TAKING ADVANTAGE OF REAL-TIME PRICING TO OPTIMIZE PLANT OPERATIONS

CHRISTOPHER J. RUSSO, RICHARD D. TABORS & PETER L. COOPER
TABORS CARAMANIS & ASSOCIATES / MIT
CAMBRIDGE, MA

CRUSSO@ALUM.MIT.EDU
TABORS@TCA-US.COM
PLCOOPER@MIT.EDU

Keywords: real-time pricing, plant optimization, energy storage

ABSTRACT

The Massachusetts Institute of Technology recently conducted a study to determine the potential benefits of becoming a wholesale market participant in New England and operating in a real-time (hourly) spot electric energy market. The study, utilizing historical hourly New England market prices, analyzed the MIT system based on the availability of state-of-the-art information technology. The analysis showed that operating in a competitive wholesale electric market would result in savings and revenues ranging from $150,000 to $600,000 per year. MIT could achieve these results simply by making a number of relatively minor operational modifications such as the testing of back-up generators at times of high spot-market prices as opposed to constant, prescheduled times. In addition to detailing the MIT study, we review some necessary components of response to spot price electric markets as they apply to large institutional consumers. Further, we provide insights into additional system optimization strategies for large institutional consumers that can be applied to either time-of-use or flat rate structures.

The MIT Cogeneration Plant

MIT's cogeneration plant was conceived in the early 1980s in response to the energy crisis of the previous decade. The decision to undertake construction was driven by MIT's growing energy needs and the need for an economically stable energy supply. MIT required a stable and disruption-free electrical and thermal power supply to support its extensive research organization; past history had revealed inherent inadequacies in the local electric utility distribution system. To address these concerns, a redesign of the utility interconnection and on campus electric system was a part of the plant development. In addition, the new state-of-the-art cogeneration plant resulted in an environmental benefit of a 45% reduction of emissions.

The plant began operation in September of 1995 and represented collaboration between the administrative and academic at the Institute. The existing boiler and chiller plants were repowered with a gas turbine from Alstom Power[1] of Finspång, Sweden (with a

[1] Formerly ABB STAL

combustor designed, in part, by the MIT Combustion Research Facility) and an advanced heat recovery boiler, the combination of which led to an overall thermal efficiency for the facility in excess of 80%. The facility is monitored through a state-of-the-art IT system (PI) that allows MIT to continuously track and store data from real-time operations of the main cogeneration plant, a satellite campus chilled water plant and the high-voltage campus distribution system.[2]

MIT's plant operates in a load-following mode, but because of MIT's higher electrical demand and the reduced plant output during hot weather, it still must purchase some supplementary power during periods of peak usage. MIT's current pricing structure is a relatively standard time-of-use, or TOU, rate. MIT wished to investigate the possibility of becoming a market participant in New England, and thus decided to analyze what potential savings and increased revenue might be achieved under a variable pricing system. Working in conjunction with Tabors Caramanis & Associates,[3] MIT examined historical New England price information and its own operational data over the past two and one half years. These data were then compared with data resulting from alternative scenarios of proposed plant operational modifications that MIT could make to capitalize on some advantages of the new pricing structure.

Traditional energy savings strategies have focused on a simple reduction of energy use and/or an increase in the ability to self-generate. To some extent, this was an artifact of the pricing structures available to large industrial customers; typically, a flat or a simple time of use rate was used to calculate energy costs and thus favored "across-the-board" strategies. Demand reduction/peak shaving strategies were more challenging due to long (8-12 hour) peak demand periods in the tariff and contract ratchets sometimes lasting as long as eleven months. With the advent of increasingly sophisticated systems that can calculate electricity prices temporally as well as spatially, significant opportunities exist for sites to monitor this information and plan their operations accordingly, and in almost all cases, decrease their overall costs.

What we attempted to do at MIT was to apply some minor operational changes and see their result on the cost of supplying the campus with energy. While the tactics we propose are relatively simple, there are increasingly sophisticated methods that can be employed to take advantage of real-time pricing. One of the authors, in fact, has performed work on various computer programs that can, with some process knowledge, availability of energy storage and a crude pricing forecast, plan plant operations such that total energy cost can be significantly lowered and asset usage maximized.

---

[2] For further details of this system, please see Using Real-Time and Historical Data over the Internet, Christopher J. Russo, 2000 IDEA Annual Conference Proceedings, http://cogen.mit.edu/unified/

[3] Mr. Russo and Dr. Tabors both have academic appointments at MIT as well as professional association with Tabors Caramanis & Associates.

A Review of Electricity Markets

Market Mechanics

Historically, the pricing of electrical energy has been regulated and sold at a fixed price to retail and large consumers. This system has been replaced by markets where there are multiple market participants on both the demand and supply side and where the price of electricity can change on an hourly (or even more frequent) basis in response to consumer demand and the marginal cost of production. Notwithstanding recent developments in such markets as California, the objective is to, through competition on the generation side, reduce the cost of supply whilst simultaneously inducing price-responsiveness on the demand side.

The market for much of the electrical power in several markets, including New England and California, is run by a central authority that uses a bid system to determine the energy price on an hourly basis. In this "day-ahead" system, generation owners enter a bid for their resources for the next day, specifying how much power they are willing to supply for what price. The market operator, commonly known as a power exchange, forecasts the amount needed for the following day and decides which units will be dispatched, or called upon to generate, for each hour. All suppliers then receive the price bid by the last unit dispatched, known as the *market-clearing price.* This price can vary significantly from hour to hour, and, as recent events in California have demonstrated, is not entirely immune to manipulation by participants wielding market power. Participants, however, do not know the market-clearing price until after the time interval in question has passed, thus this mechanism is known as *ex-post* pricing.[4]

We wish to make clear a semantic distinction here; we refer to the day-ahead hourly market as the spot market, but in reality, a "true" spot market for power exists with time horizons as short, in some cases, as 10 minutes. Often, much of the activity in this market is the result of short-term purchases by the system operator (ISO) for system stability purposes. This paper does not consider this market, but rather the aforementioned hourly market. Hereafter, the use of the term RTP will refer to a day-ahead hourly spot market for electrical energy.

Historical Market Conditions in New England

Prices in the New England market are calculated on an hourly basis, but because of market regulations, only traders and utilities are impacted by these prices; the typical retail consumer (including MIT) does not yet see them. This may change in the near future, however, as new rules allow large retail consumers to more fully participate in the market. One beneficial effect of California's power problems may be to accelerate development of real-time price signaling to consumers, which even politicians now realize is a necessary stabilizing element of a market.

---

[4] Separate markets exist, notably in PJM and in California, for the ability to transmit power from one region or point to another. This paper does not consider that market, but for a thorough treatment of the subject, please see Spot Pricing of Electricity, Bohn, Caramanis, Tabors, Schweppe, Kluwer Academic Press, 1988

Figure 1, below, shows the hourly prices per MWh from 1999 until 2000. A $1,000/MW price cap has been imposed on the New England market, thus perhaps artificially mitigating the intensity (if not the duration) of the price spikes seen in the graph below. The data show that the New England market has, to date, been for the most part relatively stable, having experienced few of the gyrations characteristic of California in recent months. One sees a typical diurnal variation in prices as well as a modest seasonal variation.[5]

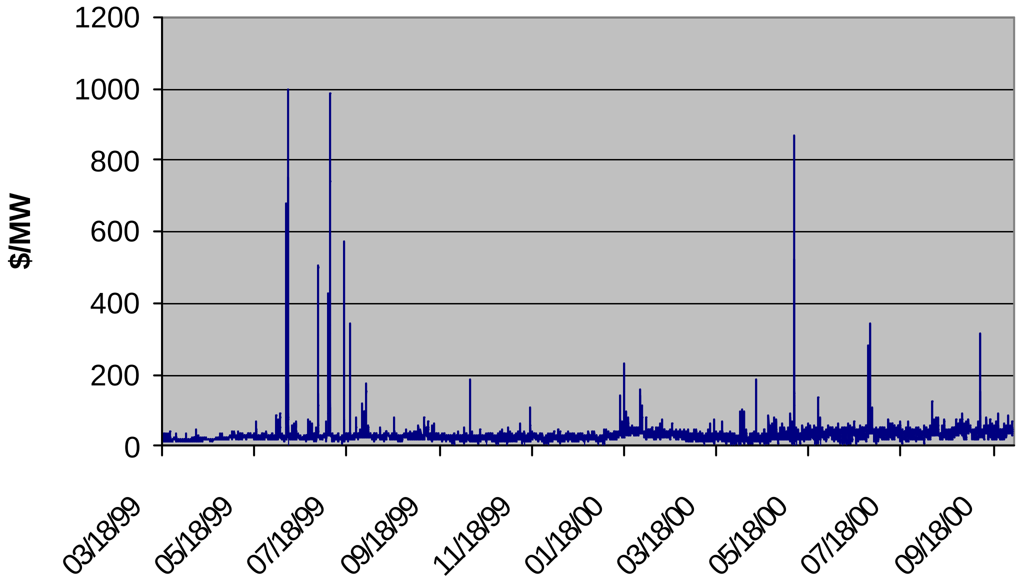

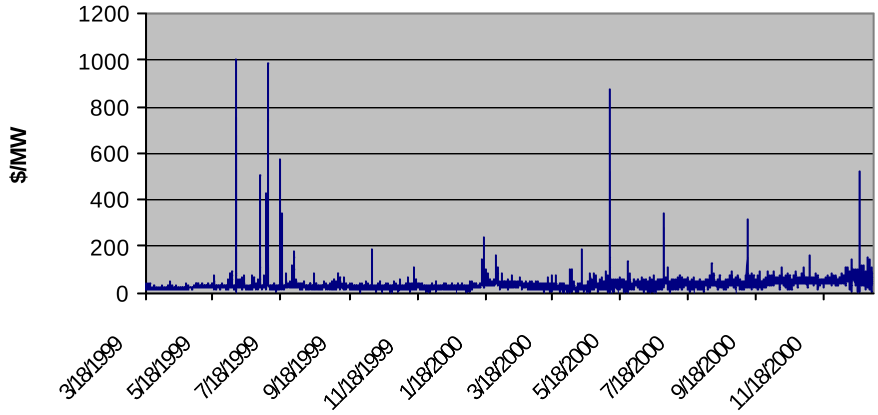

[5] One interesting facet of the New England (and other) market is the fact that the price may, at times, become negative for purchased power. The cause of this seemingly impossible economic situation is the fact that at times there may be too much capacity available in the form of large hydro and nuclear plants that must be used lest it be dumped.

Figure 1 - ISO New England Hourly Prices 1998-2000

Price spikes can and do occur, however, and when they do, prices can be more than an order of magnitude above normal levels, owing to outages or simply unexpectedly hot weather.[6] The typical duration of price spikes tends to be on the order of two to six hours, making short-term response to prices a significant component of total energy expenditures.

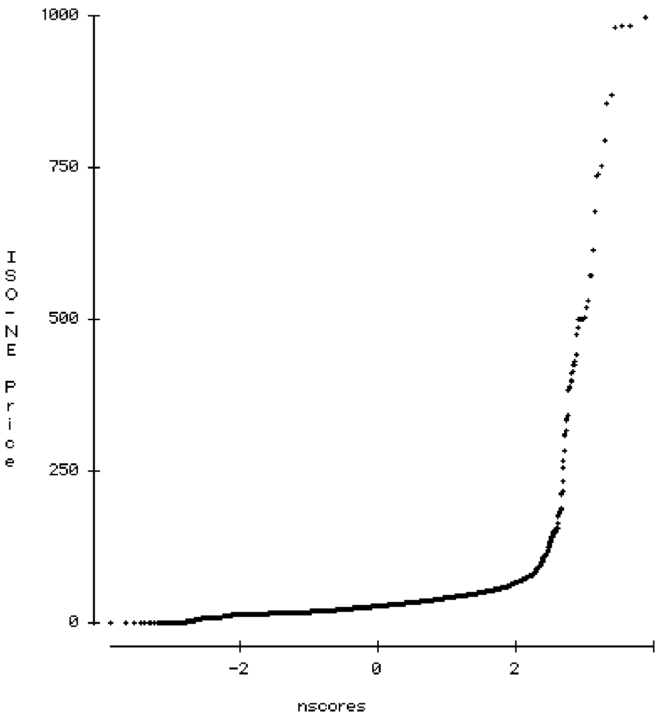

Figure 2 - Normal Probability Plot for ISO New England Hourly Prices 1998-2000

Figure 2 shows a normal probability plot for the hourly market prices during the time period studied. Examining the graph, one sees that when elevated prices occur, they are indeed "spikes" rather than extended periods of moderately higher prices.[7]

Indeed, a greater range in price only enhances the utility of scheduling and optimization to real-time pricing. In the optimal case, one would encounter a highly predictable and very different price throughout the diurnal cycle. Put another way, optimal RTP response can be thought of as "buy low, sell high"; the greater the price difference for

---

[6] Indeed, the highest prices of the year 2000 occurred on May 24th, the reason being that New England experienced anomalously high temperatures concurrent with many plants' schedule maintenance outages.

[7] A normal probability plot has been used here to illustrate the severity of the price spikes instead of a histogram or cumulative frequency chart because the level and duration of the spikes could not be well represented on a common scale with the remainder of the data. A normal probability plot, if the data have a normal (i.e. gaussian) distribution, should be a straight line. The sharp uptick in the graph above indicates that there are some data points that deviate very significantly from the remainder of the data

each unit of energy that one can switch from one time period to another, the greater the financial benefit.

Necessary Elements for a Spot-Market Energy Strategy

To fully participate in a real-time wholesale market for electricity, some necessary components are required. The first, of course, is the existence of a real-time market for electricity, which may not yet exist in many locations. The second element is the capability to capture and store the large amount of data necessary to predict with some measure of accuracy loads and prices. There are, however, some additional pieces of the puzzle that must be assembled before becoming a market participant.

One necessary component is some form of energy storage. The use of storage allows users to, put simply, shift energy and electricity usage from high-priced hours to low-priced hours. In addition to the obvious effect on energy costs, intelligent use of storage mechanisms can yield additional benefits in the form of reduced capital costs. Separate research by one of the authors has shown that with careful planning, installation of an energy storage system can reduce capital expenditures by approximately 23% by reducing the need for peak capacity. Often, in systems with a cyclic demand pattern (typical of colleges and universities), the marginal cost for the "last megawatt" (or ton, or pound of steam) of capacity often is considerably higher than the average cost per megawatt. With the installation of a smaller baseline capacity and off-peak storage systems, the generating equipment can be more fully utilized and peak capacity increased.

An obvious application already in widespread use among district energy systems is thermal energy storage in the form of ice. While this is not precisely electricity storage *per se*, it may be used to offset electricity usage during peak hours. If used in conjunction with electric driven chilling equipment, the effect is even more pronounced and apparent.

Another necessary component of plant operation in response to real-time pricing is some sort of forecasting capability. Price forecasting is notoriously difficult to do,[8] but because of the manner in which real-time prices are calculated under a central exchange system, common to many markets today, it is quite possible to determine the relative value of a price, but not necessarily its exact value. Put another way, it is difficult to forecast what the price will be, but not necessarily as difficult to determine whether the price during a given hour will be high, low or somewhere in the middle.

Some manner of price forecast, be it coarse or fine, is a necessary component of any type of optimization to real-time pricing. While it is conceivable that a site could react in "true" real-time, the infrastructure to allow this type of information transmission as well as the ability on the consumer side to react to it is not yet sufficiently developed.

In addition to sending forecasting and energy storage, an effective method of reacting to real-time prices is to switch energy sources for self-generation depending on real-time

---

[8] The authors are reminded of a quip attributed, perhaps apocryphally, to Yogi Berra: "Predictions are difficult to make – especially about the future."

prices. One example that has brought the issue into sharp focus recently is the spike in natural gas prices. If one were able to forecast the prices for each fuel, or more importantly, the relative price of electricity to gas (or oil, for that matter), one could choose to change the level of self generation or to switch fuels. One example is the possibility of installing a gas/electric chiller that can switch "on the fly" from one source of power to another. In the event that the cost of one hour of operation on natural gas exceeds that of electricity, then electricity would be used to fuel the machine or vice versa.

While the prices of electricity and gas have historically been highly correlated, there is often a difference to be exploited. Contrapuntally, while economic concerns alone may provide sufficient impetus to consider switching fuels, the power woes of California may provide yet more; not only are prices high, but there also exists a very real possibility of curtailment and interruption. Many large consumers are on interruptible contracts with their supplier, meaning that they are sold power at a lower price in return for agreeing to be interrupted in the event of a power emergency. Typically, these contracts have specified several days worth of hours per year to be curtailed, and, recently, the number of actual interruptible hours for each consumer was very close to zero, leading to the expectation that interruptions were a possibility only in theory.

With the incidence of Stage Two and Stage Three power emergencies increasing alarmingly, many customers have chosen to remain online during these interruptions, paying penalty rates of up to $10,000 per MW, instead of interrupting their operations. Were a customer better able to predict when an interruption event might occur,[9] they could make provisions to use an alternate fuel source and thus avoid interruptions in service or large bills.[10]

RESULTS FROM APPLICATION TO MIT

Our intent at MIT was to develop minor operational modifications that, with additional though given to their timing and method, would yield large energy savings for the Institute. To allow us to more easily analyze the possible scenarios, we selected three methods of real-time pricing response that could be implemented relatively easily. We analyzed scenarios in which MIT rescheduled the weekly testing of its emergency generator, sold excess power into the New England market, and reduced consumption very slightly during periods of high prices. All three have the advantage of being easily monitored and centrally controlled and entail very minor operational changes – usually an operator in the control room could make the changes literally single-handedly.

---

[9] It is necessary to concede that while our methods can predict whether an emergency may be declared, it is decidedly, if not infinitely, more difficult to predict whether a specific customer might be interrupted, although one reasons that the former surely foments the latter.

[10] Since the submission of this paper for initial review, the hitherto incredible possibility of involuntary curtailments in California has, in fact, been realized.

Rescheduling of Emergency Generator Testing

Like most large industrial consumers and self-generators, MIT has an on-site emergency generator (rated at 2,000 kW) that must be periodically tested to ensure its proper function. Also like most self-generators, stringent environmental regulations prohibit the use of the generator for peak shaving and limit the number of hours of testing per year.[11] While MIT has traditionally tested its generator during off-peak hours, we decided to determine the potential savings were it to test its unit during the optimal price hours each week. To do this, we assumed total past prescience; that is, we assumed that, at the beginning of any given week, MIT would know in advance which two hours during the week would have the highest hourly price, and that MIT would then schedule its generator to be tested during those two hours.

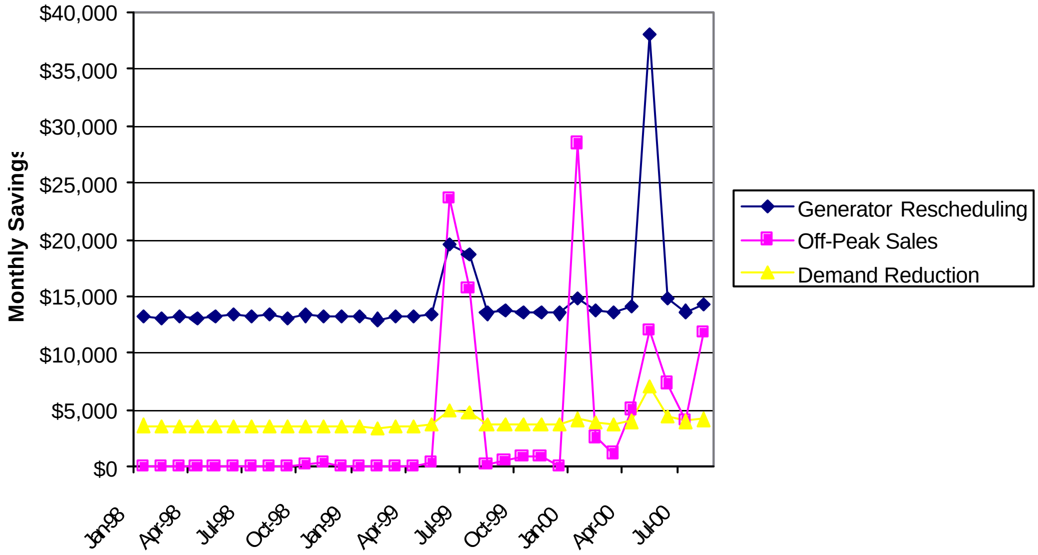

Figure 3 - Monthly Savings from RTP Optimization

As discussed previously, this claim of prescience is not as far-fetched as it might appear; determining *when* a peak might occur is not as difficult as determining *what* the peak would be. The results of our analysis show that had MIT tested its generator during the optimal hours over the last two years of operation and had real-time pricing been fully implemented, the savings would have been on the order of $150,000 per year. Even if one assumes that we would have been wrong in predicting when exactly the peaks would have occurred (a not unreasonable assumption given the notoriously volatile nature of electricity prices), the savings are still significant. This change in operations typifies the type of change we are trying to engender: minor operational changes that result from operating the plant more intelligently, not necessarily just increasing production.

[11] One large corporation in the northeast was recently sanctioned for, in fact, "testing" its emergency generators for several hours per day during peak days.

Off-Peak Power Sales

The second operational change that we examined was the possibility of off-peak sales of power on the New England market. To analyze this scenario, we again used historical prices from ISO-NE and assumed that we had run the MIT plant at the maximum possible electrical output as determined by ambient conditions and then sold all available excess power at market prices. We also estimated the plant efficiency at those points and used historical fuel prices to determine our cost of generation.

From 1998 through the spring of 1999 market prices in New England were most likely depressed due to low prices for natural gas combined with an abundant supply of hydroelectric power in the market. Consequently, we found that it would have been difficult for MIT to have had made a profit during this period. The price for power was simply too low to justify the additional cost of generation. .

Beginning in the spring of 1999, however, New England saw an upswing in prices. The principal components of this increase were not the off-peak prices, but rather peak prices and temporary price spikes in the market. From this time onward, MIT's potential savings would have increased dramatically as prices in the New England market began to experience spikes. When we examined where MIT realized the majority of its savings, we found that there was a relatively constant "baseline" level of revenue from off-peak sales under normal operation, but during periods of elevated prices, revenue increased by at least an order of magnitude in a given hour.

The overall savings to MIT during the period from 1998 through 2000 was approximately $827,000, but this amount is not entirely evenly distributed; savings ranged from $3,000 per month during some winter months to more than $30,000 per month during warmer months when price spikes occurred.

Demand-Side Peak Reduction

The third possibility for real-time pricing response studied for MIT was the possibility of demand reduction during periods of high prices, commonly referred to as peak shaving.[12] In this case, we decided to model the results if MIT had been able to predict the highest priced hour of each day during the summer months (May through September inclusive) and then reduce load by 500 kW during that one hour per day.

Performing that analysis yielded relatively constant savings of approximately $3,800 per month. (Shown again in Figure 3) Somewhat surprisingly, this number was considerably lower than the savings that MIT achieved using the other methods, but represents the "low-hanging fruit" of operational changes that MIT and other such institutions might implement.

For MIT, controlling building consumption is more challenging than the aforementioned strategies, which could be implemented at the central plant itself. Building management systems range from current technology to 20+ year old systems; making integrated data

[12] In this case, the peak shaving refers to load reduction, not necessarily to increased power output. This distinction is necessary because MIT is indeed a self-generator, but it does not have enough generation capacity to satisfy its entire load all of the time.

collection and control system more difficult to implement. Also, changing space temperatures is not generally supported by the members of the MIT community, which may experience some mild discomfort but may only see the savings indirectly.

Application to TOU and flat rates

The optimization techniques and operational modifications presented in this paper apply primarily to those sites that are able to buy and/or sell power in a market with variable prices. Even many large consumers and self-generators, however, are still unable to take advantage of these rates, and many of them are using so-called time-of-use, or TOU rates. These rates coarsely model the cost of electricity by assigning a different price to the electricity purchased (or sold) depending upon the time of day. For example, many TOU rates have an off-peak block and one or more peak blocks where power is more expensive.

Operational modifications to adapt to a TOU rate are indeed possible, but more difficult because the differences between the rates are likely not as extreme as those of true market prices, and the length of these time blocks is often considerably longer than an hour, requiring operational changes that may be more pronounced. Furthermore, demand charges are calculated on a month-long basis and may involve an 11-month ratchet. This requires month-long or year-long performance to capture and retain demand cost reduction benefit.

Concluding thoughts

From a campus operational perspective, the interesting aspect of real-time pricing is that cost of power can be much more expensive for periods of much shorter duration as compared to the peak demand periods and structure of TOU tariffs. This makes almost any demand reduction scheme easier to justify as long as real-time pricing data is available. And for schemes that involve reduction of usage in buildings by voluntary action of occupants or by temperature changes via building management systems, there will be much higher rewards over shorter periods to motivate voluntary cooperation or acceptance of inconvenience or discomfort.

One aspect of adaptation to a real-time market that was not examined in this study but may have significant effects is the participation in a real-time market for ancillary services. Ancillary services are those services that are necessary for operation of the power grid but are not real power, that is, megawatts, in an electrical sense. Rather, they are those services such as spinning reserves and regulation and that are, in some markets, being treated as commodities. While not studied here, participation in these markets could potentially also yield large benefits for those considering moving to a real-time rate.

Overall, our findings show that there is a potential for dramatic savings on energy costs and increase in energy sales revenues for those customers who may choose to participate in a real-time rate. Far from introducing a measure of unpredictability and volatility into campus operations, real-time rates provide significant opportunities for those equipped to turn them to their advantage. While a customer must make efforts to closely monitor and analyze information, we have found that even modest changes in operations can yield savings that more than justify the cost.

Acknowledgements

The authors wish to thank Joshua Bartholomew, Meredith Miller, and Renee Rushnawitz of TCA, Roger Moore of MIT, David Rogers of Rogers & Associates, and Bernardo Fernandes of Maxpower for their assistance and insight to this paper.