March 31, 1990, price information was sent via fax from the NYSEG control room to Cambridge. Weather information was received electronically for the duration of the experiment, beginning December 1, 1989. The second-season experiment covered the period from late December 1990 to the end of April 1991.

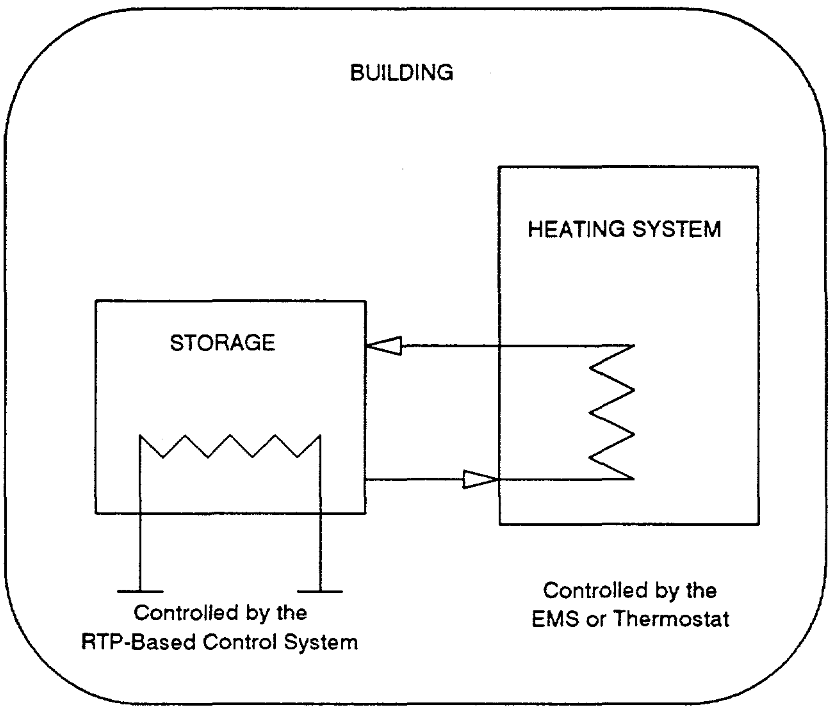

Figure 3-2. Independent Control of Storage Charging and Heating Operations.

OPERATION OF RTP-BASED CONTROL SYSTEM

The central computer in Cambridge supervised and automatically implemented all site data retrieval, storage, control computations, and site scheduling. It coordinated placement of calls to the sites and weather service, and to a price subdirectory/file which resided in the computer itself. The computer functioned as the final clearinghouse by processing and compiling data files and placing them in the appropriate databases. The optimization programs, also stored in the central computer, determined the near-optimal storage charging schedules which were uploaded to the sites.

The central computer initiated execution of required programs using an automatic event management program, according to a time-stamped schedule in an "action" file. The required files and the timings of their executions were automatically entered into the action file at various stages of the daily operations. Table 3-1 provides a brief description of the programs used in the operations.

Table 3-1. Programs Used in RTP-Based Control Systems

Program	Description
P1MAIN	Ran the main event management program.
DLPRC	Downloaded price forecasts.
DLTMP	Downloaded temperature forecasts.
DLBR1	Retrieved data from the Brewster office building.
RCBR1	Calculated optimal schedule for the Brewster office building.
PRBR1	Transmitted schedule to the Brewster office building.
DLBR2	Retrieved data from the Brewster storefront.
RCBR2	Calculated optimal schedule for the Brewster storefront.
PRBR2	Transmitted schedule to the Brewster storefront.
DLPLA	Retrieved data from the Plattsburgh building.
RCPLA	Calculated optimal schedule for the Plattsburgh building.
PRPLA	Transmitted schedule to the Plattsburgh building.

A typical action file is shown in Table 3-2.

Table 3-2. Typical Action File

03/02/90	07:02	*DLTMP
03/02/90	08:02	*DLBR1
03/02/90	08:03	*RCBR1
03/02/90	09:02	*DLBR2
03/02/90	09:03	*RCBR2
03/02/90	10:02	*DLPLA
03/02/90	10:03	*RCPLA
03/02/90	11:32	*DLPRC
03/02/90	12:02	*DLBR1
03/02 90	12:03	*RCBR1
03/02/90	13:02	*DLBR2
03/02/90	13:03	*RCBR2
03/02/90	14:02	*DLPLA
03/02/90	14:03	*RCPLA
03/02/90	16:02	*DLBR1
03/02/90	16:03	*RCBR1
03/02/90	17:02	*DLBR2
03/02/90	17:03	*RCBR2
03/02/90	18:02	*DLPLA
03/02/90	18:03	*RCPLA
03/02/90	19:02	*DLTMP
03/02/90	20:02	*DLBR1
03/02/90	21:02	*DLBR2
03/02/90	22:02	*DLPLA
03/03/90	00:02	*DLBR1
03/03/90	01:02	*DLBR2
03/03/90	02:02	*DLPLA

The programs whose names appear after the asterisk were executed at the specified times. Figures 3-3, 3-4, and 3-5 diagram data flow amongst the various components of the RTP-based control system.

Data Retrieval and Storage

At each site, the data loggers continuously collected the measured data, performed the initial processing (time averaging, integration, etc.), and placed the data in half-hourly tables. The data logger at the Brewster office building kept the 14 most recent days of data. The data loggers at the Brewster storefront

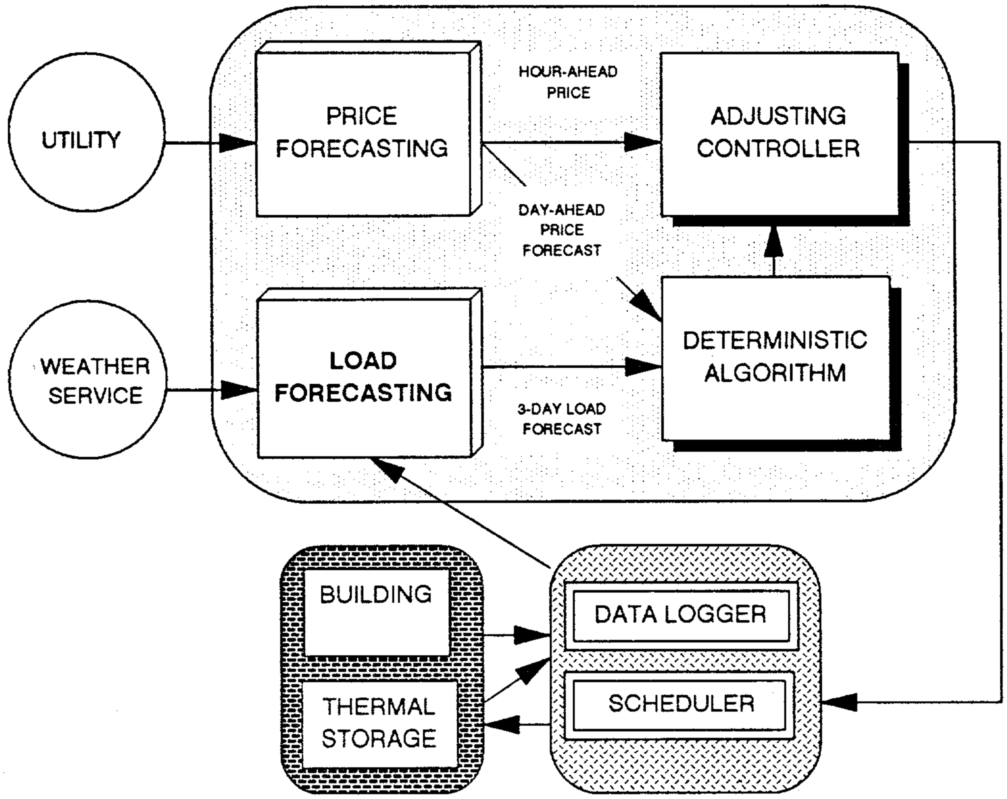

Figure 3-3. System Data and Calculation Flow Diagram.

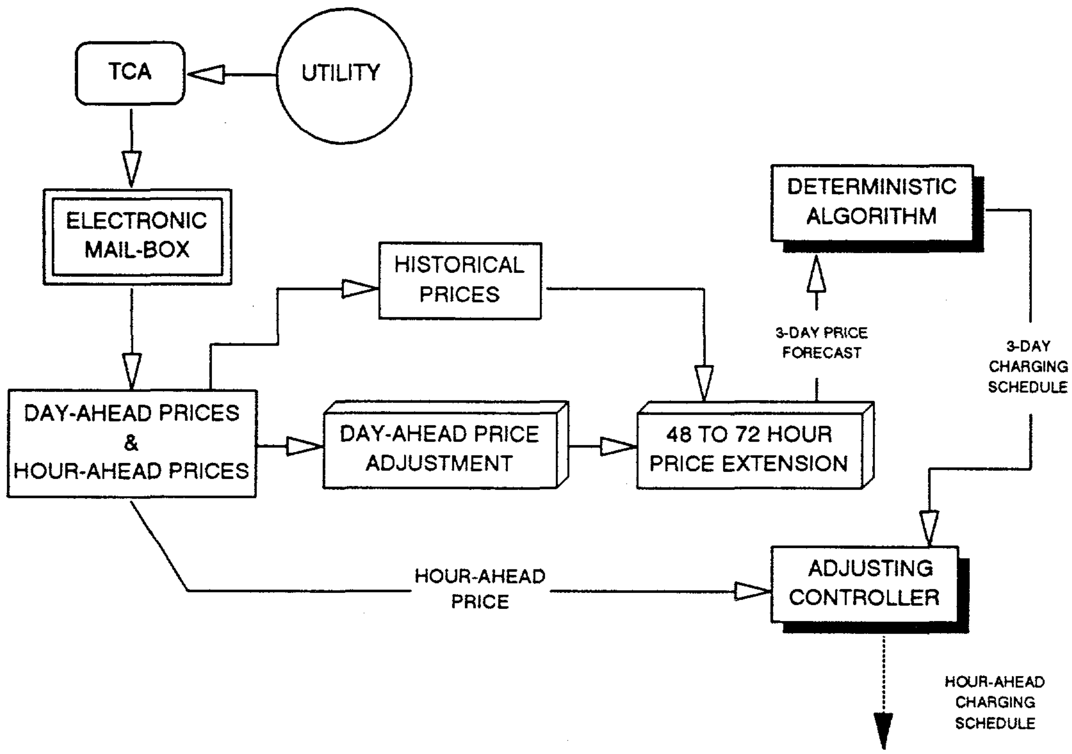

Figure 3-4. Price Forecast Data and Calculation Flow Diagram.

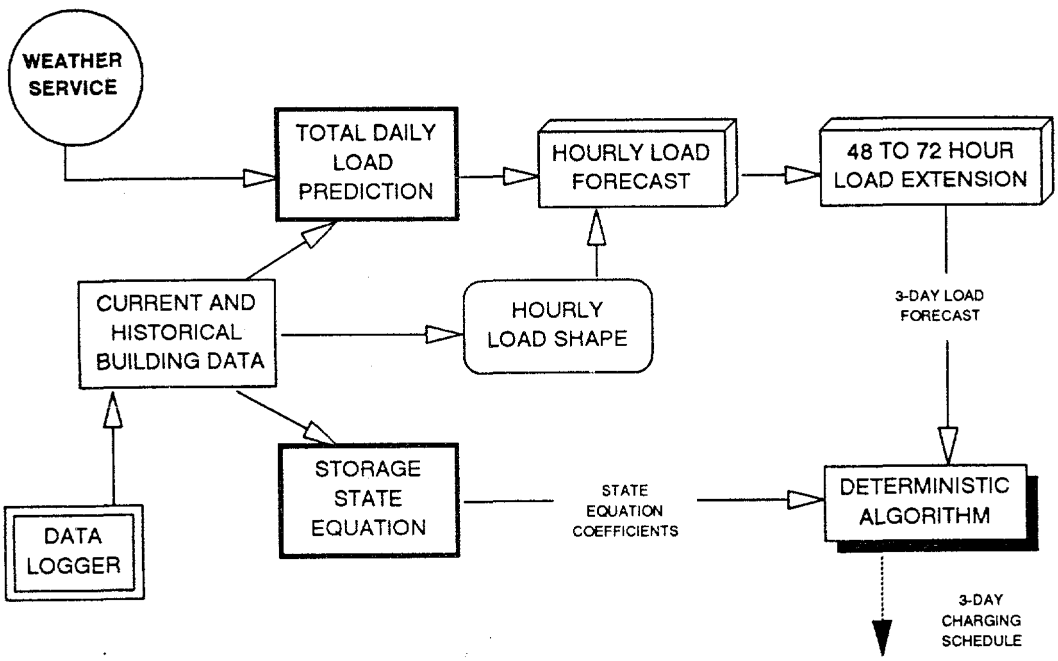

Figure 3-5. Load Forecast Data and Calculation Flow Diagram.

and Plattsburgh office building could store only 24 hours of data. In a typical data logging and control session, the event manager executed a data retrieval program (DLBR1 or DLBR2 or DLPLA) at a predetermined time which placed a communication call to the specified site and received a table of half-hourly data--the most recent 48 hours for the Brewster office building (retrieval of all 14 days of data was not necessary), and the most recent 24 hours for the Brewster storefront and Plattsburgh office building.

Immediately after downloading data from each site, the data were reformatted and appended to a data file which provided a historical performance database for each site. The data were time-stamped, accessible on-line when needed, and used as inputs to the calculations. Based on the schedule of the action file, but independent of communications with the sites, the central computer executed programs DLPRC and DLTMP, which retrieved price and outdoor temperature forecasts from the price subdirectory and weather service company, respectively. The most recent price data received by facsimile were manually entered into the price subdirectory. Weather information was downloaded to the central computer in Cambridge twice a day, at approximately 7 A.M. and 7 P.M. The price and weather programs reformatted and time-stamped the data and entered them into a data file that was accessed when the optimization programs were executed.

<u>Calculations and Scheduling</u>

Immediately following successful downloading of data from a site, the event manager called the corresponding optimization program (RCBR1, RCBR2, or RCPLA) that determined a least-cost storage charging schedule. This program accessed the historical data files, price file, and forecasted site temperature files. The current storage status, the initial condition for the control algorithm, was input from the recently appended historical data file. The program also extracted the hourly heating load and site temperature data of the last day, and, in conjunction with the predicted future site temperature and predetermined load-pattern shapes, computed forecasts of the next two days' hourly heating loads.

The price and building load forecasts were used in a two-stage procedure to calculate a storage-charging schedule. The first stage assumed deterministic conditions and considered the load and price forecasts to be accurate, and produced an optimal schedule for the following 48 hours or more as a solution of a deterministic algorithm. The second stage considered the next available hour-ahead price and the most recent reading of the storage level, and modified the schedule by means of an adjusting controller algorithm.

The charging schedule was actually a schedule of maximum and minimum storage temperature limits, which acted as an adjustable hysteresis band. The storage heating elements were activated or deactivated depending on the value of the actual tank temperature relative to the temperature band. The heating elements were turned on by setting the temperature band above the maximum allowable storage temperature. Conversely, setting the temperature band below the minimum allowable storage temperature turned off the heating elements.

After a charging schedule for storage was determined for the next two or three days (depending on the time of the scheduling), the first 18 hours of the schedule was entered into a file. The 18-hour limit was imposed by the memory constraints of the Trinet data logger units. All optimization results were saved in files for future analysis. A schedule of future data logging and control events for the following three days was written into the event manager's action file.

Next, commands were written to the event manager's action file to immediately execute a communication program (PRBR1 or PRBR2 or PRPLA) which called the site and programmed the charging schedule into the data logger/controller hardware. The controller switched relays on or off according to the residing schedules. These relays, in turn, controlled electrical heating of the storage.

General Features

The central computer in Cambridge called each site six times daily (every four hours) on a predetermined schedule, retrieved the stored data from each data logger, ran the optimization algorithm, and sent the control schedule back to the site scheduler. If no connection was made, the computer immediately tried to call again. After a limited number of tries, the computer terminated the process and tried again the following hour. Site communications were established such that only one site was called in any given hour.

The operation was sufficiently robust to repeat calls to a site if no data or bad data were received. In case of three repeated failures, calls were automatically rescheduled for the next hour. Experimental failures were usually due to problems with communication, data loggers, or sensors. An operations log maintained by the computer recorded all communications (downloading of data or uploading of control schedule) with each site, and noted the occurrence of any errors. It also logged the same communication information regarding temperature and price retrievals. In addition, any error message produced during the computations was written into the operation log.

During the experiment, the site data and control calculations were manually backed up onto floppy disks once a day. This backup was followed by an intensive review of all the information about the site. The analysis included checking the load forecast, the charging time, and information from temperature sensors. Much of the analysis was done using graphs that contained all the variables used in the control algorithm. If the charging behavior could not be understood from the graphs or printed data, the algorithm was examined to determine the reason for the unexpected behavior.

Running Time

The data logging took about two minutes. The data reformatting, checking, and storage took about one minute. The running of the optimization algorithm took one and one-half minutes. Sending the control schedule from the central computer to the site scheduler took from three (Campbell system) to twelve (Trinet system) minutes. These times can be shortened considerably in the future by using faster computers, modifying the control algorithm and communication software to make them more time-efficient, and using modems with higher baud rates.

Incorporation into Commercial Building EMS

Eventually, the RTP-Based control system can be incorporated into a commercial-building EMS if the EMS is based on a PC mother board with similar programmability and data storage capabilities. There would also be requirements for additional sensors and integration of price and weather communication within the EMS. An initial option would be to keep the RTP-based control system separate from the commercial building EMS, and to communicate with the EMS through a simple interface. This should require only a minimum amount of EMS programming. In this way, a single RTP-based control system design could be used with a multitude of EMS's currently in use in commercial buildings.

CONTROL ALGORITHMS

Two-Stage Optimization

The control algorithm consisted of two stages. First, using the multi-day price and load forecasts, a deterministic linear program was solved to calculate an optimal schedule, assuming the price forecast and load forecast were completely accurate. A previously developed non-simplex algorithm took advantage of special features of the problem and was much faster than a standard linear program. Second, an adjusting controller algorithm took into account the uncertainties and inaccuracies of all forecasts and determined, hour by hour, whether charging should be done during that hour. If the forecasts were perfect, then the adjusting procedure simply gave the same results as the deterministic optimization. More often, the procedure reacted to unpredictable changes (increases or decreases) in the prices and loads. It calculated a "threshold" price and compared it with the hourly price (hour-ahead forecast), and depending on the actual storage status, adjusted the schedule determined in the first stage.

Deterministic Algorithms

The deterministic algorithm was run every four hours for each site. In principle, the algorithm could be run each time significant new forecast information was available. The algorithm was executed in one and one-half minutes on a 12-MHz IBM-compatible 286-based PC, this computational time was insignificant. (The term deterministic, as opposed to stochastic, refers to a class of algorithms that assume the input information about the future will hold with certainty for the given period.)

Brewster Office Building and Brewster Storefront. The hourly load forecasts, the day-ahead price forecasts, and the current storage temperature were inputs to the deterministic algorithm. Embedded in the algorithm was a state equation for storage at each site. The state equation related the energy in the storage system at each hour to the storage energy level and the energy inputs and outputs of the previous hour. The coefficients for the state equation were determined using regression analysis of the accumulated site data. The state equations were updated twice during the experiment. The state equations and their coefficients are described later in this section.

Although the optimization problem was structured as a linear programming problem, the algorithm used a fast non-simplex method. The solution of the algorithm was the optimal storage charging schedule for the remaining hours of the current day plus the full hours of the following two days. Therefore, depending on the time of the optimization run during the day, the time horizon ranged from 48 to 72 hours. Any savings or costs after the end of the time horizon were ignored. A time horizon of longer than 24-hours was chosen in order to:

- avoid a full discharge of storage at the end of the 24-hour period; and
- take advantage of scheduling across days if the size of storage compared to building load was large.

The algorithm computed an initial solution which was based on meeting the heating requirements at each hour by heat generation at the same hour without storage action. Next, the algorithm searched in structured fashion for the most expensive hour in the time horizon and substituted that hour's electricity usage with a less expensive hour's, without violating any storage or production constraints. It continued searching for the next most expensive hour and repeated the scheduling until no further savings were possible. A closely related algorithm is discussed in Daryanian, Bohn, and Tabors, 1989[8]. The input and output flow to the deterministic algorithm are shown in Figure 3-6.

Plattsburgh Office Building. Due to the difficulty of measuring the heat outflow from the earth storage, the deterministic algorithm for Plattsburgh had a different and simpler form. The problem arose from the fact that the earth storage was not enclosed, which resulted in continuous heat transfers between the ground, the storage, and the building based on the their respective temperatures. Historical data were used to estimate a linear relationship between desired average daily temperature difference between indoor and outdoor, and the number of hours required to activate the heating elements or mats below the sand and cement floor. The algorithm used this relationship to decide the number of hours needed for charging by using the forecast of the outdoor temperature for the next day. The algorithm then selected the least expensive hours for charging. The control was based on the measured temperature of the bottom of the floor slab. Similar to the water storage cases, the heating elements were activated by setting the temperature hysteresis band above or below the maximum or the minimum allowable temperatures.

Adjusting Controller Algorithm

Brewster Office Building and Brewster Storefront. The deterministic algorithm would have sufficed if both the day-ahead price forecasts and the heating-load forecasts were completely accurate. A second-stage algorithm, the adjusting controller, was developed to account for the inherent uncertainty of future prices and heating loads.

The deterministic algorithm chose the charging periods for the time horizon based on day-ahead prices. A threshold price, $P_{threshold}$, was found below which the deterministic algorithm requested storage charging. The deterministic algorithm also provided a forecast of tank temperature T_{Detr} for each future hour in the time horizon based on day-ahead prices and forecasted heating loads.

[8] Daryanian, B., Bohn, R. E., and Tabors, R. D.,"Optimal Demand-Side Response to Electricity Spot Prices For Storage-Type Customers," *IEEE Transactions on Power Systems*, Vol. 4, No. 3, August 1989, pp. 897-903.

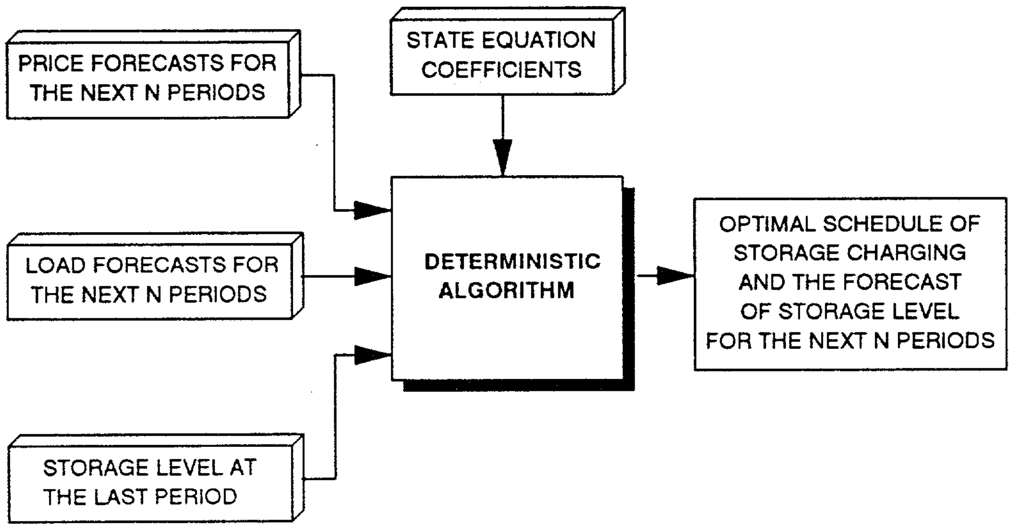

Figure 3-6. Input and Output of the Deterministic Algorithm.

For each hour, the adjusting controller used $P_{threshold}$ and T_{Detr} and other pre-selected constant parameters to construct the combined parabolic decision curve. A decision curve is shown in Figure 3-7. Since $P_{threshold}$ and T_{Detr} changed every hour, the decision curve also changed. The objective was to adjust the desired deterministic tank temperature T_{Detr} when the hour-ahead price and actual tank temperature deviated from their forecasted values. At each hour, the hour-ahead price, $P_{Hour-ahead}$, was applied to the decision curve to determine a new tank thermostatic setting, T_{Heur}. The operation of the storage charging depended on the relationship between T_{Heur}, T_{Detr}, $P_{threshold}$, $P_{Hour-ahead}$, and the actual tank temperature. For example, if $P_{Hour-ahead}$ was greater than $P_{threshold}$, then the tank thermostat setting would be lowered from T_{Detr} to T_{Heur} to prevent storage charging. However, charging would occur if the actual tank temperature, which was forecasted to be T_{Detr}, was actually below T_{Heur}. Similarly, if $P_{Hour-ahead}$ was lower than $P_{threshold}$, then the tank thermostat setting was increased in order to force storage charging. Again, charging occurred only if the actual tank temperature was below T_{Heur}.

If the price communication from the utility was automated and hour-ahead prices were available every hour, then only the first hour would be implemented, and the entire scheduling computation would be repeated every hour. However, for this experiment the scheduling computation was repeated once every four hours for each site in order to minimize the time and the cost of communication with the sites. Therefore, every four hours, the deterministic algorithm was run to determine a charging schedule for the current day and the next two days (a time horizon of 48 to 72 hours). The prices used in the deterministic algorithm were the current day-ahead prices adjusted and extended using the previous price data and the hour-ahead price of the current hour. Then the adjusting controller was used to adjust the schedule using the hour-ahead prices of the future hours. This process mimics the case in which a charging schedule based on the deterministic algorithm using the day-ahead price data and the current hour-ahead price is calculated for the corresponding time horizon. Then, instead of hourly recalculation of the schedule for the entire time horizon, the schedule from the deterministic algorithm is adjusted every hour as the hour-ahead price becomes available. If there are no communication or computational constraints, it is possible to adjust and extend the day-ahead prices every hour as the hour-ahead prices become available, run the deterministic algorithm once every hour, and implement the storage charging plan of the first hour in the time horizon. Under the scenario just described, there would be no need for the adjusting controller.

Plattsburgh. In the case of Plattsburgh, the adjusting controller had a slightly different form. Again, similar to the water-storage system, a threshold price was calculated based on the highest price below which the deterministic algorithm allowed storage charging. However, instead of using temperature, the earth-storage status was measured using the actual number of hours charged compared to what was

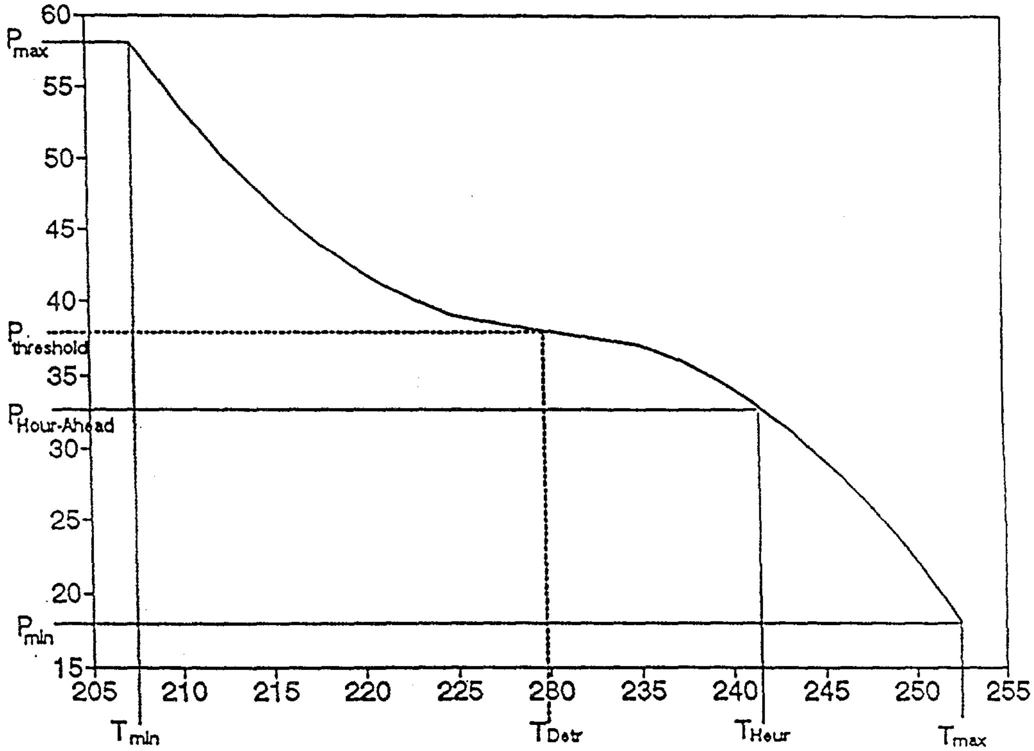

Note: The objective is to adjust the desired deterministic tank temperature (T_{Detr}) when the hour-ahead price and actual tank temperature deviate from their forecasted values. At each hour, the hour-ahead price, ($P_{Hour\text{-}ahead}$), is applied to the decision curve to determine a new tank thermostat setting, (T_{Heur}). However, the operation of the storage charging depends on the relationship of T_{Heur} and the *actual* tank temperature.

Figure 3-7. Decision Curve for the Adjusting Controller.

determined by the deterministic algorithm. Therefore, at each hour, the adjusting controller compared the hour-ahead price to the threshold price, and, depending on the actual number of hours already charged, decided whether it should charge during the next hour. If the answer was in the affirmative, the hysteresis band was set above the maximum allowable slab temperature. Because of the very slow dynamics of the earth-storage system, and because of the time delay in transferring the heat from the heating mat to the slab, heating elements would sometimes not be activated for hours.

PRICE INPUTS

Prices

For each optimization run, the deterministic algorithm used the day-ahead price and load forecasts to solve for a deterministic optimal storage charging schedule. The hour-ahead price forecasts were then used as inputs to the adjusting controller, which modified the storage charging schedule by considering the deviations in prices and storage behavior from their predicted values. The actual hourly marginal cost data were used in the evaluation and cost analysis. These price forecasts were provided by NYSEG, as discussed in Section 2.

As has been discussed, it was determined to be operationally unfeasible to send the utility price forecasts in real time to the central computer in Cambridge during the experiment. At NYSEG, the day-ahead price forecasts were made every day before noon for the 24 hours starting at midnight of the next day. Similarly, the hour-ahead price forecasts were made about 45 minutes before the hour in which they would take effect. At the same time, the utility knew the actual marginal costs for the previous hour.

The NYSEG utility control room sent price information by facsimile to the TCA office in Cambridge before noon each weekday. On each Monday, the actual prices and hour-ahead forecasts were received for the previous Friday through Sunday, and the day-ahead forecast was received for Tuesday. For the period Tuesday through Friday, the actual prices from the previous day and price forecasts for the next day were received. Additionally, TCA received forecasts for Friday through Saturday on Thursday, and for Sunday through Monday on Friday.

Every Monday and Friday, the price information used by the control algorithms was placed into a single, formatted ASCII file in tabular form and put into the price subdirectory in the central computer. The file included the time and the hour at which the prices would be in effect during the experiment. Since the prices from the previous week were assumed to be in effect for the current week, the time stamp of the prices was advanced by seven days. Twice a day, the event manager executed a pricing program that looked at the price subdirectory for the latest price file and appended the "current" price data to the historical price data file. When required, the central computer accessed the price data file and chose the necessary price data. In this manner, the process of on-line price data retrieval during the optimization phase was divorced from the actual receipt of the prices from the utility. Therefore, the manner in which the price information was used in the experiment would be the same as under automatic price communication.

Day-Ahead Price Forecast Shift

During the experiment, each time a site was rescheduled, the day-ahead price forecasts used in the deterministic optimization algorithm were adjusted to reflect the most recent information contained in the hour-ahead price forecasts. A simple trend-averaging method was used, which consisted of finding the average of differences between the day-ahead and hour-ahead forecasts from the first hour of the day to the current hour of the day, and then shifting the day-ahead price forecasts by that value for the remainder of the day. In a future, non-experimental use of RTP, this 24-hour forecast could be revised by the utility several times a day, obviating the need for day-ahead price adjustment as performed in this experiment.

Day-Ahead Price Extensions

When a site reprogramming was done before noon, the day-ahead price forecast available covered only the remaining hours of the same day. As a result, the time horizon of the deterministic algorithm was less than 24 hours. However, a day-ahead price forecast for the next day became available around noon. Therefore, in the afternoon the time horizon available was somewhere between 24 and 36 hours. In both cases, the time horizon based on the available utility forecasts of day-ahead prices was too short to enable the algorithm to account for the inter-day opportunities for electricity load shifting, which could be particularly significant on weekends. In other words, a shorter time horizon did not properly address the end effect problems, since it completely ignored any benefits or costs beyond the last period in the time horizon. Thus, it became necessary to extend the price forecasts by two days beyond the available time horizon of the utility's price forecasts.

Since it had been decided that the experiment should operate in a mode that received the prices under simulated real-time conditions, the end effects problem had to be addressed. In the absence of historical trend or future information, the simplest assumption was that the prices would remain the same day after day. However, it was probable that in the future the utilities would have greater ability to make qualitative judgments about the overall trend for the prices of the following few days. Given that TCA had access to prices for the following few days, it was decided to use a future day's information to qualitatively assess the "magnitude" of the prices for that day and then assign a typical price pattern to that day.

In practice, this was done by classifying the hour-ahead price forecasts of the preceding four weeks into expensive, normal, and inexpensive days, and determining to which category a future day belonged. It was assumed that even under RTP, most of the storage charging periods corresponded to the TOU off-peak hours (10 P.M. to 7 A.M.). Hence, the classifications considered only these periods.

During each optimization run, the average and standard deviation of the off-peak prices for the previous four weeks were computed. Weekends were treated separately from weekdays. Days were classified as medium if the average of off-peak period prices for those days was within 40% of one standard deviation of the total average. If the average fell above or below the medium band, then the day was classified as expensive or inexpensive, respectively. The average of hourly prices of days in each category (inexpensive, medium, expensive) was taken to represent a typical day of that category. Next, the average of the off-peak period of the hour-ahead prices for the future day (into which the day-ahead prices were to be extended) was compared to the classification. Then, according to the classification, the associated typical day's hourly prices were taken as the forecast of that day's prices. In this manner, the effective time horizon of day-ahead price forecasts used in the deterministic optimization algorithm was extended by 48 hours.

LOAD INPUTS

Weather Forecast

The weather forecast used was from the limited-area, fine mesh model of the National Weather Service. All forecasts were produced using the model output statistics technique and were supplied by WeatherBank, Inc., Salt Lake City, Utah. The forecasts contained precipitation probabilities, thunderstorm probabilities, tri-hourly temperatures and surface dew points, and surface winds. The control algorithm, however, only used the temperature forecast. The weather information was updated twice a day and contained temperature information for the next 48 hours. It was obtained via modem without human participation. Once the information was transmitted to the central computer in Cambridge, the temperature forecasts were extracted and converted into hourly information through linear interpolation.

Weather forecasts were not available for the cities of Brewster and Plattsburgh. Therefore, the Albany, New York, weather forecast was used for Brewster, and the Burlington, Vermont, weather forecast was used for Plattsburgh.

Building Load Forecast

The control algorithm required an estimate of the hourly heating load for the duration of the time horizon to perform the optimization. The forecasting method used a two-step process. First, the following day's total load was estimated based on the following day's forecast of outdoor temperature, previous day's total load, and the day of the week (weekdays and weekends). Second, the forecasted total load was distributed into forecasted hourly load according to the percentage of the total load that historically occurred during each hour. The hourly percentages formed a normalized daily load pattern. The load forecasting process was automated in the second season as described later.

Brewster Office Building. The total heating load (measured when the circulating pump was operating) was proportional to the difference between the storage outflow and inflow temperatures, multiplied by the water flow rate. The total daily heating load for the following day could be accurately forecast based on a forecast of the day's average temperature, the previous day's actual load, and the day of the week. According to several regression analyses, changes in the outdoor temperature accounted for about 80% of the variation in load. It was found that the previous day's load was also a predictor of the current day's load. Data indicated that weekends typically had higher loads than weekdays[9]. Regression using three independent variables (previous day's load, average outdoor temperature, weekday or weekend) produced an R^2 of 0.95 and all independent variable coefficients were significant. Holidays were treated as weekend days. Other regressions were performed using different combinations of variables but none yielded better results. In the second season, the least-squares estimation technique was used automatically on a daily basis.

Hourly load shape used in the optimization was determined by first examining the hourly load of about 40 days. The general heating pattern peaked in the morning, indicating initial heating of the building, leveled off during work hours, and dropped to a lower level after work hours. This pattern held regardless of the average outdoor temperature and the day of the week. The area under the curve, however, was affected

[9] This is believed to be because the building is kept at the same temperature schedule on the weekend as during the week. During the weekend, the heating system has to supply the heat that is given off by people, computers and other machines, which would account for the large weekend loads.

by those two factors. The load shape was determined by taking the average percentage of total daily load during each of 48 one-half hourly periods over several days. Figure 3-8 shows the average load shape for the site.

Brewster Storefront. The load forecast for the Brewster storefront was different from that of the Brewster office building because of the different nature of the businesses. The office building was occupied by a relatively large number of people for a large portion of the day, while the storefront operated on a strict schedule typical of a retail store. The factors that influenced the total load were average outdoor temperature for the day and the day of the week. Examination of the heating load graphs showed that the heat was on only when the store was occupied; thus, the heating load was driven by the store's operating schedule (Monday to Friday 9:30 A.M. to 8 P.M., Saturday 9:30 A.M. to 6 P.M., Sunday noon to 5 P.M.). Other factors studied in addition to the day's average outdoor temperature and the day of the week included the previous day's heating load, the previous day's average outdoor temperature, and the current average temperature squared. These were not significant predictors of heating load.

It is believed that some of the variation of daily heating load was caused by the number of customers that passed through the store's door. When the door was opened, heat escaped. This heat was then replaced by the heating system. The number of customers depended on the day of the week, season of the year, and sales promotions.

The Brewster storefront load shape was basically the same from day to day. The only differences were due to the number of hours the store was occupied. The load shape held regardless of the average outdoor temperature. Because the storefront was open for only a few hours on Sunday, two load shapes were calculated, one for Monday through Saturday and one for Sunday. The load shapes were determined by taking the average percentage of total daily load during each half-hour over several days. Figure 3-9 shows the two average load shapes.

Load Forecast Extensions

Since load forecast data were inputs to the deterministic algorithm, it was necessary to extend the load forecast 24 hours to a total of 72 hours. It was assumed that the outdoor temperature would be repeated for the future days and, consequently, the load forecast was simply assumed to be in effect for another day, with explicit consideration for the day-type.

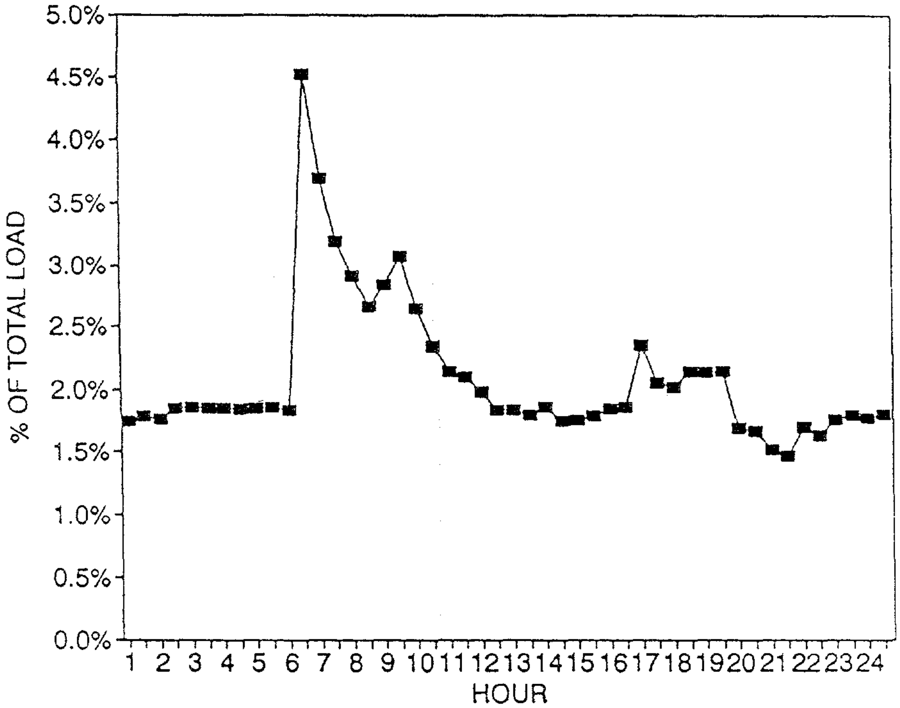

Figure 3-8. Heating Load Pattern for the Brewster Office Building.

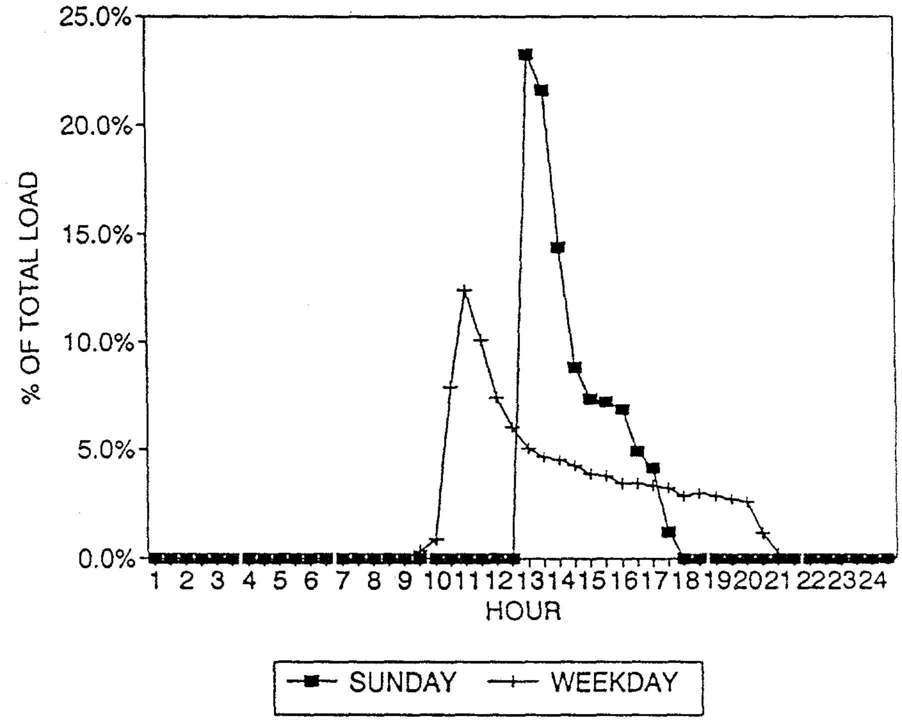

Figure 3-9. Heating Load Pattern for the Brewster Storefront.

ESTIMATION AND PREDICTION ENHANCEMENTS

In the winter of 1990-91, a number of enhancements were implemented which contributed significantly to the robustness of the system and made the control and forecasting algorithms adaptive and responsive to changes. However, further work in the development of better load forecast models is encouraged.

The major enhancements included the daily linear least-squares estimates of state equation coefficients and total load parameters. In addition, the daily normalized load patterns were calculated daily instead of relying on a one-time calculation. These algorithms were automatically run once a day to update the state equation coefficients and the total load parameters. Before implementing the daily update feature, the coefficients and parameters used in the forecasting and control algorithm were based on a one-time regression analysis of historical data.

The continuous daily recalculation produced a more adaptive system model where the daily or seasonal changes were reflected in the evolving values of the coefficients and parameter. Results of the daily calculations were automatically entered into a database. The control and load-forecasting algorithms accessed the database each time they were run for the most recent values of these coefficients and parameters.

Linear Least-Squares Estimation of State Equation Coefficients

The state equation provides a model of the energy input and output from storage. It is a discrete linear equation which relates the storage level at a given period to the storage level and electricity input and heat withdrawal during the previous period. Because the storage systems were enclosed, the buildings' heating loads were computed from the data provided by the temperature sensors affixed to the supply and return circulation pipes, together with the data from the flow-rate sensors. The linear discrete storage state equation has the following form:

$$X[k+1] = a\,X[k] + b\,U[k] + c\,W[k] + d$$

where $X[k]$ (in °F) is storage value at period k, $U[k]$ (in KWH) is electrical charging at period k, and $W[k]$ (in KWH or proportional to it) is the building heating load at time k. Storage tank temperature was chosen as a proxy measure for the storage energy level. The values of a, b, c, and d (state equation coefficients) are estimated from the chronological data. For the winter season of 1989-90, the coefficients used in the control algorithm were based on a one-time linear regression of the historical data for storage tank temperature, electricity input, and heat withdrawal from the tank. State equations were:

a) for the office building:

$$\boldsymbol{X[k+1] = X[k] + 0.025\ U[k] - 0.017\ W[k] - 0.40}$$

and

b) for the storefront:

$$\boldsymbol{X[k+1] = X[k] + 0.47\ U[k] - 0.45\ W[k] - 0.32}$$

The differences in coefficient values reflect the unit conversion factors in addition to the static errors in the measurement sensors. These models were sufficiently well-behaved to be used in prediction and control. The same state equations were used in the RTP- and TOU-based control simulations. The state equations were constructed using data from both TOU- and RTP-based control operations.

The daily updating of these coefficients in the winter season of 1990-91 was based on the linear least-squares estimation (LLSE) technique. Initially, a recursive linear least-squares estimation (RLLSE) technique was used, in which the last set of evaluated coefficients was input into the estimation algorithm. However, after testing, it was shown that a moving LLSE technique provides a satisfactory and stable set of coefficients, given an empirically derived time window for estimating 13 days. The effect of using a moving LLSE instead of RLLSE was equivalent to ignoring past data or assigning a zero weight to the old estimates beyond a certain time window in the RLLSE technique.

Once every day, the central computer automatically ran the LLSE algorithm and appended the time-stamped results to the state equation coefficients database in the central computer. The algorithm accessed the site data log and extracted the half-hourly storage temperature, electricity input, and heat withdrawal data for the previous 13 days. It then used the LLSE technique to compute the state equation coefficients.

Actual LLSE results for state equation coefficients for the Brewster office building and storefront sites are provided in tables 3-3 and 3-4, respectively.

Table 3-3. State Equation Coefficients for the Brewster Office Building.

a	b	c	d
1.0000	0.0319	-0.0185	-0.1545
1.0000	0.0319	-0.0238	-0.0678
1.0000	0.0327	-0.0218	-0.1178
1.0000	0.0326	-0.0224	-0.0941
1.0000	0.0325	-0.0239	-0.0754
1.0000	0.0323	-0.0251	-0.0696
1.0000	0.0320	-0.0274	-0.0346
1.0000	0.0317	-0.0271	-0.0513
1.0000	0.0321	-0.0277	-0.0375
1.0000	0.0307	-0.0262	-0.0534

Note: Rows represent coefficients calculated on a daily basis for 10 sequential days. Each day's calculation is based on the storage and energy data of the previous 13 days.

Table 3-4. State Equation Coefficients for the Brewster Storefront.

a	b	c	d
1.0000	0.4679	-0.5004	-0.0269
1.0000	0.4716	-0.5091	-0.0127
1.0000	0.4685	-0.4455	-0.1732
1.0000	0.4679	-0.4320	-0.2251
1.0000	0.4709	-0.4316	-0.2444
1.0000	0.4696	-0.4281	-0.2524
1.0000	0.4680	-0.4058	-0.2947
1.0000	0.4674	-0.3902	-0.3269
1.0000	0.4630	-0.3965	-0.3029
1.0000	0.4572	-0.3664	-0.3628

Note: Rows represent coefficients calculated on a daily basis for 10 sequential days. Each day's calculation is based on the storage and energy data of the previous 13 days.

Linear Least-Squares Estimation of Total Load Parameters

The load input data used in the control algorithm were based on projecting the next day's predicted total daily thermal load into an hourly pattern. The next day's total load was predicted from an empirical equation which related the next day's total load to the previous day's total load, the next day's average outdoor temperature, and the next day's day-type. The linear discrete total load prediction equation had the following general form:

$$L[i+1] = a\, L[i] + b\, T[i+1] + c_w\, D_w[i+1] + d$$

Where $L[i]$ (in KWH or proportional to it) is the total daily load for day i; $T[i]$ is the average outdoor temperature for day i; and $D_w[i]$ is the day-type of day i. The coefficient c_w is a row vector of seven values, one for each day of the week. The day-type variable D_w is a seven-element column vector, where each element signifies a day of the week, the first element being Sunday. For each day, the corresponding element of D_w is one, and the other elements are zero. In effect, this is a vector representation of having seven dummy variables for each day of the week. The day-type variables do not necessarily contribute equally to the variation in the total daily load.

The hourly load patterns are normalized patterns of thermal load at the building where each hour's value determines the fraction of the daily load that occurs at that hour. In the winter season of 1989-90, the daily total load predication was computed from a total-load equation whose parameters were computed from a one-time regression analysis of the historical data for each building. For the office building, two day-type variables were used, one for weekdays and one for weekends. For the storefront, Saturdays affected the load in a manner similar to weekdays. Thus, day-type for the storefront was defined as Sunday and non-Sunday. The normalized daily-load patterns were also based on a one-time analysis of historical load patterns classified into weekday-type and weekend-type patterns (Sunday and non-Sunday for the storefront). These load patterns and total load equation parameters were updated once during the experiment for each water storage site.

In the winter season of 1990-91, both the total load equation parameters and the normalized daily load patterns were calculated daily. Again, the technique was based on LLSE. The RLLSE was also considered, but it was decided that an LLSE algorithm using the most recent 25 days' data provided sufficiently stable parameters for forecasting. Thus, once a day the central computer automatically accessed the logged data for each site for the most recent 25 days, and after extracting the required data and reformatting them, it calculated the total load equation parameters and stored them in the total load parameter database. Each time the control algorithm was run, this database was accessed for the latest calculation of the total load prediction parameters, which were then used in conjunction with the following day's outdoor temperature prediction to estimate the following day's total thermal load. Tables 3-5 and 3-6 illustrate the results for a few days' data.

The hourly load patterns for each day of the week were recalculated each time the control algorithm was run. The calculation was based on an averaging window technique. The hourly load data for a particular day was estimated to be the average of the hourly load data of the same day of the week for the most recent 4 weeks.

The overall effect of regular and frequent recalculation of total load prediction equation parameters and the daily normalized load pattern curve is to reflect the gradual shifts in the seasonal conditions or load consumption patterns. An unexpected and abnormal load pattern--a non-scheduled holiday, for example--might temporarily skew the parameters and the load pattern, but the impact gradually would be smoothed over.

Table 3-5. Total Load Prediction Coefficients for the Brewster Office Building.

a	b	c_d	c_w	d
0.5189	-11.0340	1003.6206	1004.9238	818.4678
0.0672	-40.5338	2966.1586	3045.8725	3133.0052
0.2393	-33.5947	2422.4363	2456.8128	2593.7284
0.2934	-32.2857	2307.6342	2337.9928	2399.9336
0.3088	-31.9493	2274.5513	2306.2709	2370.0312
0.3107	-31.8325	2262.6294	2298.0229	2362.0931
0.3361	-31.2398	2223.1579	2243.5141	2310.6297
0.3577	-31.2277	2196.5101	2194.7081	2294.7295
0.3768	-29.9233	2103.9047	2107.8073	2223.6210
0.3789	-31.2218	2170.9051	2165.2396	2288.8564

Note: Rows represent coefficients calculated on a daily basis for 10 sequential days. Each day's calculation is based on the total daily load, day type, and outdoor temperature data of the previous 25 days.

Table 3-6. Total Load Prediction Coefficients for the Brewster Storefront.

a	b	c_d	c_w	d
0.5980	-1.0979	104.9969	96.1717	81.6553
0.2464	-3.1569	247.0489	249.0453	265.0150
0.3403	-2.7817	216.3329	218.9208	232.9192
0.3585	-2.8787	220.8385	222.0691	235.0508
0.3672	-2.9922	227.6439	227.3006	240.5156
0.3659	-2.9524	224.2506	225.2512	238.4144
0.3639	-2.9240	222.9486	223.9142	237.0231
0.4162	-2.6928	204.2690	212.6570	218.4371
0.4328	-2.5540	194.3210	203.4017	211.6016
0.4276	-2.5895	196.4641	209.0177	214.0370

Note: Rows represent coefficients calculated on a daily basis for 10 sequential days. Each day's calculation is based on the total daily load, day type, and outdoor temperature data of the previous 25 days.

FILE AND DATA-HANDLING ENHANCEMENTS

The file and data-handling capabilities of the RTP-based control system were enhanced in the 1990-91 season to introduce more robustness and failure recovery functions. The following list summarizes the improvements:

- Automatic price and weather data checking, which checked the price data for date, time, and completeness. If the date and time of the data were not in order, the program sorted the data according to the correct date and time. In case of redundancy, it dropped the superfluous data. In case of missing data, it substituted an error number (-999) which could later be identified as an error.

- Control under incomplete data, which enabled the control system to continue running when there were errors in the values of stored price, weather, and site-specific data. In effect, the program filled in the missing periods with interpolation of adjacent data. Therefore, a previous error in data retrieval or storing did not automatically result in the later disruption of control.

- Automatic start of a new data file, which resulted in the creation of a new data file if the current data file was lost.

HARDWARE

Computer

The central computer in Cambridge was an NEC PowerMate SX Plus. This computer featured the Intel^{R} 386 SX^{TM} microprocessor, a 16-MHz clock speed, and a 40-Megabyte hard disk. As a substitute, a Zenith Supersport 286 was used. This had an 80C286 processor at a 12-MHz clock speed.

Data loggers/Controllers

The Brewster office building was previously equipped with a Campbell CR10 data logger as part of another project. A controller option and the necessary relays were added for this experiment. The Campbell unit could store two weeks of half-hourly data.

The data logging and control in the Brewster storefront and the Plattsburgh office building were implemented with Trinet (Trimax) Powersense 432 Energy Management Controllers. The Trinet unit could log up to 16 channels and maintain the latest 48 values, which was equivalent to one day of half-hourly data. Output relays were set to determine output action if the control unit failed.

Communications

Communications between the central computer in Cambridge and the three sites were maintained with a Practical Peripherals PM2400SA MNP modem. This modem was chosen for its reliability, speed, and hardware error-correction capability. The Brewster office building had previously been equipped with a Hayes compatible modem.

Almost all communications were accomplished with the Procomm Plus software package. Procomm Plus was selected for the project because it supports error-correcting file-transfer protocols, and it can be programmed using a simple script language. The exception was for the Brewster office building in the data-logging mode, which required special communication software (PC208: TERM and TELCOM) from Campbell. Communication protocols had already been written for this site by The Fleming Group. However, downloading of the control schedule was accomplished using Procomm Plus.

Section 4

SITE PERFORMANCE AND RESULTS UNDER RTP-BASED CONTROL

This section discusses the results of the experiment and site performance under RTP compared to conventional TOU-based control. Week-by-week and total savings during the experiment are presented.

All three sites were under NYSEG TOU-based controls before inception of the RTP-based control experiment. The TOU-based control system was based on a time clock that enabled charging of the storage system during NYSEG's off-peak rate period (from 10 P.M. to 7 A.M. weekdays and any time on weekends). Thus, at 10 P.M. each week-night, the storage heating elements would turn on and begin heating the storage medium (water for the two Brewster sites, and sand for the Plattsburgh site). Heating continued until a target temperature corresponding to a target energy level was reached. This target was based on the outdoor temperature reset formula, in which the contemporaneous outside temperature was used as a proxy for the next day's heating load. Thus, as the outdoors became colder, the storage tank was heated to a higher maximum temperature. These systems were conservative, i.e., the storage systems were generally heated more than was necessary to meet the next day's demand.

In this experiment, the TOU- and RTP-based control schemes were the same during discharge (removal of heat from storage for use in the buildings). At the Brewster sites, the existing buildings' heating systems determined when and how much heat was needed, and pumped hot water out of the storage tanks to meet these needs. The fundamental differences between RTP- and TOU-based control were:

- The amount of energy stored each night was determined in a more complex way under RTP, based on a weather forecast[10]; and

- The timing of energy storage was determined in a more complex way, based on the hourly real-time prices and an optimization program. Customer comfort was not adversely affected, since the building heating system was controlled by the building occupants[11]. The goal of RTP-based control was to use less expensive electricity without adversely affecting comfort.

---

[10] A weather forecast could also have been used under TOU-based controls; however, the forecast would have provided little benefit to the customer, since prices were the same each day. The only penalty for completely filling the storage would have been additional thermal losses.

[11] Customer comfort was improved in Plattsburgh.

For the experiment, weeks started on Saturday. For the first season, week one was January 6 to 12, 1990. For the second season, week one was January 5 to 11, 1991.

BREWSTER OFFICE BUILDING/MEGATHERM SYSTEM

Operations History

For the first season, data logging for this site started in late October 1989. The storage system was put under TCA's manual TOU-based control on December 1, 1989. The control was implemented by adjusting the storage temperature dead band so that charging occurred only during off-peak periods. The outdoor temperature reset was computed using coefficients evaluated by regression analysis of the data collected the previous month. TCA directed TOU-based control was automated on December 11, 1989. RTP-based control, using only the deterministic algorithm, was started on January 4, 1990. The adjusting controller was added on January 10, 1990. An improved adjusting controller with price and load extensions was put into effect on January 21, 1990. Gradual modifications and improvements continued to be made to the end of the experiment on March 31, 1990. The gradual phasing-in of RTP-based control facilitated identification and remedy of initial shortcomings in hardware and software.

Data logged prior to the control experiment indicated that, under conventional TOU-based control, weekends were treated like weekdays; thus, the conventional control system did not take advantage of off-peak prices during the weekend.

The final settings for the minimum and maximum allowable tank temperatures were 170 °F and 270 °F, respectively. These settings were established on January 27, 1990. However, for system security, the tank temperature in the deterministic algorithm had a range constrained between 180 and 255 °F. This range was extended by 10 degrees in each direction in the adjusting controller. For control purposes, the final maximum and minimum dead band storage temperatures were set at T_{Heur} (target temperature from the adjusting controller) and (T_{Heur} - 4) °F, respectively, on February 13, 1990.

An interesting characteristic of the Brewster office building's Megatherm system was that the number of active heating elements depended on the status of the stored tank energy. The heating elements were activated by a pressure sensor, which was more sensitive than a thermostat. The tank temperature was measured by a single point sensor near the middle of the tank. This temperature was used as a proxy for the value of the stored energy in the tank. Due to initial disturbance and movement of water in the tank, the measured tank temperature usually showed a steep drop during the first hour of each charging interval.

There were some variations in the values of charging rate, even for the same tank temperature. When the tank temperature was relatively low, the measured charging rate was about 450 KW, considerably less than the factory rating of this system of 600 KW. This value dropped as the tank became hotter. Analysis showed that there was a cut-off point of 225 °F, above which the charging rate dropped to around 280 KW. Prior to February 12, 1990, the control algorithm assumed that the charging rate remained constant independent of the tank temperature. The algorithm was adjusted to account for this behavior. Some daytime charging was due to internal controls in the Megatherm system, and these could not be accounted for by the RTP-based charging schedule.

The second-season results cover the period from mid-December 1990 to the end of April 1991. The data logging started a few weeks earlier, but actual control commenced with the start of price communication from the utility in mid-December, 1990. The earlier part of the season was spent on implementation, testing, and troubleshooting automatic estimation and prediction of state equation coefficients and total load parameters.

Hourly Behavior of the Brewster Office Building

Figure 4-1 shows the hourly behavior of the Brewster office building for Tuesday, March 6, 1990. The weekly behavior for week nine (Saturday, March 3 to Friday, March 9, 1990) is shown in Figure 4-2. During week nine, the site charged 10,944 KWH[12]. The exact times of charging are shown by the solid line in figures 4-1 and 4-2. Figure 4-2 shows some small charging occurring between 7 A.M. and 9 A.M., due to independent actions by the Megatherm's internal control. The dotted line shows the tank temperature in degrees Fahrenheit, a measure of the tank energy level. Figure 4-3 shows a comparison of the actual load with the forecasted load for week nine.

Each night, as electricity was used, the tank temperature rose quickly. Then later, as the building was heated, the tank temperature slowly dropped. The rate of decline of tank temperature during the day is a function of building load, which in turn is a function of outdoor temperature and other factors. Although building load is not shown on this graph, it was tracked, forecasted, and used as input in the optimization.

During the weekend, the highest temperature, approximately 245 °F, was reached Sunday morning, and the lowest temperature, about 200 °F, was reached late Sunday. The daily working range (maximum and

[12] This was a medium load week. The previous week had the highest load of the experiment, 13,000 KWH. Five out of the 12 weeks had loads equal to or greater than that of week 9.

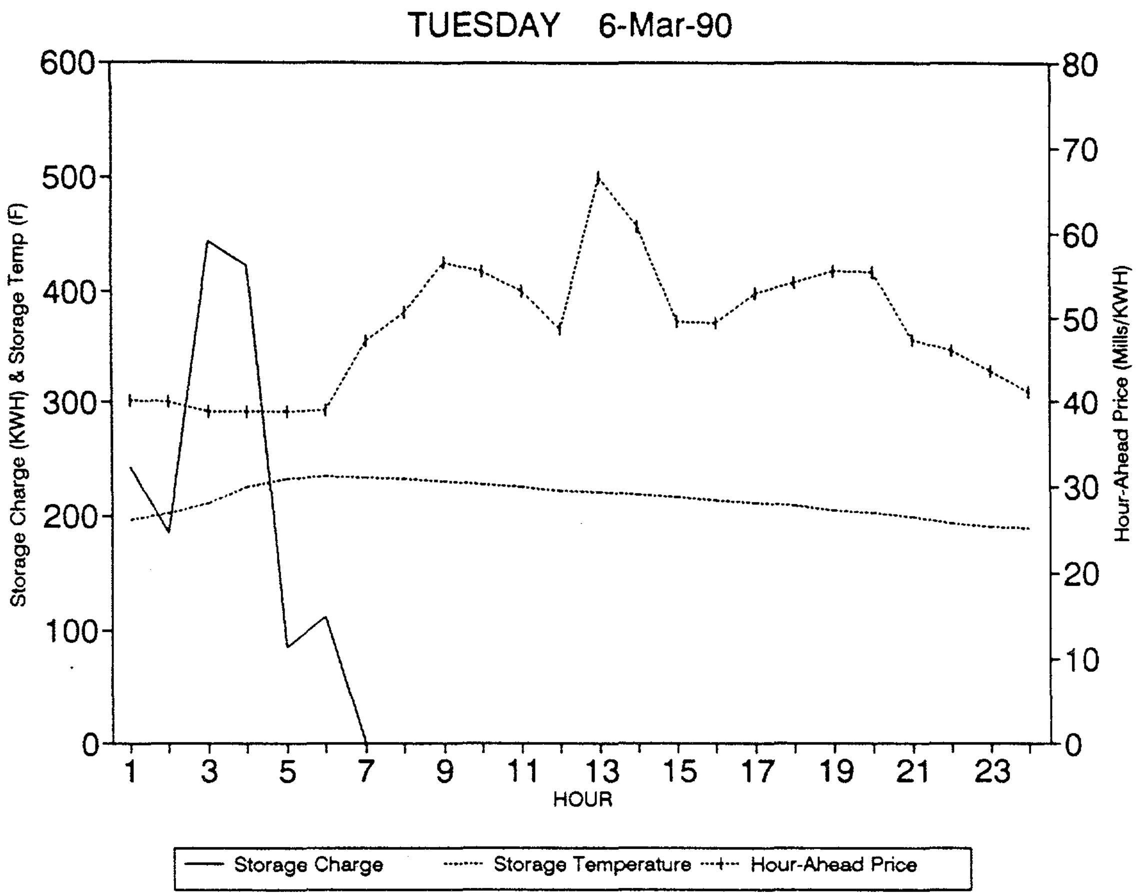

Figure 4-1. Hourly Behavior of the Brewster Office Building.

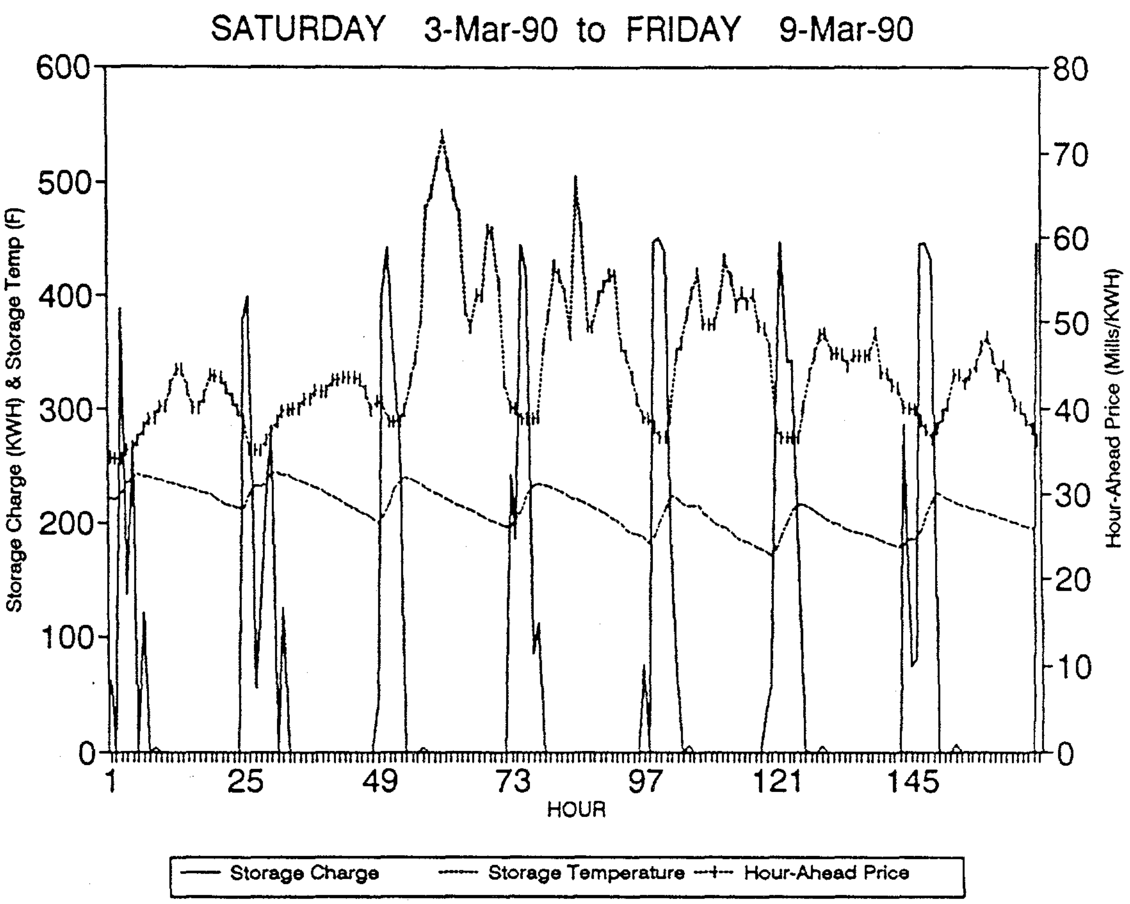

Figure 4-2. Weekly Behavior of the Brewster Office Building.

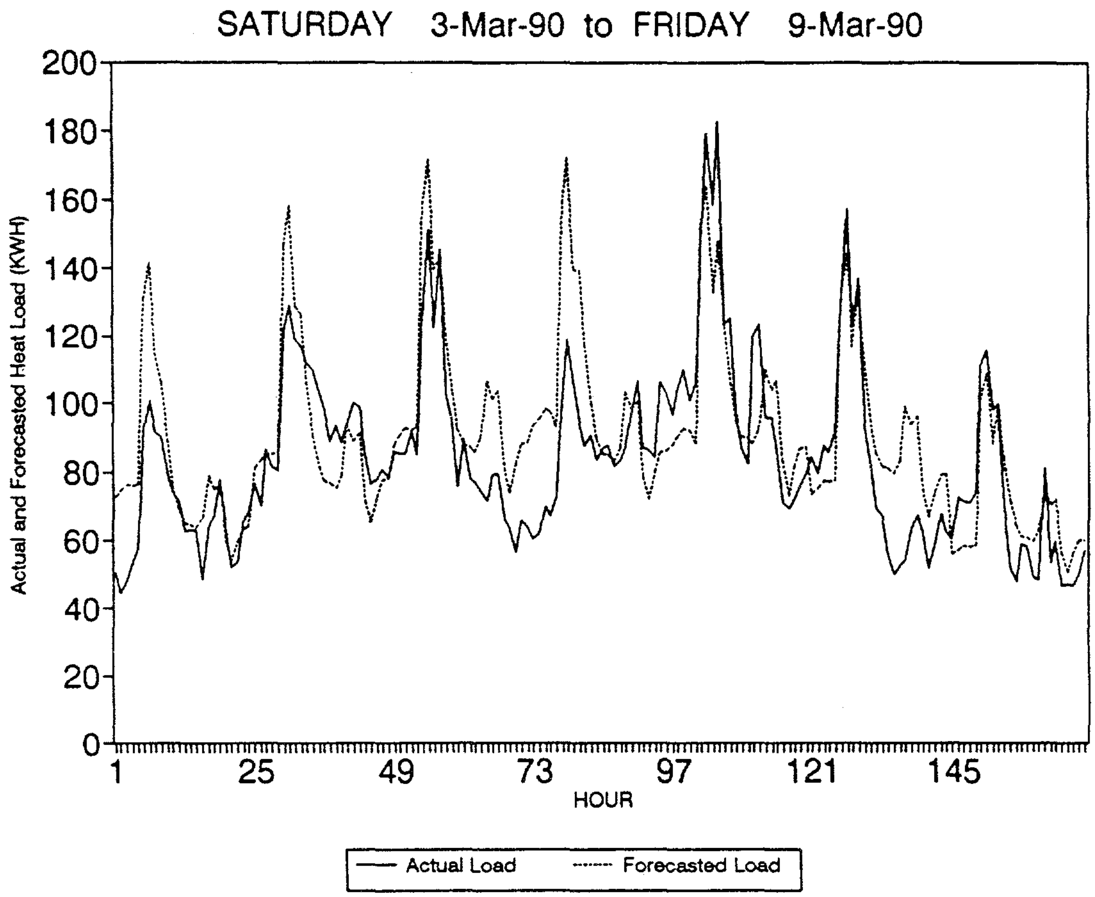

Figure 4-3. Comparison of Actual and Forecasted Heating Load for Brewster Office Building.

minimum) of tank temperature decreased progressively from Monday, until the lowest tank temperature was reached, on Wednesday evening (170 °F) just before charging started. The previous peak of tank temperature was reached on Wednesday morning at 223 °F. Therefore, due to the large amount of the storage, there was significant inter-day storage carryover from the less expensive weekend hours to expensive weekday hours.

Figure 4-2 shows the hour-ahead prices as they unfolded during the week. Charging decisions closely tracked the times of lowest prices. For example, on Saturday and Sunday the system waited until the early morning hours for the lowest prices instead of charging at 10 P.M. The system also acted to optimize as much as possible across days. For example, Tuesday night/Wednesday morning had only four low-priced hours, so the system charged for only a few hours that night. This led to the decline in tank temperature to its lowest level the following day, followed by an increase to 220 °F on Thursday morning when prices were lower than Wednesday morning.

Decisions of this type were based on hour-ahead price forecasts, and therefore turned out to be not quite optimal after the fact. Charging schedule was also influenced by the load forecast and the status of the storage system at the time of scheduling.

Some of the unscheduled daytime charging can be attributed to inaccurate load forecasts due to imprecise sensors or the rigidity of the load forecast algorithm. Therefore, further work on the development of a better self-tuning and feedback-based load forecast algorithm, together with the installation of multiple and more precise sensors, would minimize the incidence of unscheduled charging.

Performance of the Brewster Office Building

Two measures of performance of RTP-based control were the improvement in performance of RTP-based control compared to conventional TOU-based control and the improvement compared to a conventional heating system with no storage (i.e., NSC). TOU-based control and NSC behaviors were simulated during the weeks in which RTP-based control was actually running. The simulation method slightly overstated the effectiveness of conventional TOU-based control, producing a conservative estimate of performance. Results are shown in figures 4-4 to 4-9.

The NSC results were based on simulating an equivalent heating system with no storage system. The underlying assumption was that in a NSC, the heating loads were met concurrently when they occurred. Thus, the NSC weekly electricity usage pattern was proportional to the heating load pattern. In the NSC

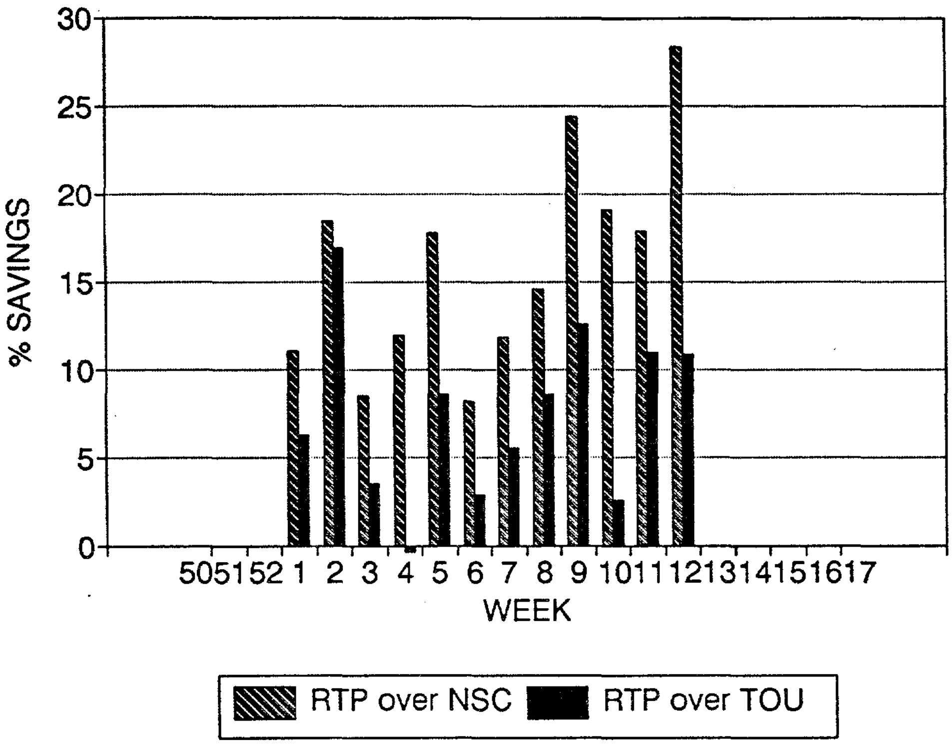

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Actual Hourly Prices (AH).

Figure 4-4. Weekly First-Season Savings for the Brewster Office Building (Based on Actual Hourly Prices).

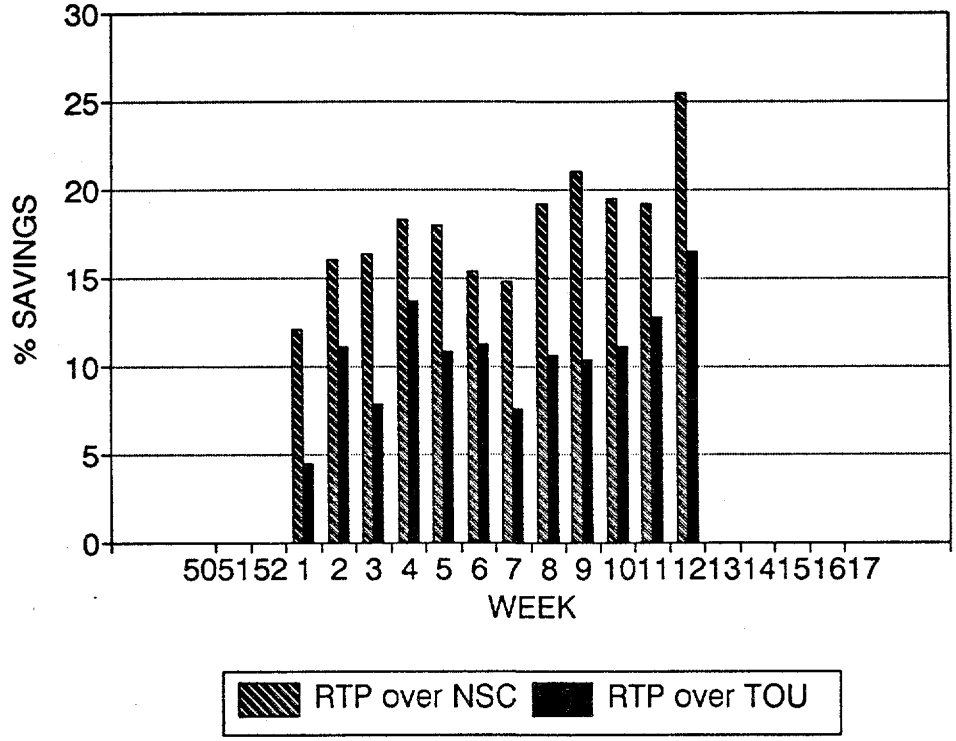

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Hour-Ahead Prices (HA).

Figure 4-5. Weekly First-Season Savings for the Brewster Office Building (Based on Hour-Ahead Prices).

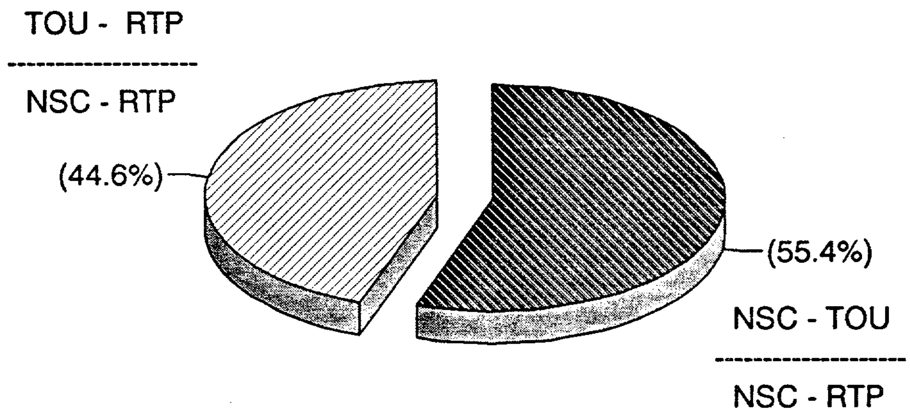

Figure 4-6. Share of First-Season RTP and TOU Savings over the NSC for the Brewster Office Building (Based on Actual Hourly Prices).

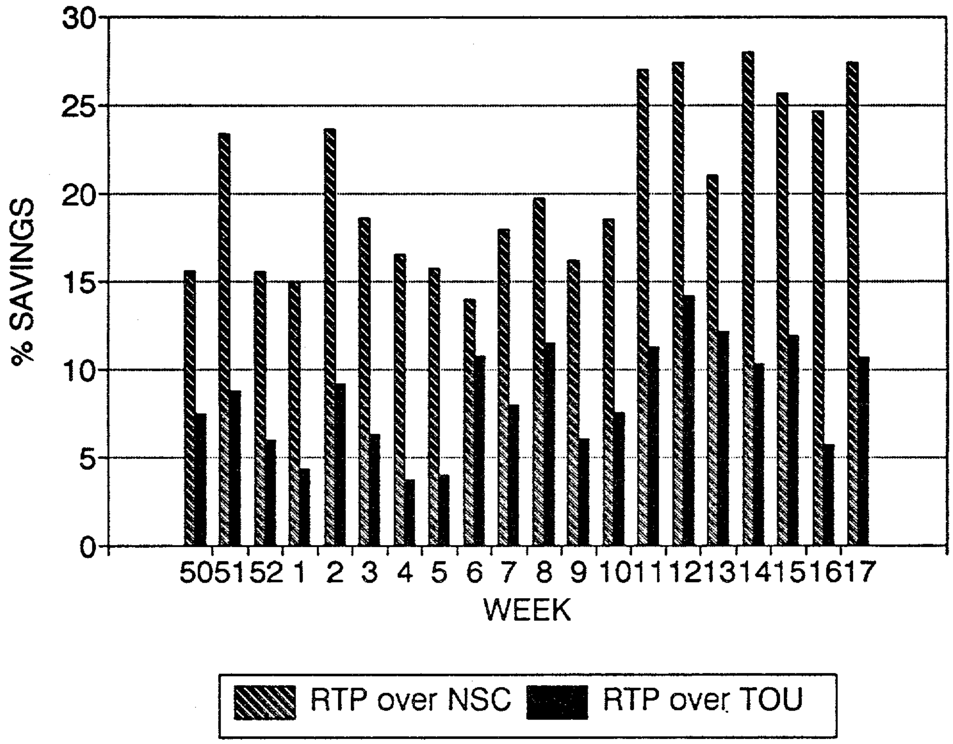

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Actual Hourly Prices (AH).

Figure 4-7. Weekly Second-Season Savings for the Brewster Office Building (Based on Actual Hourly Prices).

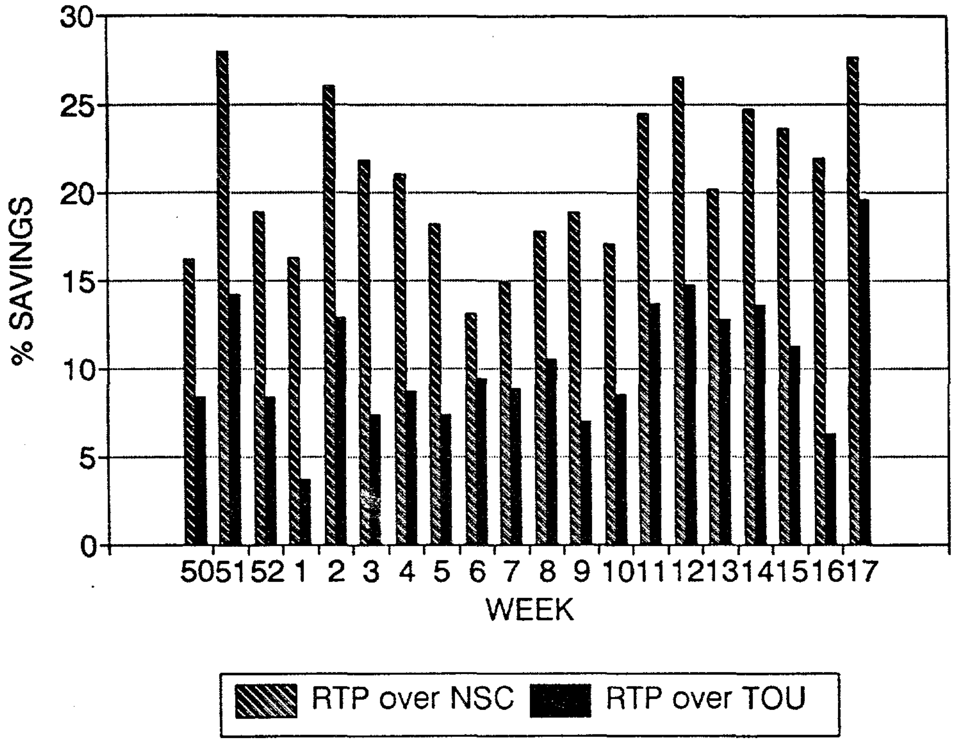

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Hour-Ahead Prices (HA).

Figure 4-8. Weekly Second-Season Savings for the Brewster Office Building (Based on Hour-Ahead Prices).

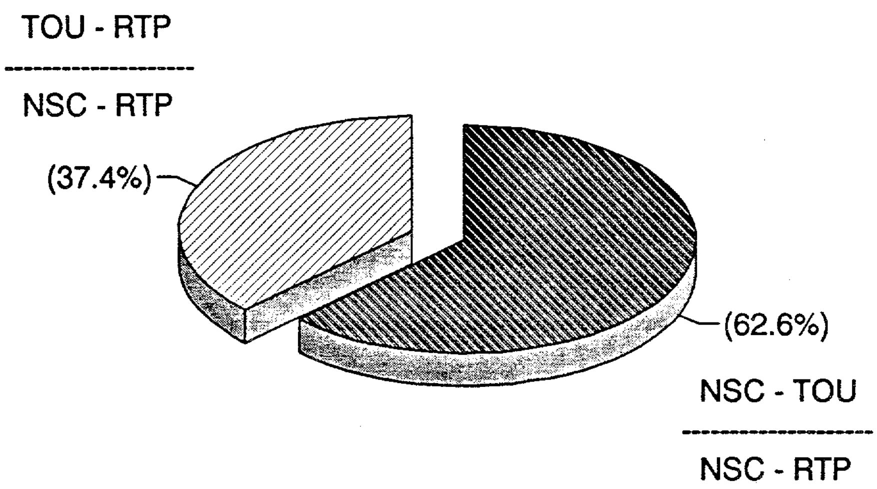

Figure 4-9. Share of RTP and TOU Second-Season Savings over NSC for Brewster Office Building (Based on Actual Hourly Prices).

simulation, the heating load pattern for a week was scaled so that its total for the week equalled the total actual electricity used under the RTP-based control.

The TOU-based control simulation assumed TOU off-peak operation at maximum hourly charge when the storage maximum temperature was based on the outdoor temperature reset, and a 10 °F temperature dead band. The outdoor temperature reset rule used in the simulation was based on the regression analysis of the actual TOU operations early in the experiment. The heating load used in the simulation was based on the actual heating load measured during RTP-based control. The time horizon was a week, 10 P.M. Friday to 10 P.M. the following Friday. The initial tank temperatures used in the simulation were equal to the actual tank temperature measured at the start of the corresponding weeks. The resulting electricity-use pattern was again scaled so that the total simulated electricity use was equal to the actual total electricity used under RTP-based control.

The TOU percentage saving was calculated by subtracting the total RTP cost from the total TOU cost and dividing it by total TOU cost. The calculation for NSC was similar. The total costs were the costs to the customer if the corresponding price was used for billing. Demand charges were not considered. Prices were adjusted for losses but not for revenue reconciliation. In the actual RTP operations, a 10 mills/KWH revenue reconciliation adder was used. This shifted the pattern of prices upward by raising the absolute level of prices used in the control algorithm. However, since the control algorithm is based on the relative level of prices, it did not affect the actual control schedule.

During the second season, on some occasions the RTP-based control was overridden to test the estimation and prediction algorithm enhancements. As a result, some daytime charging occurred. In the performance evaluations, the data for these days were adjusted to show the real impact of RTP-based control for those days.

A useful comparison is how much was saved by RTP-based control over conventional TOU-based control of the same storage system, measured using the prices seen by the customer (hour-ahead forecasts of the actual hourly RTP). These savings are expressed as a percentage of the cost of operating the system using TOU-based control. Thus, in week nine of the first season, RTP-based control saved approximately 11% over TOU-based control. This site was on a simplified control system (deterministic algorithm only, no adjusting controller) for week one, which accounts for the fact that the savings that week were the lowest of the 12-week experiment. Major improvements to forecasts and other equations used in the control of the Brewster office building were made during the first three weeks. Minor improvements continued to the last week of the project.

The week-by-week savings show variations for a number of reasons. Most importantly, some weeks had RTP prices which gave greater opportunities for savings. For example, if the prices were perfectly flat for a week, then neither conventional TOU- nor RTP-based control of the storage system would produce a benefit for that week. Similarly, if the least expensive hours fell between 10 P.M. and 2 A.M. every weekday, the RTP-based control would have chosen the same hours as conventional TOU-based control. Simulating the same week with slightly different starting conditions might have given somewhat different results due to the effect of temperature dead-bands. These results would have averaged out over time.

For evaluation of performance under RTP-based control, total seasonal savings should be more meaningful than weekly savings, because the different factors affecting the results would have averaged out[13]. The savings resulting over the two seasons of RTP-based control are shown in tables 4-1 and 4-2.

As can be seen, the savings to the utility of RTP-based control of the electric thermal storage over TOU-based control constituted more than one-third of the total savings compared to the NSC. Since the cost of electronics and communications was a small fraction of the hardware costs of storage plumbing and tanks for this site, the results show that RTP-based control was a cost-effective enhancement to the storage system.

[13] Overall patterns of weather and prices will be different from year to year, leading to differences in the benefits of RTP from year to year.

Table 4-1. First-Season Performance of the Brewster Office Building

Electricity used under RTP control Week 1 of 1990 to Week 12 of 1991	112,000 KWH
Costs under No Storage Case (NSC)	$3979
Costs under TOU control (determined by simulation)	$3636
Costs under RTP as incurred (measured by actual RTP)	$3360
(NSC costs - TOU costs)/(NSC costs)	8.6%
(NSC costs - RTP costs)/(NSC costs)	15.6%
(TOU costs - RTP costs)/(TOU costs)	7.6%
(TOU costs - RTP costs)/(NSC costs - RTP costs)	44.6%

Table 4-2. Second-Season Performance of the Brewster Office Building.

Electricity used under RTP control Week 50 of 1990 to Week 17 of 1991	178,879 KWH
Costs under No Storage Case (NSC)	$5411
Costs under TOU control (determined by simulation)	$4748
Costs under RTP as incurred (measured by actual RTP)	$4352
(NSC costs - TOU costs)/(NSC costs)	12.3%
(NSC costs - RTP costs)/(NSC costs)	19.8%
(TOU costs - RTP costs)/(TOU costs)	8.3%
(TOU costs - RTP costs)/(NSC costs - RTP costs)	37.4%

BREWSTER STOREFRONT/HYDROKINETIX SYSTEM

Operations History

Data logging at this site started on December 1, 1989. The site was put under TCA's control December 12, 1989. Initially, the storage charging schedule algorithm was based on a time clock operation that mimicked conventional TOU-based control. The outdoor temperature reset rule was implemented as a subroutine in the central computer in Cambridge. The RTP-based control was phased-in beginning on January 4, 1990. By January 21, the site was under RTP-based control with the adjusting controller and price extension.

The Brewster storefront RTP-based control system underwent continuous modifications due to problems with the averaging tank temperature probe, which started to function erratically at higher temperatures, especially when the heating elements were on. The problem was finally fixed during the week of February 2, 1990, by inserting a point sensor at mid-height of the tank. After testing and validation, the tank temperature reading used in the algorithms was switched from the averaging tank sensor to the point sensor on February 8, 1990. The tank temperature values from the point sensor were stable and followed the averaging probe values closely when the latter was functioning less erratically.

Throughout the experiment, every element of the software underwent continuous improvement, especially the control and price extension algorithms. The results for the first season showed improved performance after the system improvements were made.

The minimum and maximum temperature limits for the tank were set at 110 °F and 200 °F, respectively, on February 15, 1990. For security reasons, the algorithm was structured so that in a worst-case scenario some residual energy would be left in the tank; therefore, the tank temperature rarely got close to its minimum set point. The maximum and minimum of the hysteresis dead band were set at T_{Heur} (target temperature from the adjusting controller) and (T_{Heur} - 4) °F, respectively, on February 13, 1990.

The second season covered the period from early January 1991 to the end of April 1991. Most of this season was spent implementing, testing, and troubleshooting automatic estimation and predicting state equation coefficients and total load parameters. An important development in this season was the absence of any tenant at this location. TCA set the thermostat at a constant temperature which was in effect for every hour of every day. The effect was to increase the heating requirement for this site. Thus, the relative storage size was smaller during the second season.

Hourly Behavior of the Brewster Storefront

Figure 4-10 shows hourly behavior of the Brewster storefront for Tuesday March 6, 1990. Figure 4-11 shows the hourly behavior for week nine of the project (Saturday, March 3 through Friday, March 9, 1990). Figure 4-12 provides a comparison of the actual and forecasted heating load for week nine. For this week, the site charged 885 KWH. Most of the charging during the week was from 1 A.M. to 6 A.M., which produced a quick increase in the tank temperature. This should be compared to TOU operation, which commenced charging at 10 P.M. every week night. The demand for heat was confined to times when the store was open. The decrease in the tank temperature at these times was less steep than the tank temperature increase during the charging periods.

Only during the first three days did the tank temperature approach its maximum limit of 200 °F. This corresponded to days of less expensive prices (Saturday to Monday) than during the rest of the week. In addition, the nightly tank temperature did not fall below 155 °F from Saturday to Monday. However, during the rest of the week, the nightly tank temperature always decreased to 130 °F. Because of the small size of the storage system, the weekend charging could not be used to meet more than a day's heating requirements, i.e., there was not much *flexibility* in *intra-day* scheduling.

Although the small tank size limited improvement in performance, the high charging rate enabled the tank to require charging only a few hours each night, which provided flexibility in selecting the few least expensive night hours. Most of the savings were due to price differentials during off-peak periods. Thus, even if limited storage size does not allow much in terms of inter-day scheduling, a high charging rate makes it possible to take advantage of price variations even within a short interval corresponding to the TOU off-peak period.

Performance of the Brewster Storefront

The performance of the heating system in the Brewster storefront under RTP-based control was compared to the results of NSC and TOU-based control simulations. The comparisons are shown in figures 4-13 to 4-18. In 1989-1990, only the results of the later weeks are representative of RTP-based control due to testing and problems encountered during the beginning of the experiment. Week one started Saturday January 6, 1990. The tank temperature sensor was replaced at the end of week four. Therefore, the first four weeks of operation were based on erratic and erroneous tank temperature measurements. In the second season, access to the site for re-setting the on-site system was delayed by a few weeks resulting in sub-optional operations.

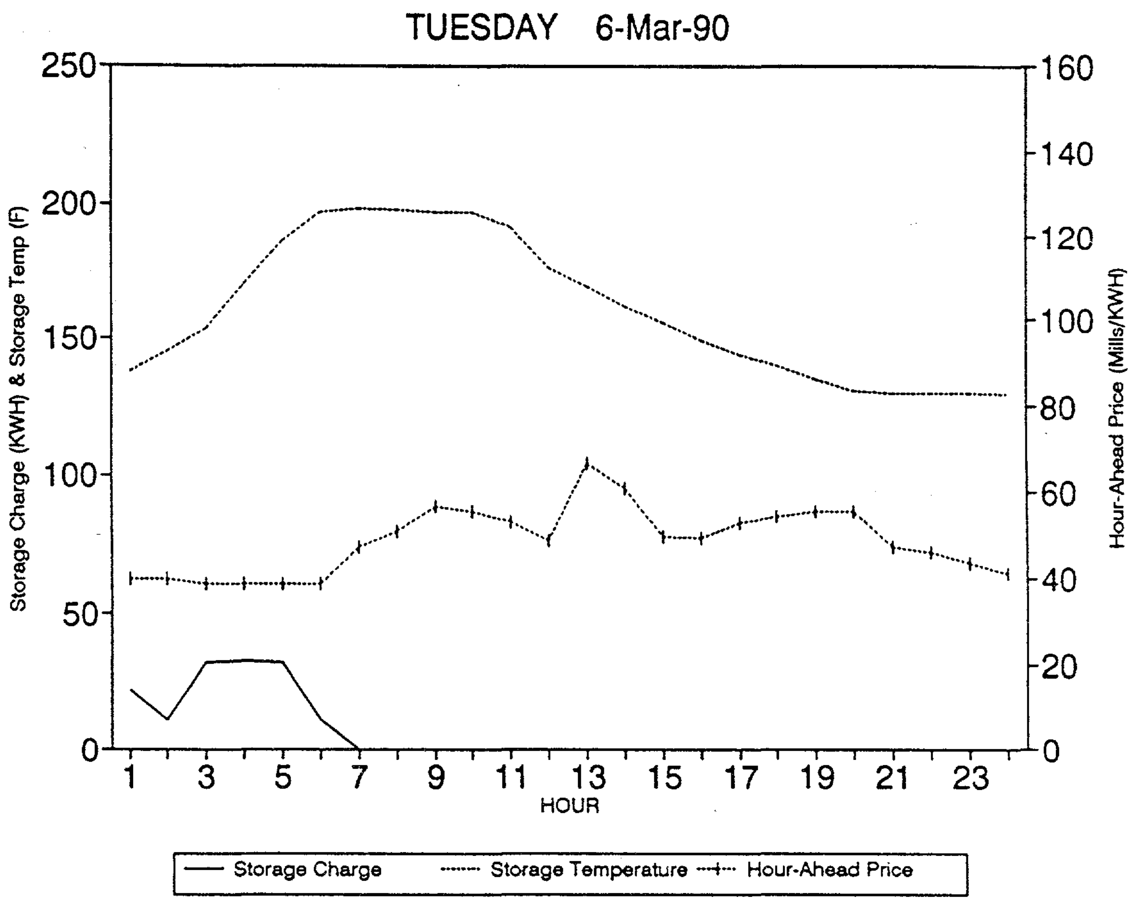

Figure 4-10. Hourly Behavior of the Brewster Storefront.

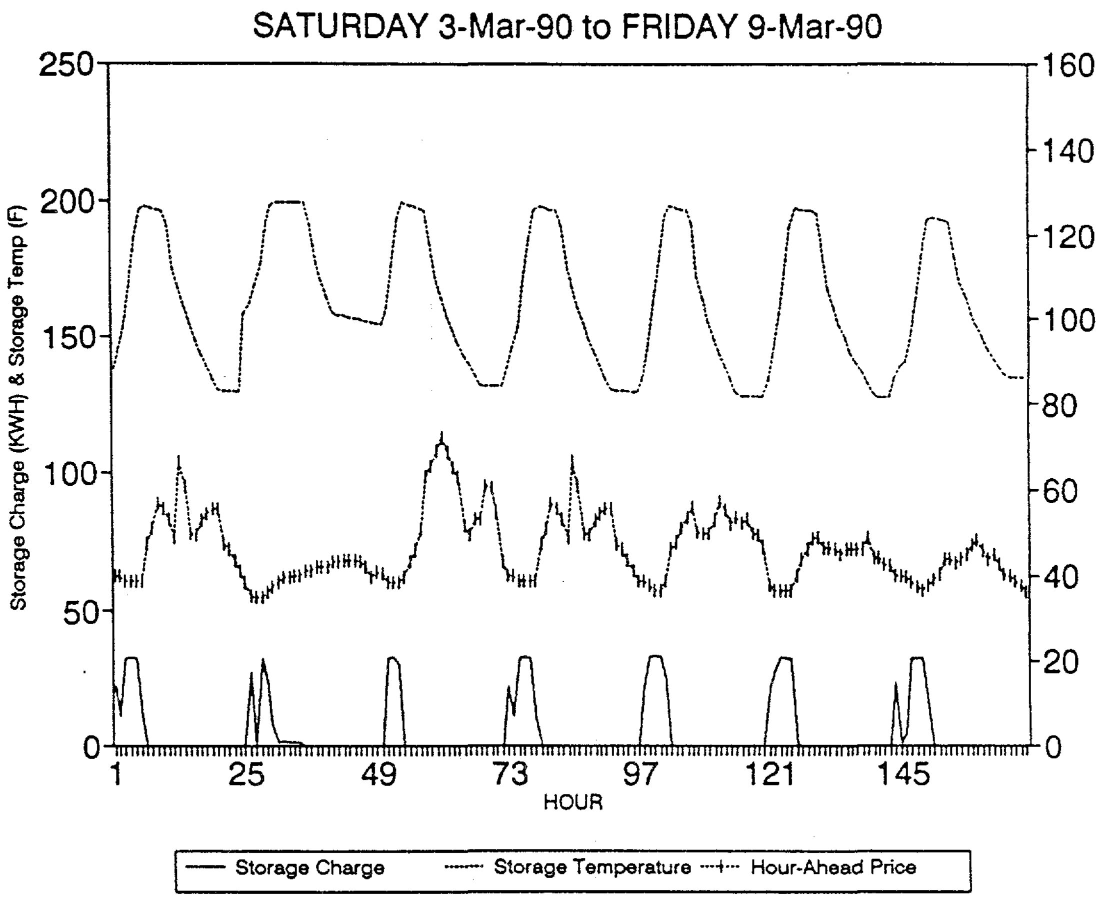

Figure 4-11. Weekly Behavior of the Brewster Storefront.

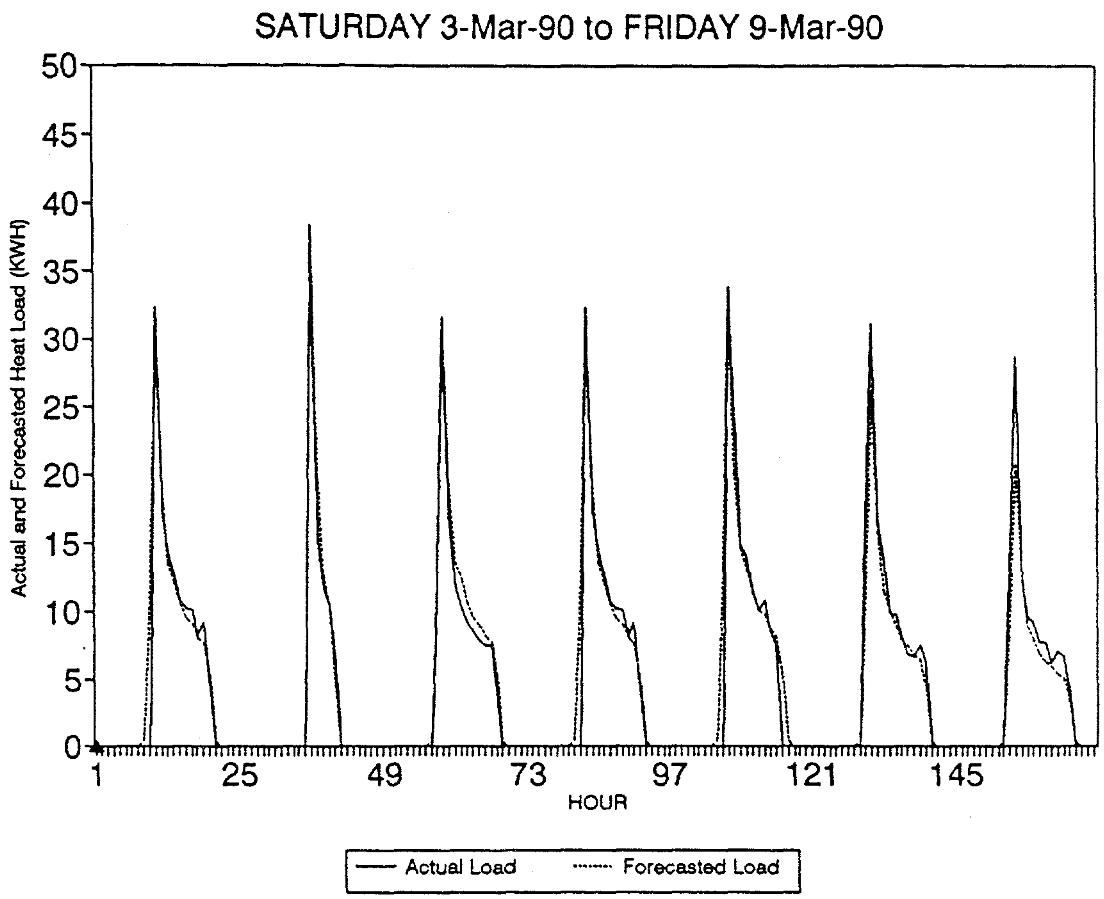

Figure 4-12. Comparison of Actual and Forecasted Heating Load for Brewster Storefront.

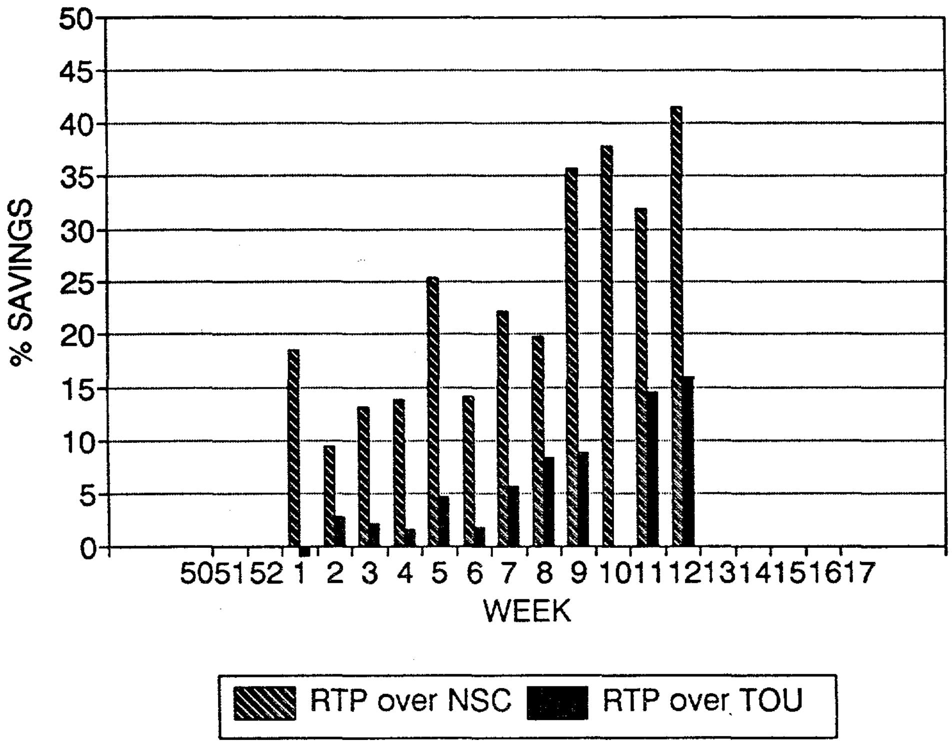

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Actual Hourly Prices (AH).

Figure 4-13. Weekly First-Season Savings for the Brewster Storefront (Based on Actual Hourly Prices).

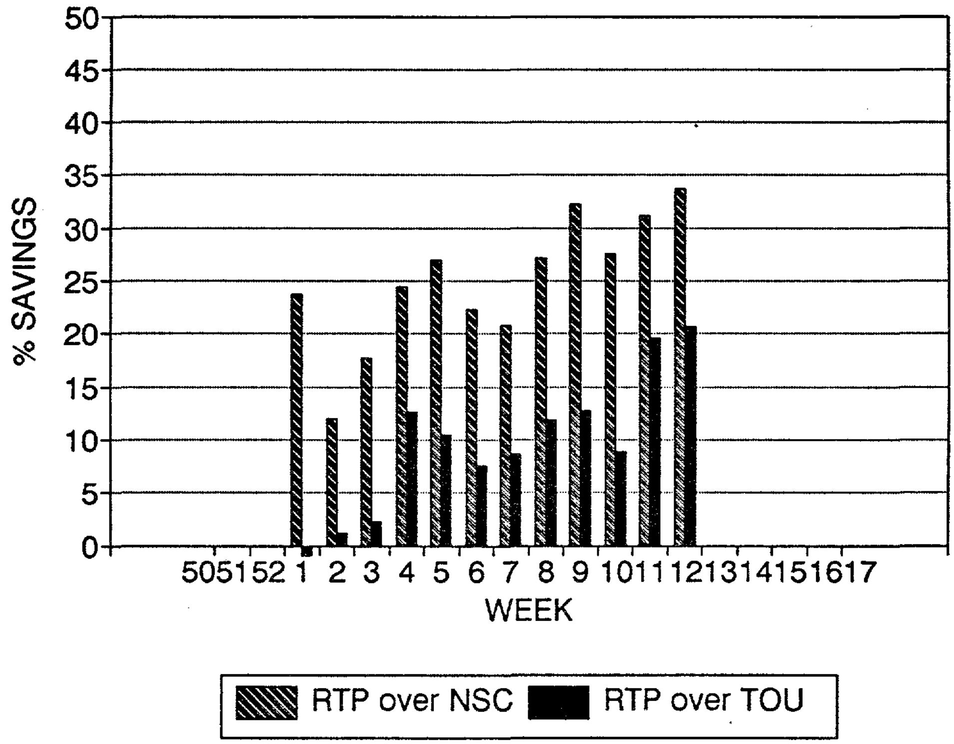

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Hour-Ahead Prices (HA).

Figure 4-14. Weekly First-Season Savings for the Brewster Storefront (Based on Hour-Ahead Prices).

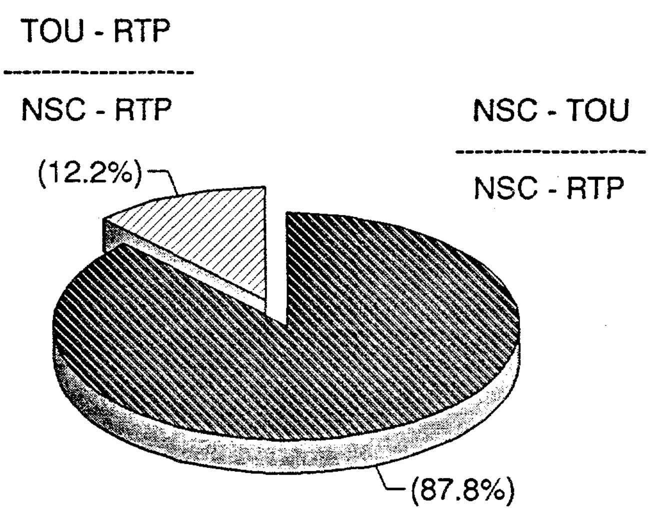

Figure 4-15. Share of RTP and TOU First-Season Savings over NSC for the Brewster Storefront (Based on Actual Hourly Prices).

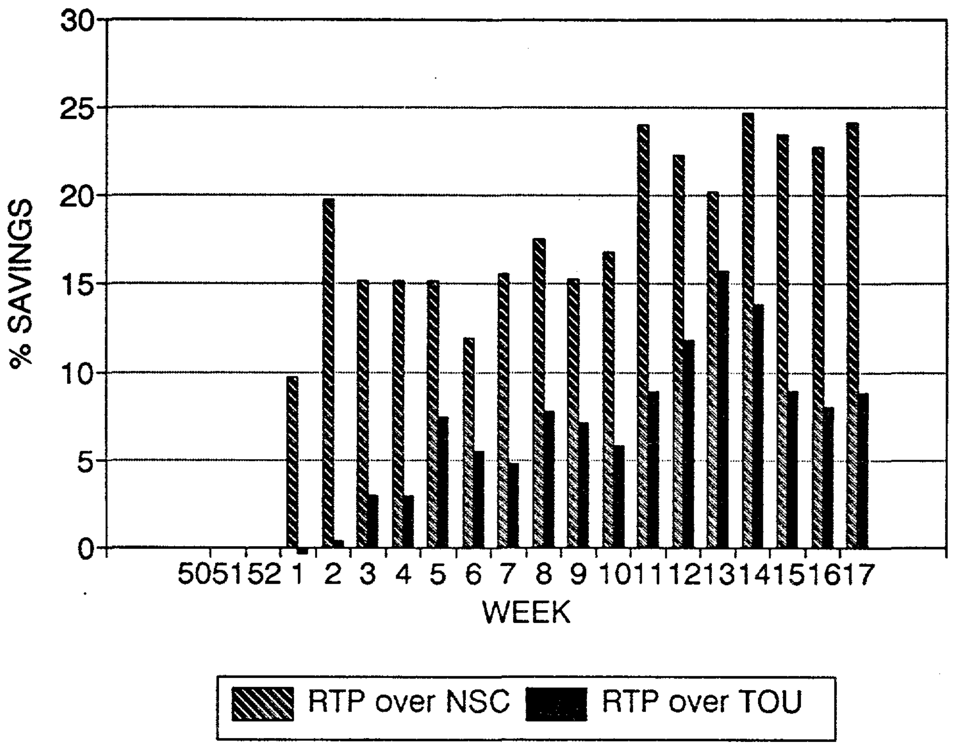

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Actual Hourly Prices (AH).

Figure 4-16. Weekly Second-Season Savings for the Brewster Storefront (Based on Actual Hourly Prices).

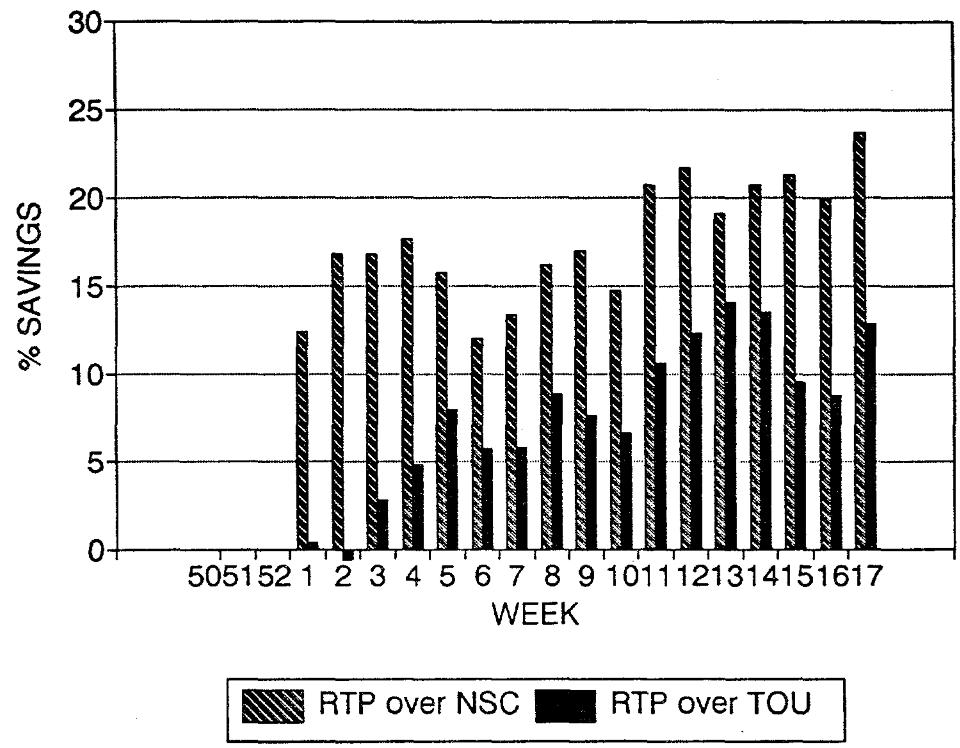

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated No Storage Case (NSC) and Time-of-Use (TOU) response. The savings were calculated using the Hour-Ahead Prices (HA).

Figure 4-17. Weekly Second-Season Savings for the Brewster Storefront (Based on Hour-Ahead Prices).

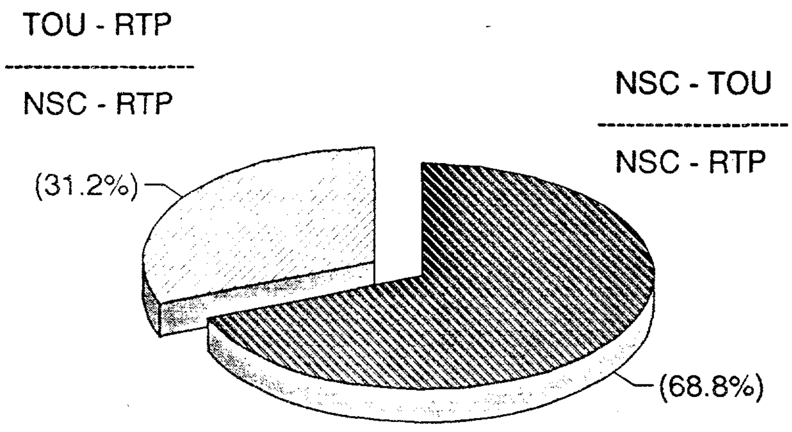

Figure 4-18. Share of RTP and TOU Second-Season Savings over NSC for the Brewster Storefront (Based on Actual Hourly Prices).

Compared to the utility's actual cost of service under TOU-based control, The RTP-based control resulted in 8.7% savings in the first season and 6.2% savings in the second season. Compared to the total savings over NSC, the savings to the utility of RTP-based control of the ETS over TOU-based control was 21% in the first season and 31% in the second season. These results are shown in tables 4-3 and 4-4.

Table 4-3. First-Season Performance of the Brewster Storefront.

Electricity used under RTP control Week 7 of 1990 to Week 12 of 1990	5,600 KWH
Costs under No Storage Case (NSC)	$152
Costs under TOU control (determined by simulation)	$115
Costs under RTP as incurred (measured by actual RTP)	$105
(NSC costs - TOU costs)/(NSC costs)	24.3%
(NSC costs - RTP costs)/(NSC costs)	30.9%
(TOU costs - RTP costs)/(TOU costs)	8.7%
(TOU costs - RTP costs)/(NSC costs - RTP costs)	21.3%

Table 4-4. Second-Season Performance of the Brewster Storefront.

Electricity used under RTP control Week 1 of 1991 to Week 12 of 1991	9,438 KWH
Costs under No Storage Case (NSC)	$270
Costs under TOU control (determined by simulation)	$238
Costs under RTP as incurred (measured by actual RTP)	$223
(NSC costs - TOU costs)/(NSC costs)	12.0%
(NSC costs - RTP costs)/(NSC costs)	17.4%
(TOU costs - RTP costs)/(TOU costs)	6.2%
(TOU costs - RTP costs)/(NSC costs - RTP costs)	31.2%

PLATTSBURGH OFFICE BUILDING/EARTH STORAGE

Operations History

First season data logging of the Plattsburgh office building started in early November 1990. The site was put under a TOU-based control emulation (TCA activated TOU-control instead of site active) on December 21, 1989. From January 9 to January 17, 1990, the site was put back on its own conventional TOU-based control to collect more data on the building's thermal characteristics. Afterward, an initial form of RTP-based control was implemented. The control system was upgraded with new methods of price extension and an empirical charging requirement/outdoor temperature relationship, was put in place February 12, 1990 (middle of week six).

The second season of RTP-based control of the Plattsburgh office building covered the period from the end of December 1990 to the end of April 1991.

Hourly Behavior of the Plattsburgh Office Building

Figure 4-19 illustrates the hourly behavior of the Plattsburgh office building for Tuesday, March 6, 1990. Figure 4-20 illustrates the behavior in week nine (Saturday March 3 to Friday March 9, 1990) of the first season. An interesting feature of this site was the slow dynamics of the storage system, as was observed in the sluggish variation of the storage variables (deep-sand, bottom-slab, and top-slab temperatures). The deep-sand temperature was the slowest to change. During week nine, it maintained an average temperature of 81.5 °F, varying within a one degree. In contrast, the top slab, which was affected mostly by the indoor or indirectly by the outdoor temperatures, ranged from 69.5 °F to 75.5 °F. The relatively small variations in temperature were due to the large thermal mass of the earth-storage system.

Time delays were another feature of earth-storage systems. It took several hours before any appreciable increase in the storage temperatures was detected in response to heat input into the ground. As a result, the storage temperature continued to increase for several hours after the storage heating had stopped. This phenomenon, and the narrowness of the temperature range, caused some problems in controlling the storage temperature by setting temperature dead bands for any of the storage temperature variables. In the case of the deep sand, the temperature variations were too small to be used as control variables. In the case of the bottom and top slab, the temperatures kept increasing above the maximum point of the dead band even after the heat input had been stopped. The time delays for the top slab were the longest, since it was located furthest from the heating elements.

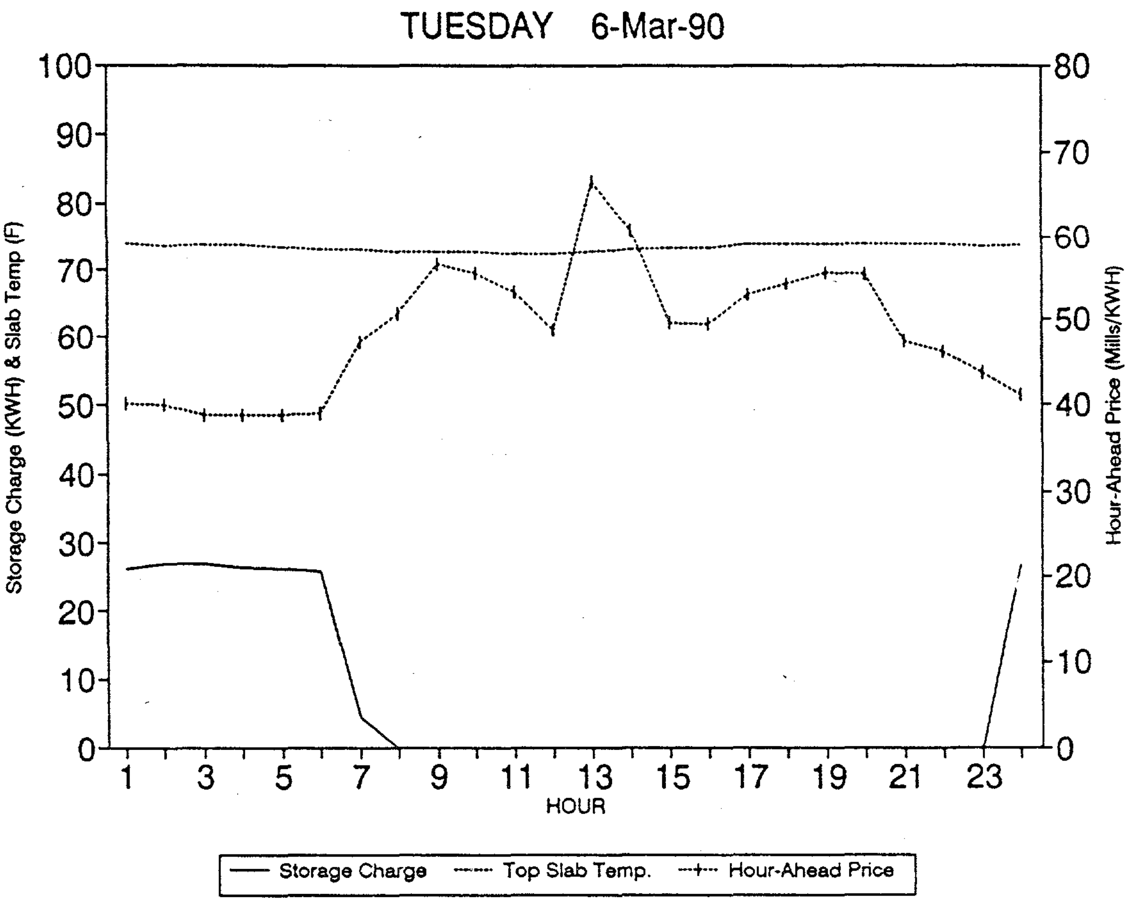

Figure 4-19. Hourly Behavior of the Plattsburgh Office Building.

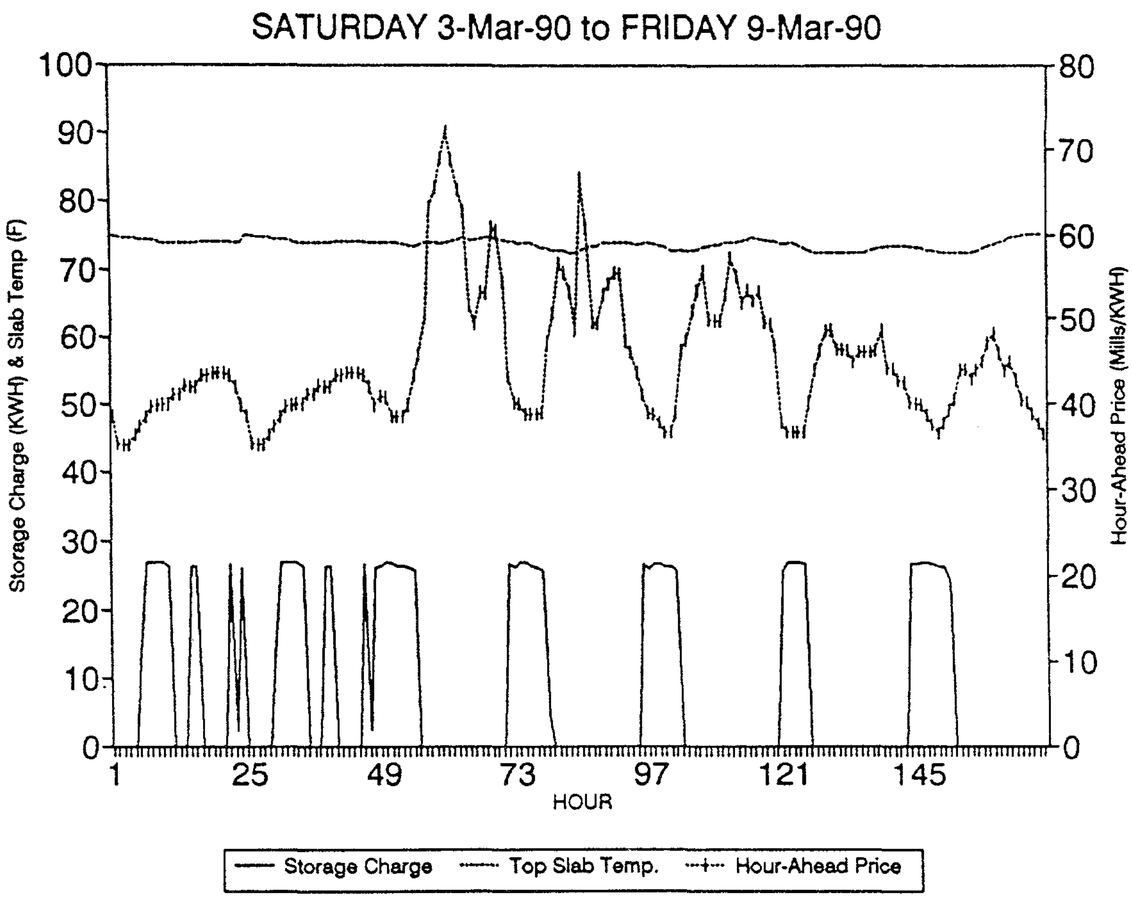

Figure 4-20. Weekly Behavior of the Plattsburgh Office Building.

Another phenomenon was that indoor temperature, on the average, was near that of the top slab temperature, although the swings in its pattern closely followed outdoor temperature. This was expected, since the heating system did not respond immediately to changes in the outdoor temperature. The amplitude of indoor swings smaller than the change in outdoor temperatures due to the action of electric baseboard heaters. The baseboard heater set points were individually established by the occupants of each room. For week nine of 1990, the outdoor temperature ranged between 2 °F and 44 °F, while the indoor temperature ranged between 68 °F and 79 °F. Occasionally, as on Friday, March 9, when the outdoor temperature reached higher levels, the indoor temperature increased above the slab temperature, probably due to decreased heat transfer to the outdoors, accentuated by internal heat generation from people and equipment.

Performance of the Plattsburgh Office Building

Because of constraints in modeling and control of the thermal behavior of the Plattsburgh office building, it was decided that, given the short time horizon of this experiment, a mathematical programming approach for RTP-based control was impossible. It was decided to use a simpler routine to determine the charging schedule as described in Section 3. Consequently, it was expected that, under implemented control rules, no significant improvement in economical performance would be achieved. The weekly results shown in figures 4-21 to 4-24 confirm this. After some improvements in the software and control rules, some slight system improvements were noted. It is possible to improve the economic performance using more sophisticated building modeling tools and control techniques. However, this would require a more detailed and protracted analysis of the site data. The performance could also be improved by using a more accurate charging requirement/outside temperature relationship in conjunction with the present control method.

The performance analysis for the Plattsburgh office building differed from those for the Brewster buildings. Since the storage system in the Plattsburgh office building was not thermally enclosed (i.e., did not have an insulated boundary), practical measurement of the building's heating load was difficult. Therefore, no corresponding NSC simulation was possible for this building. The TOU simulation for the Plattsburgh office building was based on summing up the actual number of charging hours each day (starting at 10 P.M.) incurred under RTP-based control, and finding the TOU equivalent by stacking up same number of charging hours beginning at 10 P.M. Therefore, for the Plattsburgh office building, total simulated TOU-based control charging hours equaled those of the actual RTP-based control charging hours on a daily basis. All TOU simulation charging results were scaled so that the weekly KWH for TOU-based control was equal to the actual weekly KWH for RTP-based control.

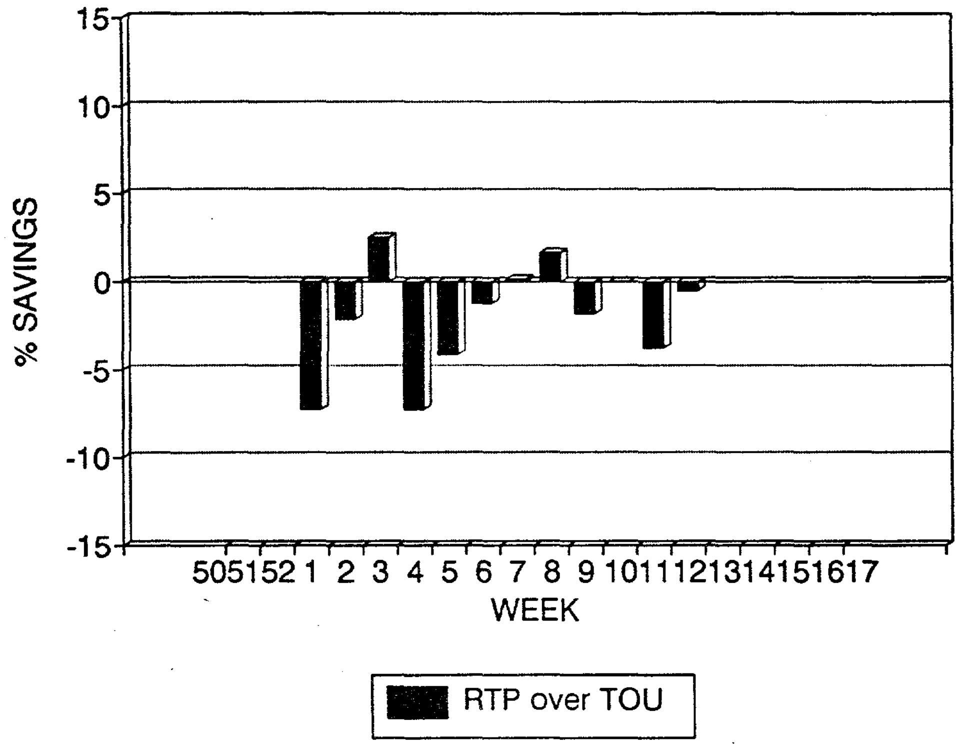

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated Time-of-Use (TOU) response. The savings were calculated using the Actual Hourly Prices (AH).

Figure 4-21. Weekly First-Season Savings for the Plattsburgh Site (Based on Actual Hourly Prices).

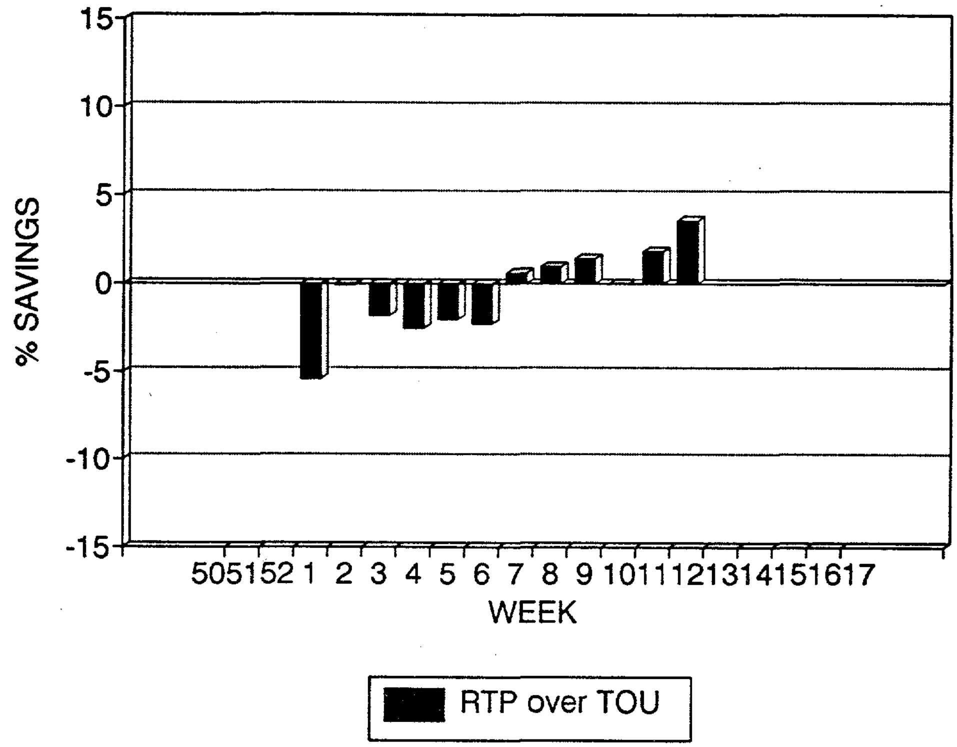

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated Time-of-Use (TOU) response. The savings were calculated using the Hour-Ahead Prices (HA).

Figure 4-22. Weekly First-Season Savings for the Plattsburgh Site (Based on Hour-Ahead Prices).

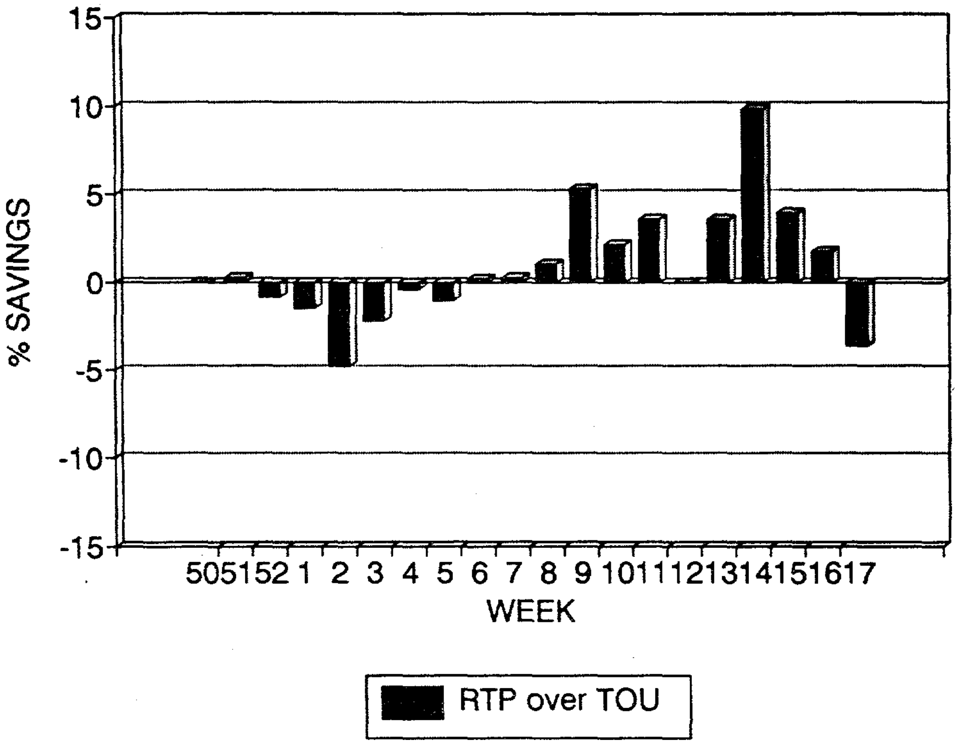

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated Time-of-Use (TOU) response. The savings were calculated using the Actual Hourly Prices (AH).

Figure 4-23. Weekly Second-Season Savings for the Plattsburgh Site (Based on Actual Hourly Prices).

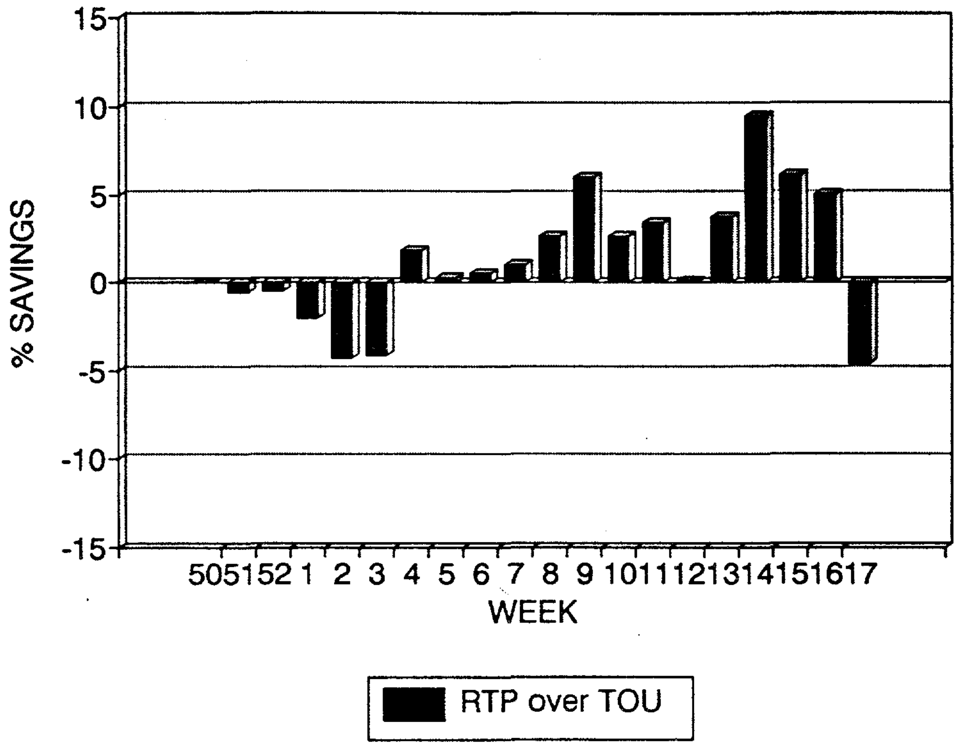

Note: Each column shows the weekly percentage of savings of actual RTP response with respect to the simulated Time-of-Use (TOU) response. The savings were calculated by using the Hour-Ahead Prices (HA).

Figure 4-24. Weekly Second-Season Savings for the Plattsburgh Site (Based on Hour-Ahead Prices).

The TOU percentage savings for the Plattsburgh office building were calculated by subtracting the total RTP cost from the total TOU cost and dividing by total TOU cost. The results were based on the logged charge values, which did not match the meter values. Therefore, the costs had to be scaled. The electricity usage by electric baseboard heaters was not measured - there were too many units - and thus, their impact on performance was not factored in.

The control rules were based only on the day-ahead and hour-ahead price forecasts. The actual hourly prices were not the basis for the storage charging scheduling since they were calculated later by the utility. Given the small percentage of savings based on hour-ahead price forecasts, it was not surprising to see negative savings based on the actual hourly prices. More accurate price forecasts would have resulted in higher savings.

It became apparent near the end of the experiment that the comfort of the inhabitants at the Plattsburgh office building had improved, due primarily to a better determination of the next day's heating requirements using a weather forecast, rather than because of RTP. The determination used a control based on number of hours to be charged instead of a control based on temperature. A similar improvement under TOU-based control would require computation of required number of hours for charging on a daily basis rather than mechanical thermostat action.

Section 5

IMPACT OF SIZING AND GENERAL DISCUSSION OF COSTS

IMPACT OF DIFFERENT STORAGE SIZES ON COST OF OPERATIONS

This section addresses the issue of storage sizing under RTP rates by presenting the results of a simulation study. Parameters of the simulation, as well as building loads and prices used as inputs, are based on the results of the experiment on the RTP-based control of ETS done in the 1989-1990 winter season. The simulation study was carried out for water-based storage systems, which included the Brewster office building and Brewster storefront. Results indicate that, under RTP-based control, increase in storage sizes lowers utility cost of providing service, in contrast to TOU-based control, where the utility's cost of service remains relatively unchanged. Because modeling the earth storage system (Plattsburgh building) proved to be difficult, no simulation was done for that type of storage.

Principles of Storage Sizing Under RTP and TOU Rates

The principles which govern sizing of storage under RTP are quite different from those conventionally used for sizing under predetermined TOU prices. These differences arise due to the more flexible way in which RTP-based systems respond to changing weather and changing prices in their day-to-day operations. The RTP system discussed here was developed to take maximum possible advantage of all available storage capacity, even if it was not needed to meet the next day's load. It also was capable of "searching" for the cheapest possible hours to do its charging, rather than beginning to charge at the same time each night. In contrast, a TOU-based system would store enough energy to meet the next day's load, then stop. Any storage capacity not needed to meet the next day's load was effectively unused and without economic value that day.

There are two types of storage size: *storage tank capacity* (X_{max}, measured in KWH) and *storage charging rate* (U_{max}, measured in KW). The function of tank capacity is to hold thermal energy across time, so that a larger tank capacity means that it takes more hours to empty the tank, and thus that it holds more thermal energy across time. In contrast, charging rate determines how many hours it takes to *fill* the storage tank with a given number of KWH.

The principles underlying TOU rates are discussed first. The method generally used for sizing TOU-based systems is the "design day" concept. Historical weather data and site characteristics are used to estimate a "worst case" load level, $W_{design\text{-}day}$. The storage tank capacity is simply set to this worst case:

$$X_{max,TOU} \geq W_{design\text{-}day}$$

The storage charging rate is set to the size needed to fill the tank during the off-peak period:

$$U_{max} = (X_{max,TOU}) \div T_{off\text{-}peak}$$

where $T_{off\text{-}peak}$ is the number of hours in the off-peak period during the week-night interval of the TOU rate, usually 10 hours. This number is fixed under typical TOU rates offered by utilities. Since almost all days are milder than the design day, these sizing principles mean that on most days neither the tank capacity nor the charging rate will be fully used. In fact, during the fall and spring, the system will often use less than half its capacity. The system will begin charging under time clock control when low prices begin (e.g., 10 P.M.), and stop charging when the tank reaches the predetermined storage level $W_{forecast}$ + safety margin. In principle, this predetermined level should be based on a weather forecast, but in practice TOU systems use a simple calculation called "outdoor temperature reset." Under outdoor temperature reset, the maximum allowable storage temperature is varied linearly as the current outdoor temperature changes. For a colder outdoor temperature, the maximum storage temperature is set at a higher level. Therefore, the current outdoor temperature is used as a proxy for the next day's predicted load.

There is no benefit under TOU rates to enlarging X_{max} or U_{max} beyond the above equations. Assuming the design-day capacity ($W_{design\text{-}day}$) was properly calculated, the additional capacity rarely would be used. In contrast, under RTP-based rates, none of these rules is optimal for either the utility or the customer. Consider optimal behavior for an existing system with a given level of U_{max} and X_{max}. Although the exact algorithm is more complex, its basic pattern on a normal week-night is simple. For a given amount of charge (X_{target}) to be put into the tank, the algorithm searches for the cheapest T hours during the night where $T = X_{target} \div U_{max}$. The smaller the T, the lower the cost of injecting X_{target} of energy, since the system is using cheaper hours in the price duration curve. This has two consequences:

1. Raising the charging rate (U_{max}) always reduces operating costs, since it allows more concentrated filling of storage, and

2. The RTP-based system will take advantage of days with low demand (X_{target}) by changing behavior to further minimize costs.

The only limit to these effects is that when U_{max} becomes so large that storage can be fully charged in a single hour, further increases in U_{max} have no incremental benefit.

Determination of the amount of energy to inject (X_{target}) is more complex, but the fundamental difference from TOU-based response is that under RTP the system tries to store energy across several days if doing so will reduce costs. For example, in the simulation runs made for this report, the algorithms usually filled the tank completely on Sunday night, even if the demand forecast for Monday was quite low. This happened because Saturday and Sunday night prices were usually lower than the prices forecast for the following nights. The larger the tank size (X_{max}), the more the system can search for low prices several days before the energy will actually be needed. Thus, *raising* X_{max} *always reduces expected operating costs,* up to the point that the extra capacity would carry past the forecast time horizon. Of course, on cold days, when demand W approaches X_{max}, this ability to store across more than one day is reduced.

This analysis confirms that increasing storage capacity and charging rate beyond TOU levels provides benefits. What happens if the system is sized somewhat below TOU sizing? Under TOU-based control, there is no penalty on most days, but on days of extreme weather and high demand, the customer loads are not met, unless charging is allowed during peak periods, where it carries the risk of incurring a high demand charge. In that case, charging would start as soon as storage runs out. Under RTP, a building with undersized storage may also require charging in the daytime. However, under RTP-based control, the control algorithm decides the required number of periods for daytime charging and plans charging for the least expensive daytime hours. For some days, the control algorithm may decide to start charging in the early hours of daytime, so that storage can be used during more expensive hours in the afternoon.

Finally, the effective size of a storage system can be raised by converting from TOU-based control to RTP-based control. Thus, if a building is enlarged or its demand increases for other reasons, the cost of retro-fitting to provide RTP-based control could be less than the cost of adding more storage tanks. Even if the resulting system violated the design day constraint, the operating costs might be lower than before, while design day weather conditions could still be met by the system.

Simulation of Storage Operation Under RTP- and TOU-Based Control

For the storage sizing simulation, the state equations of the first season were used. These are:

a) for the office building:

$$X[k+1] = X[k] + 0.025\ U[k] - 0.017\ W[k] - 0.40$$

b) and for the storefront:

$$X[k+1] = X[k] + 0.47\ U[k] - 0.45\ W[k] - 0.32$$

where $X[k]$ (in °F) is storage value at period k; $U[k]$ (in KWH) is electrical charging at period k; and $W[k]$ (in KWH or proportional to it) is the building heating load at time k. The state equations were constructed using data from both TOU- and RTP-based control operations.

The state equations together with historical RTP rates and the building heating loads were used in the simulation to determine the hour-by-hour operations of the storage systems for each site under the given loads and prices. All the main algorithms used in the actual RTP-based control were incorporated into the simulation. The simulation accessed the price and weather forecasts with no advance knowledge of future weather or price data. The simulation was done for all the weeks when prices were available, i.e., from late December 1989 to late March 1990.

The RTP-based simulation was validated using the actual results from the experiment. Results of RTP-based and TOU-based control simulations were compared to the NSC using the data for the actual heating loads gathered during the experiment. In both the actual control and in the simulation, RTP-based control was carried out by varying the maximum storage temperature setting (thermostat) for each hour. The optimal time trajectory of the maximum tank temperature was determined by the optimization algorithms.

The Weather Service forecast data are normally given for the next two and three days. These were stored in the weather database during the experiment. The simulation used the load forecasting method and data as used in the experiment.

After a desired and optimal trajectory for the tank temperature was determined by the optimization algorithms, the hourly temperature trajectory was used as the maximum temperature setting (thermostat) for the storage. The storage tank module in the simulation was then subjected to the actual heating load, which resulted in the simulated tank charging schedule, which could be different from the optimal tank charging schedule determined by the optimization algorithms. This was because the optimization algorithms used load forecasts rather than actual loads as input.

A parallel TOU-based control simulation was performed for comparison of results. The simulation did not compare operations under various rates. The only prices considered for cost calculations were the hour-ahead RTP prices supplied by the utility. These were the utility's forecast of its actual marginal costs for the next hour. At this stage, these prices did not include the quality of supply premium, and were therefore understated. RTP-based control provided optimal scheduling of ETS charging while completely meeting building heat load requirements. This required frequent optimization runs for time horizons of two to three days. The rule for TOU-based control was to schedule ETS charging during off-peak hours only. The amount of daily charge could be varied using outdoor reset mechanisms. The cost calculations for each control mode were done using hour-ahead prices. Therefore, the costs of TOU-based control operations reported were unrelated to the actual cost to a similar customer under an actual TOU rate. However, hour-ahead prices were close to the utility's actual hourly prices, and thus reflected the cost of customer operations to the utility. Costs and cost comparisons reported here reflect the utility's costs to serve its customers, not the customers' own costs. Any proper evaluation of DSM programs would require consideration of the utility's actual hourly costs to provide service.

Operational Savings

Many variables affect the relative size of the savings under RTP-based control, such as price pattern or values of control parameters that were not based on optimal design considerations in the experiment. The present study considered only the two main storage attributes. These are storage thermal capacity (X_{max} [in KWH]) and charging rate (U_{max} [in KW]). In the simulation, the present size of each storage (thermal capacity and charging rate) was chosen as the base case. Then, different values of storage capacity and charging rate were selected and simulations were run for 13 weeks. Different sizes were simulated by multiplying the state equation coefficients by the appropriate values.

For the Brewster office building, the base case U_{max} was 280 KW, and X_{max} was 75°F (the difference between maximum and minimum allowable temperatures), so that the time to fill the tank with no demand was 10.5 hours. For the Brewster storefront, the base case U_{max} was 16 KW, and X_{max} was 80°F, so that the time to fill the tank was 5.2 hours. These times were calculated from the state equation coefficients used in first season operations and in the simulation.

The utility's total cost of service was determined by evaluating the costs of RTP, TOU, and NSC behavior under hour-ahead prices provided by the utility. NSC heating load values were based on the direct and on-time charging which satisfied the actual heating loads. Results are shown in figures 5-1 to 5-4. The savings with respect to NSC operations was evaluated by dividing the difference between RTP (or TOU) costs and NSC costs by the NSC costs. For comparison, the base case total cost for the Brewster office building RTP operations was $4,964 for the 13 weeks. The Brewster storefront base case total cost for the same period was $384. These figures underestimated the benefits because the prices did not include the quality of supply premium, and design and setup of the hardware and software for the experiment were not fully optimized. Under a fully developed real-time price rate and RTP-based control system, higher savings should be realized.

Results as shown in Figures 5-1 to 5-4, refer to percent reduction in the utility's cost of service under RTP and TOU compared to NSC:

- For all cases and for both buildings, the benefits of RTP-based control were higher than those of TOU-based control. Considering the base case storage sizes, the RTP savings for the Brewster office building were 24% compared to 12% for TOU savings; therefore, the value of ETS to the utility is increased by 50% under RTP;

- Under TOU-based control, below a certain size the storage was incapable of meeting the load at all times unless charging was allowed during peak periods. Under RTP-based control, the storage size could be appreciably smaller (especially at the Brewster storefront) and still show significant savings compared to the TOU case because, although charging might be required during the daytime, this would be planned to occur during the least costly hours;

- At larger storage sizes, benefits of TOU-based control do not necessarily increase, as can be seen for both sites where a higher charging rate (U_{max}) did not improve the TOU benefits. One explanation is that higher charging rates resulted in more concentrated charging hours at the beginning of off-peak periods, which did not necessarily coincide with the least expensive hours;

- Under RTP-based control, there was a clear relationship between the savings and the storage size attributes. Increasing storage capacity (X_{max}) and charging rate (U_{max}) increased savings. Larger storage capacity provided opportunities for storing across days, and higher charging rate resulted in fewer hours of charging at low-cost hours; and

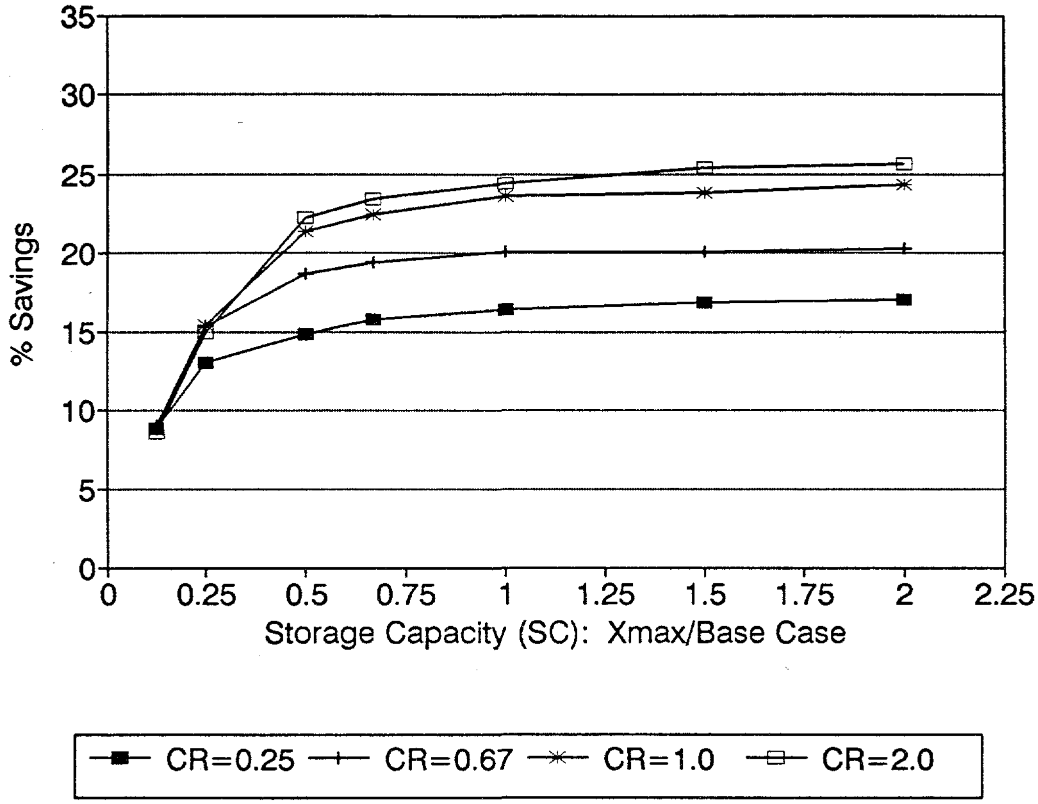

Figure 5-1. Impact of Different Storage Sizes on RTP versus NSC Savings for the Brewster Office Building.

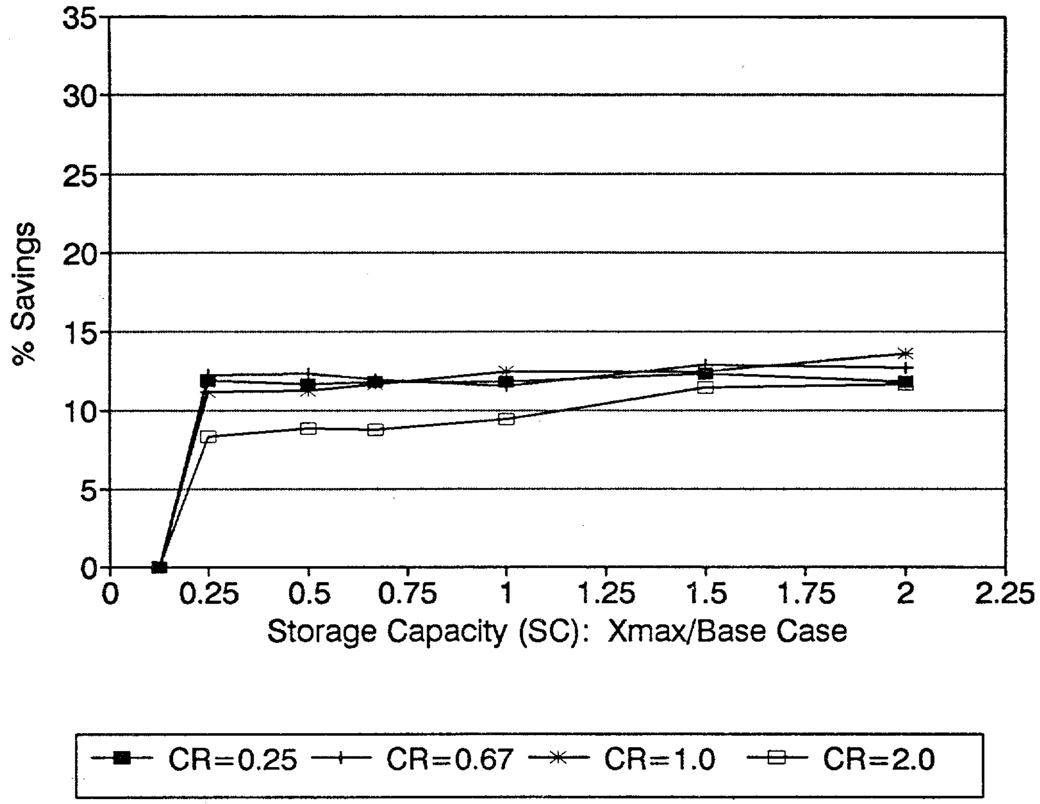

Figure 5-2. Impact of Different Storage Sizes on TOU versus NSC Savings for the Brewster Office Building.

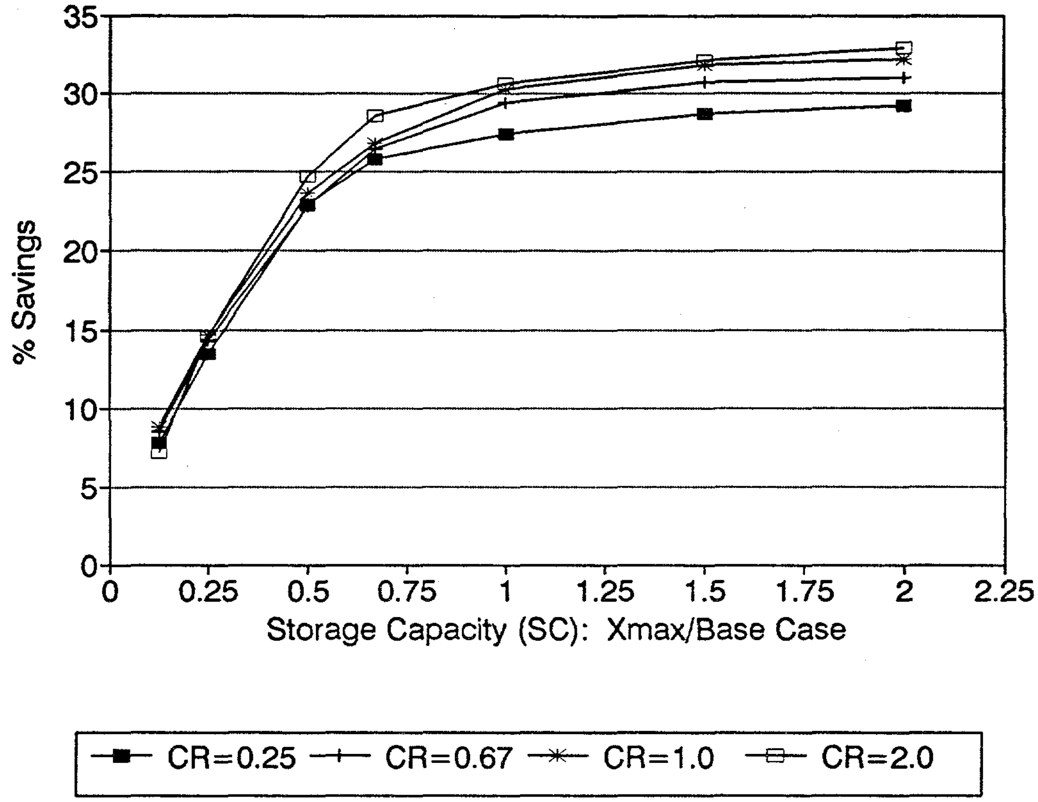

Figure 5-3. Impact of Different Storage Sizes on RTP versus NSC Savings for the Brewster Storefront.

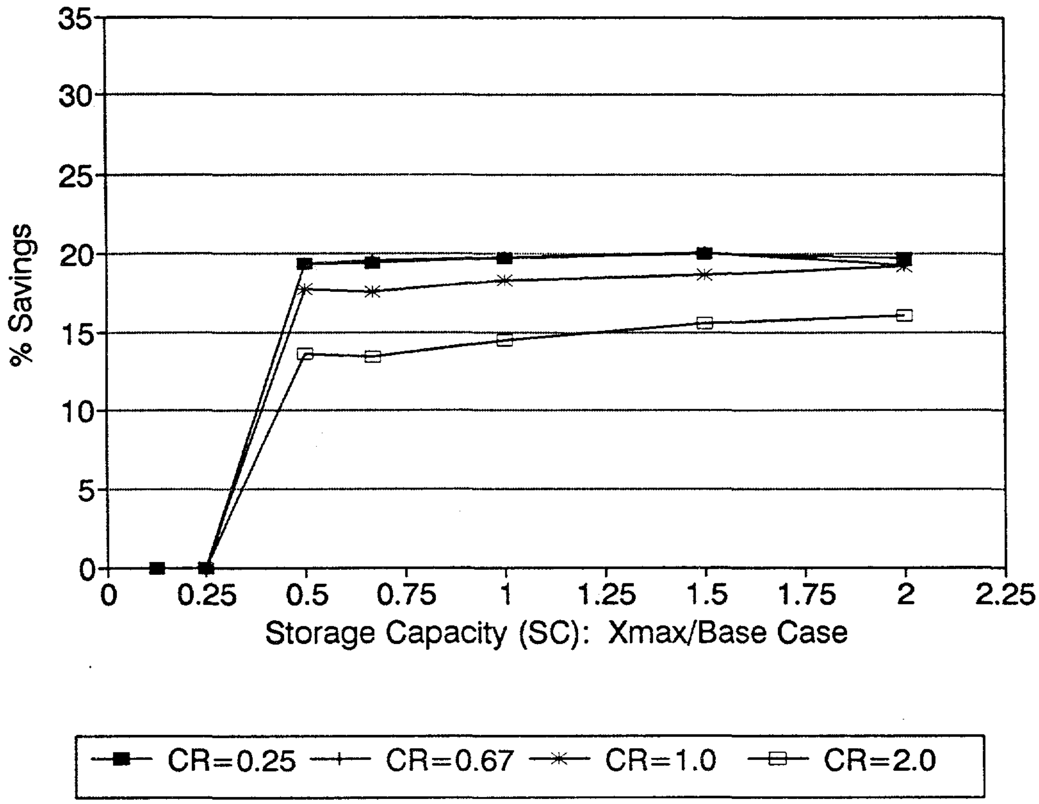

Figure 5-4. Impact of Different Storage Sizes on TOU versus NSC Savings for the Brewster Storefront.

- Results indicate that savings under RTP increase with a diminishing return at larger storage sizes. For storage capacity, the optimal size is limited by the length of the time horizon used in forecasting and optimization algorithms. For the charging rate, the savings depend on the steepness of the price duration curve for the length of the time horizon at its low price end.

These results illustrate the impact of different storage sizes on the utility's cost of providing service. Quantitative results are only valid for the particular systems and utility prices studied here. However, the results can be generalized to other systems and utilities in a qualitative manner. Further work in this area requires analysis of savings under different price and building load patterns.

Storage-Sizing Conclusions

Results of the simulation study indicate that under RTP-based control, the size of the ETS system has a substantial impact on utility costs of providing service, as opposed to TOU-based control, where the impact of ETS size is much less significant. Under RTP, higher utility savings can be realized for both undersized and oversized ETS systems. As a result, design day criteria cannot be the sole basis of ETS sizing under RTP rates. A complete analysis of economic tradeoffs would require knowledge of the individual utility's price pattern, the building load pattern, the utility's rebates and incentives, and the particular type of ETS system offered by the manufacturer and its site-dependent accessory and transportation requirements.

Capital Costs versus Size

Market penetration of ETS systems is driven by electric utility incentives and rebates. Customers seldom plan the use of thermal storage. According to ETS manufacturers (based on phone conversations), there is no simple relationship between the capital cost of storage and its size. Official price quotations are not provided. With smaller systems (e.g., the Brewster storefront), larger sizes of the same type of ETS system cost more, but cost increases are not drastic. For example, sometimes heating elements with higher charging rates can be substituted without changing the type of storage system, and at no significant cost. In fact, some control and heating components are interchangeable among systems of different size but the same type. Major cost increases are due to changes in types of equipment, where existing systems are replaced with more complex components or different configurations. In case of larger storage systems (e.g., the Brewster office building), if the increase in size is simply due to the lengthening of the storage tank with no increase in the diameter, then the cost increase may be insignificant. In addition, immersion of heating elements with higher charging rates can be relatively inexpensive. However, costs go up with an increase in the diameter of the storage tank. Larger storage systems must be custom-made, and prices are negotiated separately, so there is no simple method of illustrating the price-size relationship. Another factor affecting costs in a significant way is the location of the storage system. Costs for outdoor systems may be 35% to 40% higher than for indoor systems as a result of storage thermal losses. In addition, the amount of BTU stored versus capital costs increases with larger size. However, a limiting factor is the cost of shipment of the storage system, where field fabrication is impractical. Therefore, for larger storage systems, every unit is individually priced.

OPERATING COSTS UNDER RTP- AND TOU-BASED CONTROL

Operating savings of the larger water-based system used in the Brewster office building under RTP-based control increased by nearly 40% over savings under TOU-based control, against the base case NSC. As shown in table 4-2, TOU-based control resulted in a nearly 12% reduction in operating costs with respect to the NSC, while RTP-based control showed a nearly 20% reduction.

As shown in Table 4-4, the savings at the smaller water storage site (the Brewster storefront) were also impressive in percentage terms, but small in absolute terms because the customer's total bill was not large. Also, the effective storage size for the Brewster storefront was smaller for the second season. For this site, it is likely that any implementation of RTP would combine the control of several individual storage systems at the mall to realize economies of scale in communications and logic.

INCREMENTAL CAPITAL COSTS FOR RTP-BASED CONTROL

For new and existing storage heating and cooling installations, the cost of system enhancement with RTP-based monitoring and control should be compared to the savings that would result. The following is an analysis of costs of retrofitting, based on our experience with RTP-based control of heat storage in the two commercial buildings with hot water storage. It is expected that as more experience is gained with the methodologies and as a result of greater market penetration, the per-unit cost of retro-fitting or new system installation will decrease. A full analysis of costs and savings depends on such factors as utility price patterns and the size of thermal storage, and merits further study. However, the following presentation should serve as a starting point for such an analysis.

During the experiment, savings from the use of RTP-based control over TOU-based control for hot water thermal storage systems reached 15% for some weeks. It is expected that with further experiments and improvements in the components of RTP-based control, and better price forecasts from the utility, the total seasonal savings will approach 20%. In terms of actual customer costs, savings can be much higher if day-ahead prices are used for direct billing, because the optimization process would be deterministic and there would be less need to readjust the charging schedule every hour. Savings of more than 20% (RTP-based control versus TOU-based control) were observed for a few weeks when day-ahead prices were used for costs calculations. However, day-ahead prices are less accurate forecasts of utilities actual hourly costs compared to the hour-ahead prices. Therefore, 10% to 20% can be considered a plausible range for the savings of a mature RTP-based control over TOU-based control. For a smaller system, e.g., the size of the Brewster storefront, this represents an annual savings of $20 to $60 per year (10% of $200 to 20% of $300). For a larger system, such as the Brewster office building, this represents an annual savings of $500 to $1,400 per year (10% of $5,000 to 20% of $7,000).

Defining the "mature costs" of a hardware system in today's market is difficult. In addition to development costs (hardware and algorithms), the cost of an RTP system includes the costs of the price calculation, communications, and site (sensor, loggers, analysis and control equipment) components. Sensor requirements for RTP-based control of thermal storage are given in Table 5-1. Sensors and communication equipment are accounted for in the capital costs. Price communications also contributes to operating costs. The cost estimates in tables 5-2 and 5-3 show a range based on limited experience from other operational systems. They provide a first estimate of the cost-effectiveness of RTP in specific applications. Because the system is intended to be fully automated, building operator effort should be minimal and is not included in tables 5-2 and 5-3. However, some operator training for equipment maintenance and troubleshooting would be required.

For RTP-based control to be cost-effective, the savings compared to conventional control must be greater than the hardware and software costs. Given the utility's criteria for cost effectiveness, i.e., programs which, on a 20-year time horizon, are cost-effective at 11% of the cost of capital, an RTP project with annual savings (in constant dollars) ranging from $188 to $1,256 or more (range depends on the capital costs as shown in table 5-3) would be cost-effective to the utility. The illustrative example is given in tables 5-4, 5-5, and 5-6. If one considers a 10% increase in savings due to RTP-based control as a conservative estimate of the minimum level achievable with RTP, then customers with costs (based on the utility's actual hourly cost) greater than $1,880 to $12,560 (range depends on the capital costs) per heating season will show positive benefits with RTP. As discussed, this includes the potential applications of RTP to sites such as the Brewster storefront if multiple heat storage systems at that site are centrally dispatched and controlled. If all of the individual stores in the mall were controlled by a centralized RTP system, the total savings might be sufficient to justify the monitoring and control hardware. For the results of a shorter payback period, see table 5-6.

Table 5-1. Sensor Requirements for RTP-Based Control of Thermal Storage.

Data[a]	Units	Used in non-RTP-Based Control[b]	Transformation
Day/Time	—	Yes[c]	—
Indoor Temperature	°F	Yes[d]	average[e]
Outdoor Temperature	°F	Yes[f]	average
Tank Inside Temperature	°F	Yes[g]	average
Tank Return Temperature	°F	No	average
Tank Supply Temperature	°F	No	average
Circulation Pump Flow Rate	GPM	No	average
Pump Run Time	min	No	sum[h]
Tank Electrical Consumption	KW, KWH	Yes	integrate[i]
Charging Run Time	min	No	sum
Building Energy Transfer	KWH (BTU)	No	integrate[j]

[a]Data recording interval is every 15 minutes to 1 hour. Data storage requirements are for two to four weeks.

[b]At least one utility now requires data from supply, return, flow rate, and electrical consumption sensors as a prerequisite for utility financial support of thermal storage systems.

[c]Used under TOU rates for timing storage charging.

[d]Used for indoor thermostat control, and available if used in direct digital control.

[e]Variables are averaged for the recording interval.

[f]Often used for economizer control or outdoor reset mechanism for storage tank or supply water temperatures.

[g]Used for control of thermostat-based charging of storage.

[h]Used in demand control under demand charge rates.

[i]Variable is determined by the integration of a function of other variables during the recording interval.

[j]Cooling may require two additional temperature sensors across building cooling coils.

Table 5-2. Components of Utility RTP-Based System Cost (Illustrative).

Component	Amount ($)
Price Calculation[a] (utility cost)	
Capital cost of computer, modems, etc.	3000
Cost of development and implementation of the software systems,	5000
TOTAL	8000
Per site, based on 100 sites.	80
Per site, based on 1000 sites.	8

[a]Estimates are for a stand-alone system in the utility control room based on the extension/automation of the present experimental system. This does not include costs associated with rate development. A system fully integrated into the control room may cost considerably more.

Table 5-3. Component Costs for Mature RTP-Based Control System (Illustrative).

Component	Amount[a] ($) Low	High
Communications		
Annual operating costs (phone system)	150	250
Case 1: Enhancement of a present-day EMS system[b]		
Additional memory and sensors (such as flow meters and watt transducers)	1500	3000
Total Case 1:	1650	3150
Case 2: Stand-alone[c] RTP-based monitoring and control system[b]		
AT Class PC w/hard disc and modem	1000	1500
Sensors for KWH, flow, temperatures	2500	4000
Data logger/controller	1000	2000
Software, installation and tuning	1500	2500
Total Case 2:	6150	10250

[a]Excludes research and development cost.

[b]Excludes any utility rebates.

[c]Can be used in conjunction with a previously installed EMS system.

Table 5-4. Estimated Customer Operating Costs and Benefits.

	Amount[a] ($) Low	High
Annual Operating Savings per Site[c]		
Small Systems[b,d]	20	60
Large Systems[b,e]	500	1400
Annual Net Operating Benefits per Site (Savings minus Annual Operating Costs)		
Small Systems	-240	-90
Large Systems	250	1250

[a]Excludes research and development costs.

[b]Only sizes in the range of subject experimental sites are considered here.

[c]Savings are the cost of TOU-based control less the cost of RTP-based control computed using utility's actual hourly prices. Savings under RTP (over NSC) are appreciably higher than under TOU.

[d]Storage systems with annual costs (using utility's actual hourly cost) of $200 to $300 under TOU-based control. Savings under RTP are assumed to be from 10% to 20% of costs under TOU.

[e]Storage systems with annual costs (using utility's actual hourly cost)of $5000 to $7000 under TOU-based control. Savings under RTP are assumed to range from 10% to 20% of costs under TOU.

Table 5-5. Estimated Programmatic Net Benefits, Present Value of Savings, and Net Present Value.

	Annual Operating Net Benefits ($)	Present Value[a] of Savings ($)	Net Present Value ($) based on Capital Costs Low	High
Low Savings (large system)	250	1991	- 8009	+491
High Savings (large system)	1250	9954	- 46	+8454

[a]Assumes 11% cost of capital over 20 years. Net benefits are assumed to be in constant dollars.

Table 5-6. Break-Even Savings per Year for RTP-Based EMS for a Given Payback Period.

	Payback Periods[a]		
System Capital Cost[b]	20 Years	10 Years	5 Years
Low ($1,500)	$188	$255	$406
High ($10,000)	$1,256	$1,698	$2,706

[a]Assumes 11% cost of capital over 20 years. The net benefits are assumed to be in constant dollars.

[b]Multiply the break-even savings per year values by 5 to 10 to get the approximate heating bill of the equivalent heat storage system under TOU rates (based on RTP savings of 10% to 20% relative to TOU heating bill).

Section 6
RESULTS AND CONCLUSIONS

RESULTS

The experiments done in the NYSEG service territory during the winters of 1989-90 and 1990-1991 demonstrated the economics and feasibility of RTP for control of ETS. The experiments involved real-time cost analysis, real-time price setting, implementation of algorithms for controlling ETS loads at customer sites based on RTP.

When setting prices, forecasted and actual hourly cost information was generated by NYSEG in the control room and translated into an hourly customer price. The dollar benefits of the RTP experiment were significant. Compared to TOU-based control, it was shown that with RTP-based control, the utility's total actual cost of service for the two water-based ETS systems was reduced by 8%. The savings to the utility of RTP-based control of ETS over TOU-based control constituted more than one-third of the total savings compared to NSC. Comparison of savings based on hour-ahead forecasts with savings based on actual hourly price suggests the value of improved hourly forecasts. The more accurately prices can be forecast, the greater the savings possible with RTP.

The savings at the smaller water storage site (the Brewster storefront) were also impressive in percentage terms, but were small in absolute terms because the customer's total bill was not large. For this site, it is likely that any implementation of RTP would combine the control of several individual storage systems at the mall to take advantage of economies of scale in communications and logic. Other identical systems at the same location could potentially be independently controlled by a single on-site PC.

Control of the Plattsburgh earth-storage system did not result in significant savings in comparison with those seen in the water systems, but resulted in increased customer comfort and satisfaction with the heating system. In the earth-storage system only limited savings are possible given the slow thermal transfer constants of the system and the fact that the system is passive, i.e., the heat discharged cannot be controlled and is released when it reaches the surface of the slab. The contribution of RTP control was primarily in terms of forecasting the desired load for the next day based on a temperature forecast. With this information, the quantity of heat needed in the slab to achieve the appropriate level of dissipation the following day could be calculated and programmed. Based on this forecasting ability, customer comfort levels could be more effectively maintained, but there was little economic effect from rescheduling the charging into hours of low cost. Savings from RTP in the Plattsburgh system were hard to evaluate because of the need to simulate how the system would have operated under conventional TOU-based control.

Conventional TOU-based earth-storage control tended to overheat the system when outdoor temperatures were warm, and under-heat the system when outdoor temperatures were cold, leading to corresponding swings in indoor temperature and comfort. However in the simulated behavior of the TOU system, credit was given for correctly determining weekly total energy use and the simulation only altered the timing of that use during the week. This led to a calculation of little or no monetary savings from RTP compared to TOU, which probably understates the savings. In addition, several potential improvements to the scheduling algorithm were identified that would probably lead to savings. In the experiment, the control schedule was based on the prediction of the next day's load. As a result, overheating was avoided and customer comfort was increased.

From the perspective of utility operations, it was demonstrated that within the routine operations of a New York State utility, the information required to provide both hourly and daily forecasts of the marginal cost of operations is available and can be provided on a routine basis. For this experiment, the daily and hourly price information was generated and recorded in real time and was sent by facsimile one day late to the central computer. The logical next step would be to use a stand-alone computer system in the NYSEG control room, into which the forecast and actual costs could be input. This computer could then call either a central computer or a set of distributed control systems at customer sites. Moving from a delayed RTP to an actual hourly real-time price would then be a logical and straightforward step. The next step, integrating the cost forecasting and price calculations into an existing or future control system computer in the control room, would require more effort, but would also be feasible.

The use of hourly prices rather than day-ahead prices provided significant improvement in the benefits of RTP to the utility and the customer. These benefits resulted in part from the operating characteristics of NYSEG and in part from better matching of customer price and utility cost. Because NYSEG is a net seller into the NYPP and to other regional utility customers, its opportunity cost is not a function of its generating cost but rather of its (and the NYPP's) resource cost. This cost is closely approximated by the NYPP-based NYSEG lambda, as in the experiment (see Section 2.) This cost varies significantly between hours and can only roughly be forecast a day ahead. In the NYSEG case, forecasting the price one hour ahead is a significant improvement over forecasting a day ahead. As was seen in the more detailed discussion of the monetary benefits of RTP, the more accurate the price forecast relative to the actual, the larger the combined benefits of RTP to the utility and customers.

The second overall result of the experiment was to demonstrate the feasibility of automatic control of customer storage loads using RTP signals. The experiment demonstrated the feasibility of automatically retrieving price and weather information, implementing for each site a simplified building demand-forecasting model, developing a near-optimal charging schedule, and then modifying that schedule to reflect hourly prices. The experiments at both water storage systems successfully demonstrated each of these principles and features. Because of the inherently different storage and discharge characteristics of the earth-storage system, significant cost savings were not realized in the Plattsburgh experiment, but the RTP-based control system did significantly improve customer comfort.

The third result of this experiment was in terms of customer satisfaction with the operation of RTP. As hypothesized before the experiment, those customers using the water storage system were indifferent to RTP-based control in that their comfort level was unchanged; only their bill was reduced. The RTP-based control logic provided sufficient energy to the storage tanks to guarantee adequate supplies to the buildings at all times. In the case of the earth-storage system, customer comfort improved during the experiment[14].

During the experiment, many additional potential enhancements were identified and some were implemented in the second season. Some unscheduled charging was due to inaccurate load forecasts and imprecise sensor readings. Further enhancements should include:

- Better load forecasts, including:
    - Adaptive total load prediction; and
    - Adaptive estimation of state equations.
- Better measurements, including:
    - Adaptive correction of sensor measurement values based on engineering and statistical principles; and

[14] This improvement in customer comfort is largely a function of the use of day-ahead weather forecasts to schedule charging. To achieve this improvement at the site would require an automatic link to a weather service and the logic to forecast building load requirements based on that forecast; i.e., essentially the same hardware and software required for an RTP rate.

- Installation of multiple and more precise sensors.

- Better price forecasts, including:

  - Time series forecasting methods to supplement NYSEG calculations;

  - Price adjustments, based on a statistical model, several times per day; and

  - Enhanced price extension formulas.

- More research on parameters used in Adjusting Controller, including:

  - Temperature dead bands;

  - Time horizon length; and

  - Frequency of recalculations.

- Event-driven scheduling of recalculations.

It is not possible to estimate the impact of these changes on savings without further experimentation or simulation. However, based on the weekly results, the range of likely savings in the next generation of this experiment should be a few percentage points higher than previous results.

CONCLUSIONS

The RTP experiment carried out during the winter seasons of 1989-90 and 1990-91 at three sites in NYSEG service territory was a technical and economic success:

- It demonstrated the feasibility of developing and communicating hourly cost/price information between the utility and customer sites;

- It demonstrated the feasibility of automatically monitoring and controlling charging energy storage systems based on response to RTP; and

- It demonstrated that customer comfort could be maintained or improved with a reduction in total system cost.

The economic results demonstrated the potential for net savings for ETS customers with electric bills from $4,000 to $5,000 or more per heating season.

APPENDIX A

WEEKLY GRAPHS OF NYSEG ACTUAL HOURLY PRICES

Appendix A contains the weekly plots of NYSEG actual hourly prices for March 1990. The graphs show that prices can vary substantially from hour to hour and day to day.

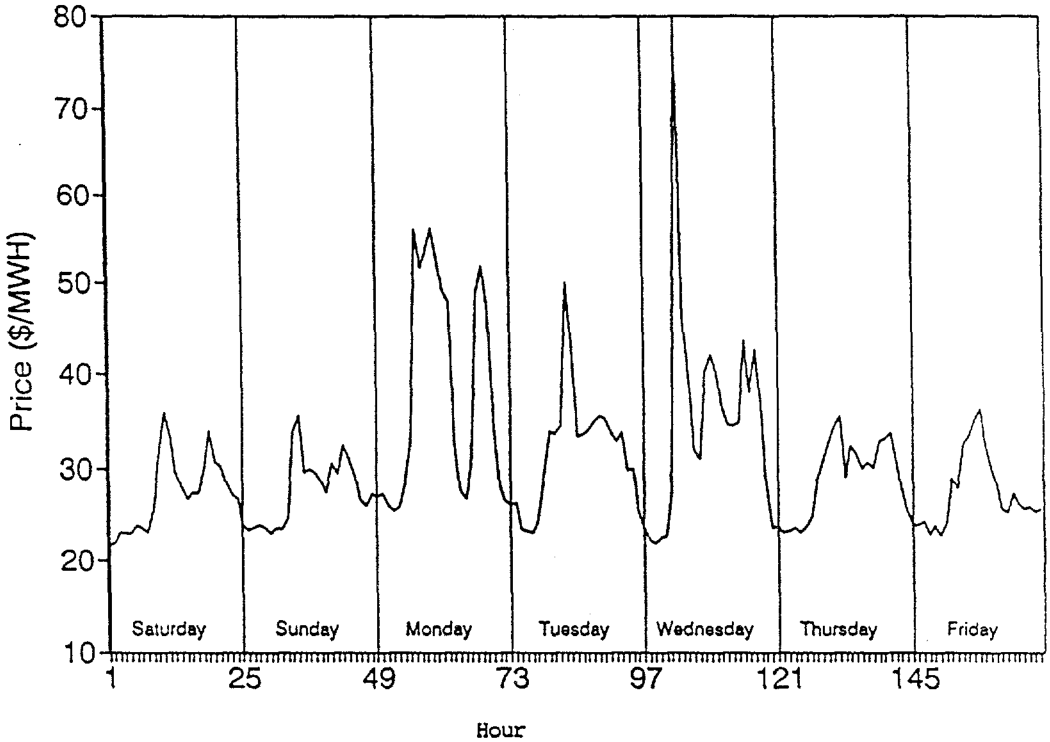

Figure A-1. Actual Hourly Prices for February 24, 1990, to March 2, 1990.

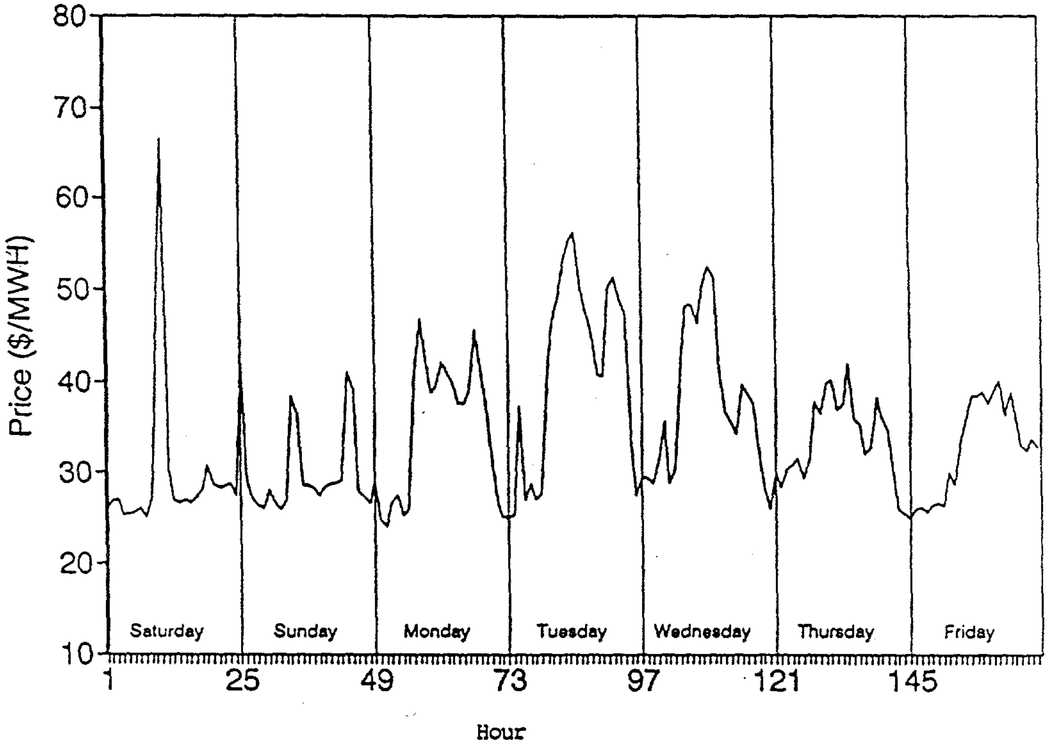

Figure A-2. Actual Hourly Prices for March 3, 1990, to March 9, 1990.

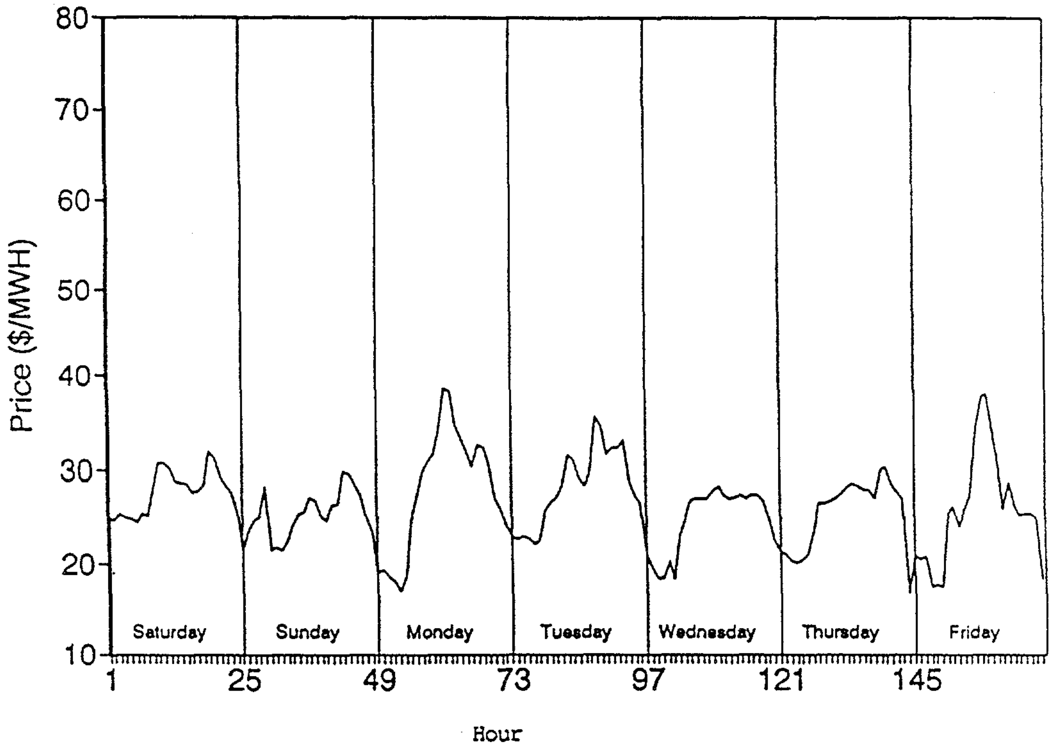

Figure A-3. Actual Hourly Prices for March 10, 1990, to March 16, 1990.

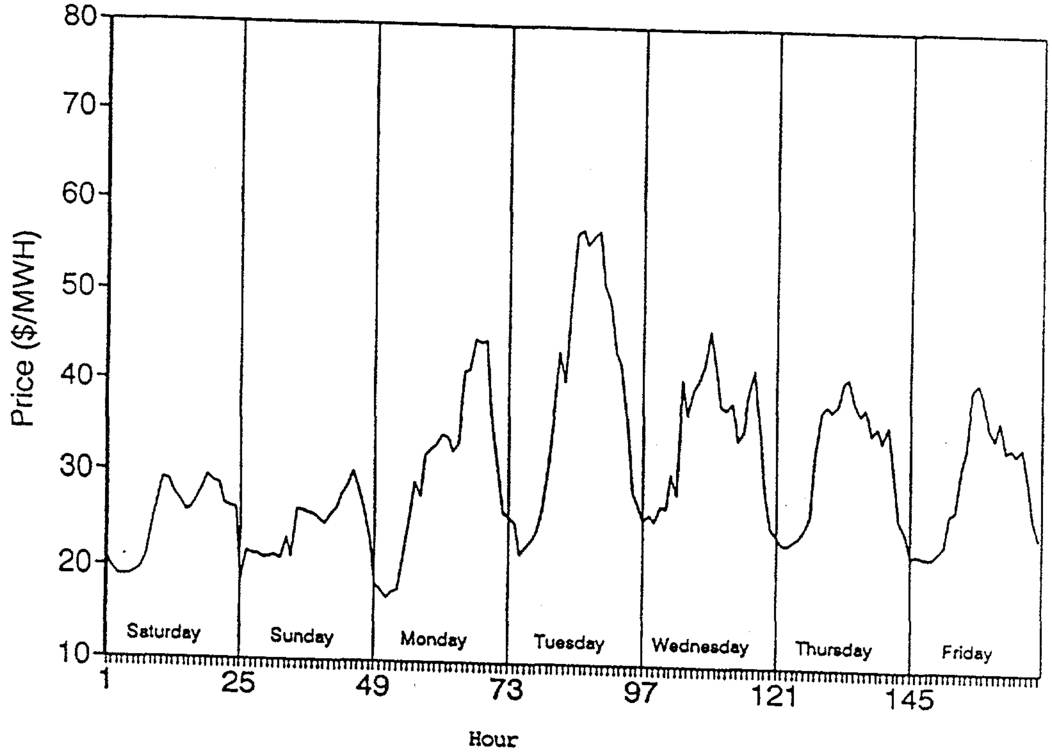

Figure A-4. Actual Hourly Prices for March 17, 1990, to March 23, 1990.

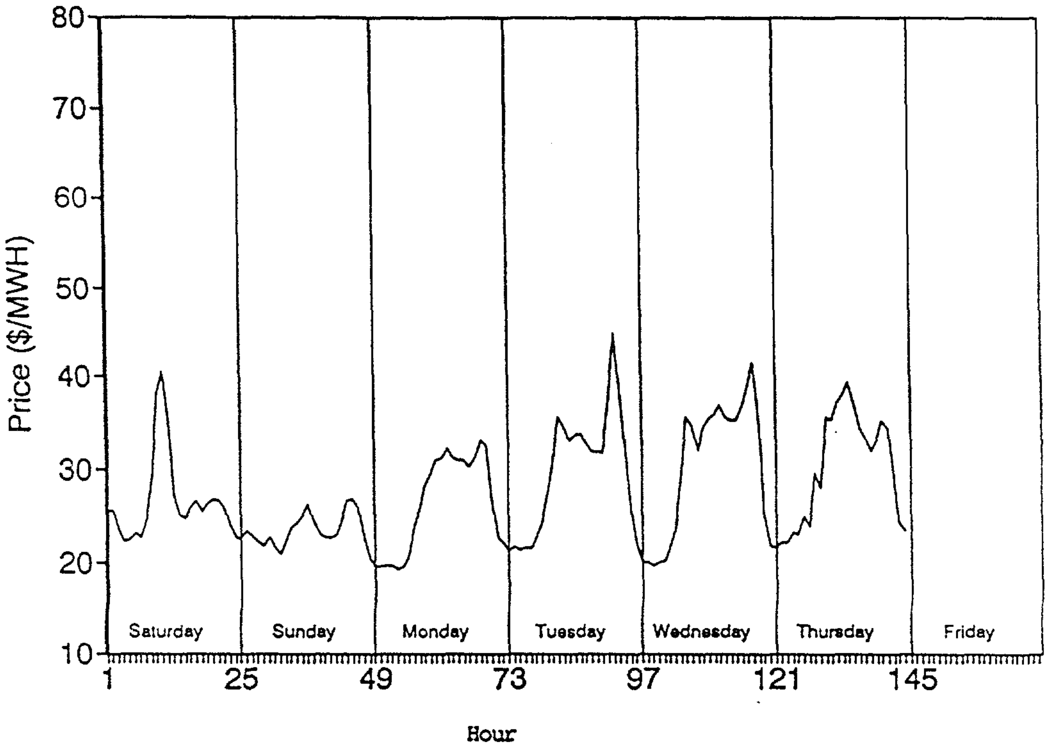

Figure A-5. Actual Hourly Prices for March 24, 1990, to March 30, 1990.

APPENDIX B

PRICE FORECAST ACCURACY

Appendix B contains the monthly plots of hour-ahead and day-ahead price forecasts versus NYSEG unadjusted actual hourly prices for March 1990. The 45-degree line represents a perfect forecast; deviations above and below the line indicate the relative forecasting error.

The plots show the variation in the accuracy of the forecasts. During days when prices varied widely, the day-ahead forecasts were generally not very accurate. The hour-ahead forecasts intermittently captured the price variations of such days, but were less successful in doing so when the variations were extreme. The inaccuracy of the hour-ahead forecasts, especially during off-peak hours, caused some suboptimal charging.

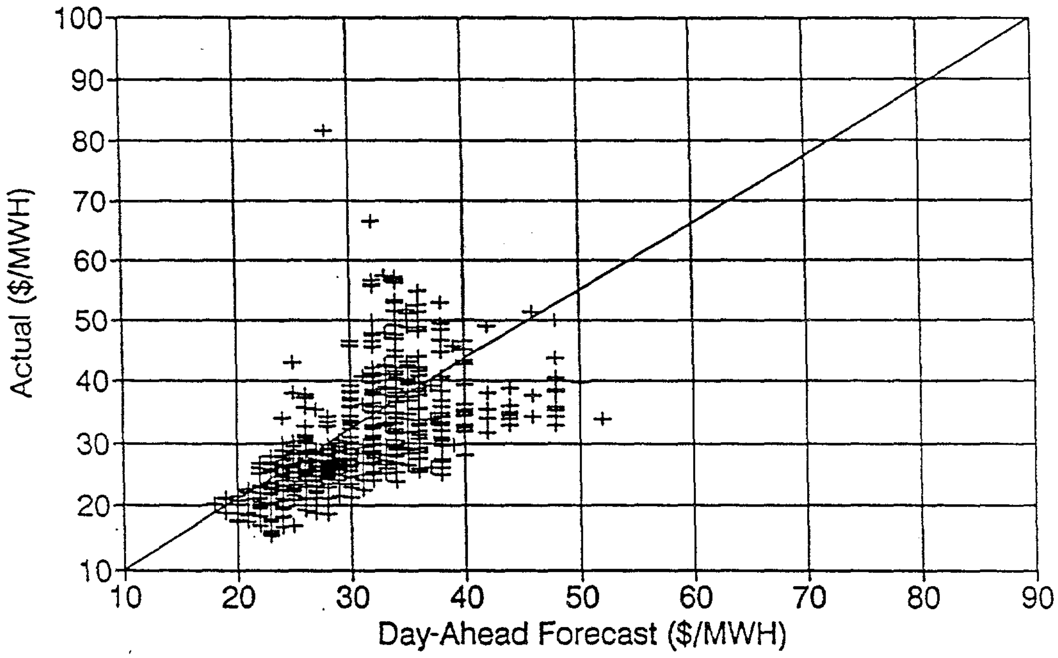

Figure B-1. Comparison of Actual Hourly Prices with Day-Ahead Prices for March 1990.

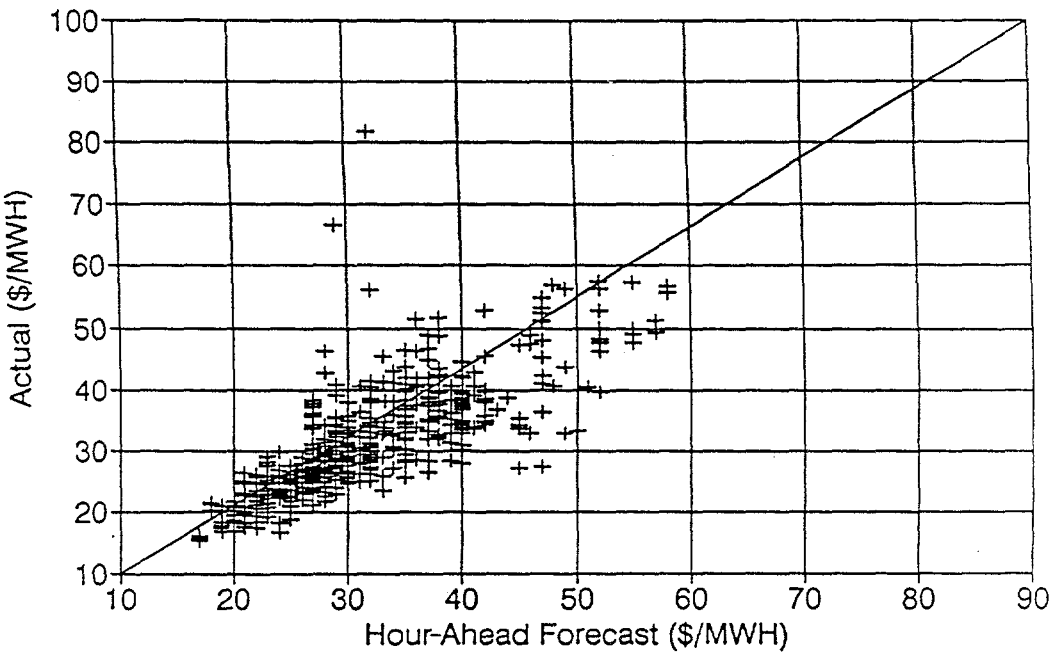

Figure B-2. Comparison of Actual Hourly Prices with Hour-Ahead Prices for March 1990.