MIT EL 01-014 WP

Energy Laboratory

Massachusetts Institute
of Technology

Understanding Demand: The Missing Link in Efficient Electricity Markets

September 2001

Understanding Demand: The Missing Link in Efficient Electricity Markets

Contributors:
Dr. Marija Ilic, ilic@mit.edu
and
Graduate students:
Jason W. Black, jwblack@mit.edu
Elena Fumagalli, elenaf@mit.edu
Poonsaeng Visudhiphan, visudhip@mit.edu
Jill L. Watz, jillwatz@mit.edu

Energy Laboratory Publication #
MIT EL 01-014WP

Energy Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139-4307

September 2001

DRAFT

Understanding Demand: The Missing Link in Efficient Electricity Markets

Contributors:

Dr. Marija Ilic, ilic@mit.edu

and

Graduate students:

Jason W. Black, jwblack@mit.edu
Elena Fumagalli, elenaf@mit.edu
Poonsaeng Visudhiphan, visudhip@mit.edu
Jill L. Watz, jillwatz@mit.edu

Based on lectures and term project work in
Course 6.689—Advanced Power Systems Control
Massachusetts Institute of Technology
Spring 2001

September 13, 2001

ABSTRACT

Two basic criteria for markets to operate efficiently are that participants have sufficient information to make informed decisions and that there is competition. The recent deregulation of the electric utility industry has witnessed a dearth of competition in retail markets due to a lack of price signals being passed to consumers, thereby creating inefficiencies in market dynamics. These problems have occurred as a result of ill-defined market structures and policies that have focused on supply side issues while giving short-shrift to demand side effects. This paper illustrates that demand responsiveness is a necessary criterion for the efficient operation of electricity markets. The technical and market mechanisms necessary to implement responsive demand are explored. In particular, it proposes a role for load serving entities (LSE) as load aggregators and market participants and establishes the need for a control system that couples electric loads to market price signals. Technologies necessary for accommodating demand response such as, real-time metering and "smart" controls are described along with examples of their potential load shaving benefits.

KEYWORDS: Demand Management, Electricity Markets, Load Serving Entity, Load Control

1. THE PROBLEM OF UNRESPONSIVE DEMAND

Conditioned by the years of electric utility regulation, the public has come to expect reliable, high quality electricity on demand. In US society this high quality, relatively inexpensive power has been taken for granted. In proportion to income and budgets, the cost of electricity to most businesses and residential consumers has been fairly small, less than 1-5% of the average budget. Under the regulated utility structure, most customers have been metered monthly and billed a flat average rate for the total amount of electricity used. A peak rate for any energy consumption above a monthly average has been charged; however, the peak usage charge has not been correlated with actual power use at the time of system-wide peak loads. One clear deficiency with flat rate, average pricing is the lack of price signals to the customer. Without this critical piece of information, consumers have no incentive to change their consumption behavior in response to price variation during the daily system peak or to seasonal high wholesale prices or supply shortages: market inefficiencies or infrastructure failure. To date, price responsive demand has not been realized to any significant degree in the new markets. As a result, demand is inelastic, and when supply is limited, generators are able to easily exert market power to drive up prices.[1]

1.1 ELECTRICITY MARKET POLICY: PRESENT ISSUES AND CONCERNS, FUTURE SOLUTIONS FOR LSEs

The lack of demand response in the new markets has been exacerbated by a lack of retail competition, as only a limited number of Load Serving Entities (LSEs) entered the market. The markets have not observed many new entrants in the retail/distribution area due primarily to the inherent market advantages of incumbent utilities, the high costs of market entry for new participants, and deregulation policies including low standard offer and default service prices. Inadequate consumer education, the difficulty for the LSE to define consumer elasticity, and the absence of an adequate policy framework are other important factors.

Incumbent utilities have clear market advantage due to the "brand" name effect attributed to decades of being the sole electricity provider for a dedicated service area. Customers are often reluctant to change from well-known providers unless offered a significant incentive. In addition, most deregulation legislation appoints the incumbent utility as the default service provider. So, if a customer does not bother to specify an electric service provider, the service is automatically provided by the incumbent. The default service and "brand" name effect provide the incumbent with a market advantage because its costs for advertising and acquiring customers is low compared to that of a new entrant that must spend money to convince customers to switch providers.

In addition, the barriers to market entry for new participants are high due to the disaggregated available customer base (mostly residential and small commercial) and the low standard offer legislated by most states. Marketing costs to new entrants are very high because the available customer base is primarily comprised of the disaggregated residential and small commercial customers, since many large industrial and commercial customers have the ability to purchase directly from the wholesale market. The low standard offer set by state legislation requiring a

[1] This issue is under serious debate in the state of California.

fixed discount on every retail bill (typically 10%) for a fixed period of time further limited the ability of new entrants. The combination of the cost of marketing to customers and the low standard offer made it almost impossible to compete directly with the incumbent utilities. For example, evidence from the CA and PJM markets indicate retailers were spending between $100-$600 per customer in marketing, representing approximately 15-90% of an average residential customers bill in CA [14].

Retail competition has also stagnated as a result of inadequate consumer education regarding the effects of electricity restructuring and customer choice. The common public expectation has been that rates would go down with competition; however, small consumers simply do not see enough impact on their household or business budget to pay attention to saving 10% off their monthly bill. Given high marketing costs, low standard offers, and the incumbent advantage of the 'brand' name, new retailers cannot offer competitive prices or savings on the electricity bill that would be significant enough for consumers to consider switching providers. After two to three years in the deregulated environment, public utility commission data from CA and MA indicate that less than 1% of the residential and small commercial customers switched providers [3], [11]. The switching rate for larger commercial and industrial is higher at 1% and 3 to 9%, respectively, which corresponds to 2.5 % of the load for residential and from 3 to 15% of the load for commercial and industrial. PA has observed greater movement in customers switching from the incumbent utility because the state offered smaller customers financial incentives to switch providers. The portion of residential load that has switched is 13.6% and it is reported to be 24% each for commercial and industrial [12][2]. After two to three years of competition in these states, the low numbers still suggest that real competition has not been achieved.

Another limitation on the entry of LSEs and competition is the ability to bid demand into the wholesale market. Greater elasticity of demand enables a more robust market; however, defining customer demand elasticity requires detailed temporal knowledge of customer loads. As explained in more detail in this paper, the ability for LSEs, to be price responsive through real-time metering and pricing, technical controls, and appropriate contracts with their customers results in lower electricity prices for all customers. Again, the incumbent utilities have a clear market advantage because they have historic data on all customers in the regulated service territory. Gathering and synthesizing the data for individual customers is expected to have high transaction costs initially, but the market and consumers will benefit from the investment. Markets that require demand bidding by the LSEs and provide proper incentives to ensure that adequate information can be obtained by all market participants will encourage competition. For example, requiring real-time metering and pricing for all customers, phased in from large to small over time, will provide a more level playing field and encourage new market participants.

Electricity deregulation policy also needs to provide a flexible market structure that allows for a transition between the current monopolistic situation with the incumbent utility serving 90+% of the market to the near-term situation with a few large LSEs and a few smaller LSEs (where the large LSEs maybe in a position to exercise market power) and eventually the long-term situation

[2] In general, the levels of switching are expected to be higher in PA because their state restructuring legislation mandated that state funds be used to pay consumers to switch. However, the higher percentages must be viewed carefully, because at least half of the customers accounted for in the state published numbers are consumers who switched to the primary incumbent utility's new retail company.

where the market is comprised of many competing LSEs. Without proper incentives to foster competition and enable the transition, consumers may be worse off after the standard offer expires. For example, if the number of LSEs and thus, competition remains limited at the time the standard offer expires, consumers will have limited choice and prices will likely be higher.

1.2 THE NEED FOR CONTROL STRATEGIES WITH REAL-TIME METERING AND PRICES

In order for demand to be price responsive, a mechanism for providing price signals to the end-use customer, as well as, a physical control scheme for responding to the price signal is necessary. illustrates the interaction between aggregate power supply from the generators, aggregate demand by the LSE, and LSE/end-use consumer response to price signals. As shown, various LSEs aggregate individual load demand ($q(x_i^n, p_i)$) at set point x_i^n on behalf of their customer base, a mix of residential, commercial and industrial loads (for n loads). An aggregate demand curve ($Q^D(p)$) is derived by each LSE based on the individual customer demand curves and contract terms and then bid into the market. A total market demand curve is matched to the supply curve ($Q^S(p)$) and a market clearing price (P_i) and quantity are determined. If the price is high, the LSE and end-user have an incentive to reduce load in the next time period (x_{i+1}) to lower costs either through manual ($q(x_{i+1})$: quantity adjustment made by the consumers given the market clearing price) or automatic control of end-use equipment (y_i: control signal of the utility that adjust the load given the market clearing price) operating set points.[3] The degree of demand reduction and necessary load control will be based on the types of contracts set-up between the LSE and its customers and the power purchasing strategy of the LSE, i.e., the mixture of bilateral (forward) contracts and real-time purchasing.

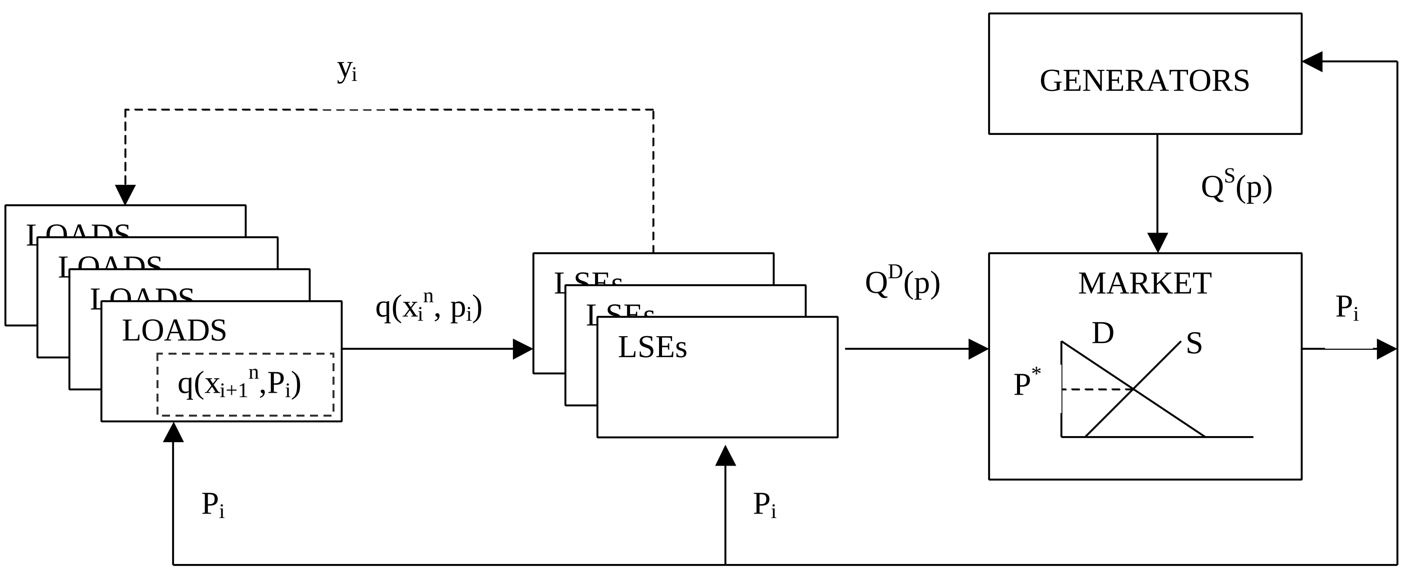

Figure 1 – Interaction between aggregate supply and aggregate demand

[3] The price signal could be an actual real-time signal when there is a multi-settlement day ahead and spot market or a day ahead price signal in a single settlement market.

The price signal provides incentives to the end-use customer to reduce load (i.e., turn the thermostat down on an air conditioning unit) at peak times of the day when prices are high. This can be done either independently (automatically or manually) by the consumer, and/or by the LSE controlling end-use customer loads through a centralized optimization scheme (with appropriate contracts). Currently available real-time metering technology can provide these price signals. In the past, this was been primarily used only by large consumers in the past due to cost and access limitations. Innovations in information technologies and power electronics, however, have driven down the cost of these metering technologies to the point where they are considered cost-effective for smaller consumers. Implementation of these technologies coupled with smart end-use equipment (such as, HVAC equipment, lighting, water heating, etc.) that can respond to price signals and shed load when prices are high, offers a control system to manage load and maintain more stable electricity market prices.

Under an appropriate market structure that allows for price responsive demand, LSEs have an incentive to enter the market to provide a variety of services. First by seeing the opportunities that derive from selling power to a price responsive demand, the LSE will have the incentive to provide consumers with the technology necessary to receive the price signals and control consumption accordingly. Secondly, given the diversity of loads and heterogeneity in consumer preference, the LSE will offer a variety of products and services that will enable it to bid competitively on the spot electricity market. For example, consumers that are more risk adverse and prefer a constant price to hedge against price volatility will opt for guaranteed rates at a higher average price, and more risk taking consumers will look for services that offer lower prices with a greater exposure to price volatility and the option to reduce load when prices are high. Finally, the LSE that has a good understanding of its consumers' demand elasticities (i.e., consumer willingness to be reduced at a certain price) can match loads to smooth out its overall load profile and determine the best hedging strategy between buying bi-lateral (forward) contracts for some portion of its load and purchasing the remaining on the spot or ancillary market. The bidding strategy of the LSE and more generally its purchasing strategy will have a strong influence on the market price. Strategic demand bidding and participation of the LSE in the bilateral market will, in fact, create an active and elastic, demand capable of driving prices down on behalf of consumers. A discussion of customer demand elasticity, the importance of load aggregation and LSE demand bidding strategies follow.

1.3 Demand-Side Economics: Consumer Demand Elasticity and Load Profiles

In order for an LSE to effectively manage load and strategically bid into the market, it is essential that it understand the way consumers respond to prices. Consumers will exhibit two types of price responsive behavior, *own-price elasticity* (electricity demand as a function of price), or a willingness to consume less electricity if prices are high and more if prices are low. Secondly, consumers have *cross-priced elasticity*, or a willingness to shift load from peak hours to off-peak hours in response to price, while keeping overall consumption the same. Own-price elasticity represents the consumers willingness to curtail or increase consumption as a function of higher or lower prices and cross-priced-elasticity corresponds to the consumers load profile and the ability to shift the peak consumption in response to price variation during a time period. Elasticity and cross-priced elasticity vary on a daily, weekly and seasonal basis.

When the LSE aggregates consumers, it will take into consideration both demand curves and load profiles in aggregating different consumer types. Consumer response, in fact, can derive from simply reducing consumption in response to price, by shifting consumption to different hours of the day, while deriving from the same overall service or output from the energy consumed, or by a combination of both. The ability of the LSE to be price responsive and its purchasing strategy will derive from an intelligent aggregation of consumer with different elasticity and cross elasticity.

Aggregate daily load profiles for various customer classes are probably well understood from years of power system management by utilities. Assuming this data can be gathered without a tremendous investment, LSEs will try to match customers with different demand peaks to smooth the aggregate load profile. For example, various industrial users that do not have continuous processes may be more flexible in load shifting even though their aggregate consumption stays the same. On the other hand, certain industries have continuous processes, particularly those of the "new economy" who have very high power loads that run processes or computer banks continuously 24 hours/day, 7 days per week. These consumers have less, if any, ability to reduce or shift load and would likely pay a premium for uninterruptible, quality power service (or opt to provide their own peak and/or back-up power if the economics are justified).

Because the LSE is a profit maximizer, it will attempt to balance its own exposure to price volatility in terms of its own power purchasing strategy and by developing differentiated services for its customers. An LSE power purchasing strategy will be a combination of forward fixed price contracts for some portion of its load (e.g., baseload) and spot market purchases. The LSE will also develop differentiated service contracts that offer customers choices of service rates based on load requirements and flexibility of reducing demand during peak periods. The optimal mix of power purchases and hence, profit maximization, will be based on the type of service contracts entered into with its customers and the ability for LSEs to be price responsive.

2. CONSUMER AGGREGATION STRATEGY FOR THE LSE

One of the roles of the LSE is to capture the value of intelligent load aggregation. The LSE creates a coalition of users (its customers) and purchases energy on their behalf. The LSE will purchase power to meet the *forecasted*, aggregate load demand of its customers. As we indicate in Figure 3 the LSE will always expect some fluctuations around the anticipated demand, within a certain range (load deviation band) because of short-term variations in load demand. Here we assume that the LSE has the responsibility to purchase power to provide for any deviations around the forecasted load profile.

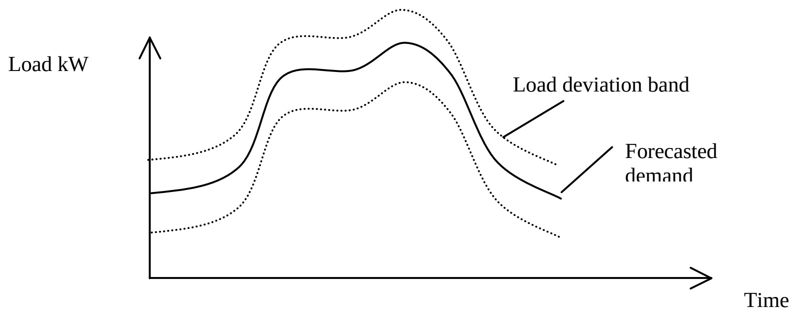

Figure 3 – Daily load forecasted demand and deviations

Since the power purchased by the LSE has to always match physical demand, contractual agreements need to be made for the different time horizons, for example: long-term bilateral contracts for forecasted demand, spot/forward market purchase for day-ahead deviations from the forecast, reserve market or frequency control market for near real time fluctuations. The contractual agreements and the purchasing strategy will depend, of course, on the structure and rules of the market[4].

Forming an adequate aggregation of electricity customers is important for different reasons. Even before we look into the value of creating a group of consumers that are responsive to price variations, there is value in bringing together end users with inversely correlated load profiles that will result in an aggregate flat purchasing profile for the LSE. In addition, there is value in choosing consumers in such a way that, in the aggregate, will reduce variations above or below the contracted purchasing profile.

Aggregate flat load profiles are beneficial to the LSE because they negate or minimize the need to purchase load following power services which are more expensive than base load power. This is very important when the LSE owns generating plants. When the LSE does not own generating assets and has to buy on the market it faces the problem that the value of not contributing to the peak load is not recognized today. In fact, contracts for electricity are based on the amount of energy purchased (in kWh) rather than the actual power profile (in kW over time). A contract based on the load profile conveys information about when the energy is needed and how much load following capacity a user requires the system to provide. A market rewarding loads or aggregations of loads that do not contribute to the system peak would create enormous value for LSEs able to intelligently aggregate load. In this case, the total cost of serving the aggregate load would be lower than the cost of serving the individual loads. The flat profile would enable the LSE to minimize total purchasing costs and, therefore to serve its consumers at a lower cost than the customer would otherwise receive on its own. Aggregating flat load profiles requires a variety of customers with heterogeneous and (some) inversely correlated load profiles. As we explain later, a combination of residential, commercial and industrial customers is necessary to smoothing out an aggregate load profile.

Even when the market design does not reward flat load profiles, an LSE purchasing long-term contracts for its aggregate load with a flat profile would not have to buy power on the spot

[4] Another important factor in the LSE purchase decision, in the mid-to-long-term, is the uncertainty about the number of its customers. Contracts between the LSE and its customers should have a sufficient duration to allow the LSE to plan in advance its aggregated power demand for the next period.

market to meet demand peaks. This reduces the exposure to price volatility for all the LSE's consumers[5]. There is therefore a value in the right aggregation of consumers that derives from reducing exposure to price risk. When buying directly from a power producer, a single load with a flat profile would be able to buy the capacity needed for a fixed price from one generation source whose marginal costs of production is low. A single load with wide variations in power needed at different time of the day or week will need to sign a contract with a power plant that has the ability to follow the load. Such generator will necessarily have higher marginal costs of production. Under another scenario, the flat load will have to make only one transaction when buying the necessary capacity, the load with variation in the load profile, instead, even if it can enter the same type of contract with a generator, will have to buy or resell on the market the surplus or deficit deriving from its actual consumption pattern. This last scenario implies the presence of a liquid enough spot market that will absorb any difference between the contract and the physical load consumption. The value to an LSE of aggregating a flat demand profile is therefore both in minimizing price risk and in better purchasing agreements.

Secondly, there is an extremely high value in aggregating consumers whose deviations bands around the forecasted profile are expected to cancel each other. In this case the value derives from minimizing costs for unanticipated spot purchases, and services for frequency deviations and Automatic Generation Control (AGC). Even though not every market provides for separate costs for these services, there is a real cost of providing AGC to the LSE. Therefore, the ability to minimize, through matching load deviations, the deployment of this service results in a lower cost power delivery and lower costs to the LSE. This effect might look negligible (depending on the market setup), but the strength of an LSE is in capturing cost savings on behalf of its customers.

The LSE can also create value by bringing together loads that have a peak during the system peak and to be able to move that consumption to other hours of the day or time of the year when demand and prices for electricity are lower. In other words, there is a high value in taking a load profile that is not flat to begin with and to make it smoother. The ability of the LSE to control loads according to price signals, is crucial in adding value to the aggregation. This is true when the LSE is a price taker and also when the LSE has a large impact on system wide demand. In the first case, the LSE's customers will benefit from shifting their consumption to cheaper hours of the day. In the second case, the aggregate load of the LSE will significantly modify the system wide load profile, directly affecting the market-clearing price. This case is more complicated in the sense that the price signal from the market to the consumers will affect the consumers behavior and that, in return, will feedback to the system demand modifying the market price itself. As a consequence the load profile of the consumers will adapt accordingly in an iterative process. As a consequence, two approaches to managing demand need to be assessed. In the first, the LSE aggregate demand will represent a negligible portion of the overall system demand. In the second, the number of aggregators in the market will be small enough for each of the aggregators to influence the system demand and consequently price variations. Note that, in New England there exist only seven distribution utilities: this means that, given the possibility to bid

[5] Of course, consumers are more or less risk adverse: some of them will require the LSE to take the risk of price volatility in their behalf, by paying for this service, some would prefer to be totally exposed to price variations and sign a contract for the LSE to pass them down real time prices. Any combination of risk sharing between the LSE and the consumer can be thought as acceptable for a possible type of contract.

their demand curve, each of their aggregate demand profiles would have a significant impact on the system wide demand profile.

The problem of load aggregation should have two formulations: a static one, where the load is assumed to be not responsive to price variations and a dynamic one where end-users adjust their load profiles according to price signals. The first formulation is the basic one, and provides an optimal portfolio of loads to serve when the load profiles of the end users are given and fixed with respect to price. The second one considers the loads' responsiveness to price and models consumers' behavior according to price changes. This section focuses on the first case. A model for load responsiveness is developed in a later section of this paper.

As stated above, the first goal of the LSE is to create an aggregate load profile that is flat. The typical case is that of a LSE producing its own power. The optimal profile (as an aggregation of the load profiles of the LSE's customers) will match a *physical generation profile*. A generator entering a contract for selling power will prefer to run its plant at a constant power for the entire day. If the output is maintained constant, and the generator is running continuously, the generator will avoid the costs of ramping up and down and of uncertainties in production. The selling price profile will be accordingly flat. The optimal power profile for the LSE's demand will match this flat price profile.

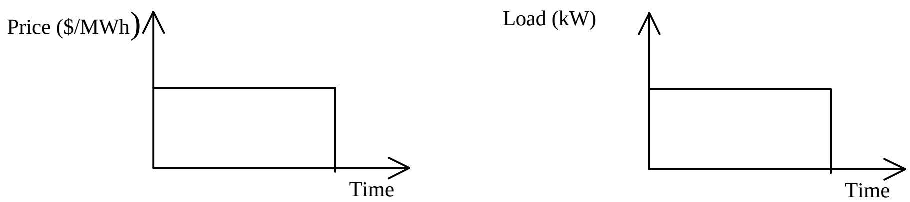

Figure 4 – Price and Load profiles

In order to have a flat aggregate demand profile, the LSE will have to combine different classes of customers. A simple example of this is the combination of a typical daily residential load profile with a commercial load profile. Together they sum up to a flat profile for most of the day and night with two likely spikes, one in the morning and another around 5pm. This corresponds to the times when people leave/arrive home for/from work, but most businesses are already/still open. In other words, the LSE will try to bring together customers whose load profiles have peaks that can be modeled as independently identically distributed.

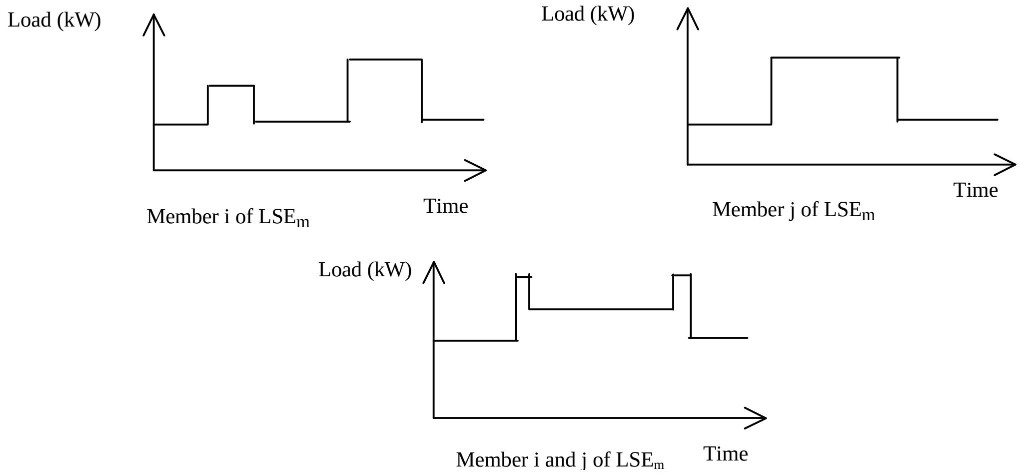

Figure 5 – Different load profiles and their aggregation

In order to keep load variations from forecasted demand within a given deviation band the aggregation strategy will be to bring together consumers whose deviations have a negative correlation with each other and therefore cancel out. The LSE needs a great deal of information about each consumer's deviation range around the forecasted profile. Each consumer's contract will permit it to deviate within specified margins, outside of which service is not guaranteed. The LSE needs information about the degree to which deviations of the loads will be correlated. To define a measure of correlation of deviations, a function needs to be introduced that will measure the relationship between the sum of the deviation bands. For example, see the function λ as defined in [10]: it will be a monotonically increasing function of load correlation, starting at λ=1/vn for n independent identically distributed loads and bounded by λ=1 for fully correlated loads.

Since demand deviations from the forecasted demand need to be purchased on the spot market, the aggregate load profile should have a negative correlation with the average price during the day, the week or the season. In order to optimize the purchasing profile, the goal of the LSE aggregation scheme will be that the sum of the aggregated deviation band will inversely reflect the market price profile. So for example in Figure 6 if the price curve has a maximum at the peak load, the load aggregation should be made to have a minimum level of deviations at that same time.

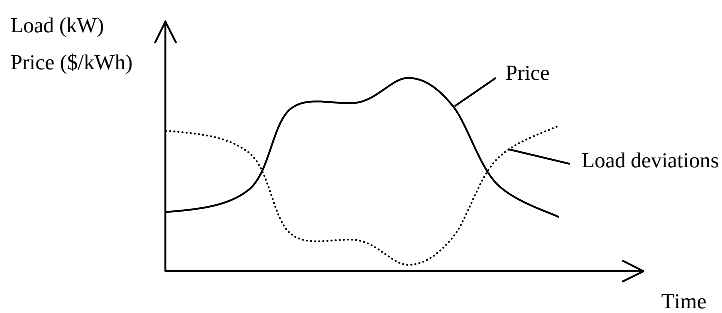

Figure 6 – Price and load profiles

3. MARKET BIDDING STRATEGY FORMULATION FOR THE LSE

An LSE needs to develop a purchasing strategy when buying energy on behalf of its consumers. Customers enter into contracts with the LSE specifying how much they are willing to pay for the amount of electricity they consume in each period. Based on this information, the LSE aggregates the demand functions of all customers and constructs its own aggregated demand function. A power purchasing strategy of an LSE must be concerned with the following issues:

- Which sources of supply: bilateral contracts, an hour-ahead market, and/or a real-time market?
- How much power it should buy from a day-ahead market and how much from a real-time market?

3.1 PURCHASING POWER IN THE FORWARD OR SPOT MARKETS

An LSE plays a similar role in an electricity market to a power producer. The differences are that it supplies negative power (or consume power) and it is exposed to greater electricity price risk, since prices of electricity can be very volatile and very high at times.

In a spot market, with demand-side bidding, an LSE, as well as a power producer, can affect the market price [by exerting market power]. Like a power producer, in which the scheduling results from the day-ahead market guides how to optimally run its generating units (under unit-commitment constraints), an LSE purchases power from a day-ahead market, which is usually held before the actual dispatch. After the market is cleared it is notified of total capacity of power scheduled for the next day and the associated total cost. This provides an opportunity for the LSE to send price signals to its customers in advance to adjust their usage accordingly.

In a forward, or futures, market or a spot market, power producers, LSEs, and marketeers, who do not necessarily own any physical assets or obligation in that spot market, purchase or sell electricity in advance. Although a day-ahead market is considered a forward market, electricity in the forward or futures market is generally traded much further in advance (NYMEX has a futures trading period up to 18 months). With a forward market, some risk-averse LSEs would seek contracts to hedge risks according to electricity price and (real-time) demand (a possible mathematical formulation can be found in [19]). Some power producers or LSEs who want to lock in selling or purchasing prices, respectively, can sell or buy a long-term contract through the forward market. The contracts traded in the forward or futures markets are solely financial obligations and are divorced from physical delivery. Moreover, prices from the forward market indicate the future market conditions. This provides signals for investors to expand new investment in generation capacity (including distributed generation), and for an LSE to have a plan dealing with its future obligation to serve demand [1][5][9][13][16][19].

3.2 MODELING A SIMPLIFIED LSE

Suppose that an LSE has no market (monopsony) power and is unable to affect the market price. The LSE takes prices as given by the market, and optimizes its objectives accordingly.

3.2.1 UTILITY FUNCTION OF A CUSTOMER

The utility function of a customer j indicates its willingness to pay for different electric power demand levels, and it varies from one customer to another. The utility function of each customer depends on their preferences and behavior, which are time-dependent. The preferences indicate the type of services it might demand. Some customers might not want to adjust their activities to consume electricity when it is inexpensive, and would rather choose to pay real-time prices for unlimited usage at any time. On the other hand, some customers might prefer lower electricity rates in exchange for a price-dependent or time-of-use service. A rational customer chooses a package of services from an LSE that will maximize its total benefits (utility minus expenses). For example, the services offered by an LSE should reflect customer preferences for:

a) Time-of-use
b) Market Price Dependent Usage
c) Unlimited Usage
d) Combinations of above
e) etc.

Depending on preferences, customers can choose to optimize the price and quantity of contracts over either a short or long time horizon.

3.2.2 UTILITY FUNCTION OF AN LSE (SERVING SEVERAL CUSTOMERS)

An LSE aggregates all customers' preferences or types of services and creates its utility function (which could simply be the summation of its contractual obligations to the customers). The objective function of an LSE is to maximize profits and perhaps to minimize electricity price risks. Based on an obligation to serve customers and its own risk preference, the LSE optimizes its decisions to buy electricity from either a spot market or through bilateral contracts.

These are difficult stochastic optimization (dynamic programming) problems, whose solutions would determine when and how much electricity to be purchase either from the spot market or through bilateral contracts. The results from this problem indicate which sources of suppliers the LSE should buy the power from, or could favor investments in small local distributed generation.

The problem formulation conforms to the profit-maximization and variance minimization objective function. (See problem formulation in [16]).

We describe here an aggregate demand function, modeled as a piece-wise constant demand. We assume here that an LSE possesses no market power so it optimizes its benefits based on given market prices. To define the aggregate demand function for an LSE, one needs to understand the method that the LSE uses to combine the demand of its customers. Suppose that the market is perfectly competitive and that an LSE could not manipulate the market clearing prices via its demand-side bids. An LSE can potentially provide multiple choices of services for its customers; for example, customers provide the maximum prices (willingness-to-pay or marginal utility) at which they are willing to consume electricity:

a) Customer j will pay r_k^1 \$/energy-unit if it consumes power less than $D_k^{max,1}$ within one period *k*, given that the market price of electricity is less than $P_k^{max,1}$.

b) Customer j will pay r_k^1 \$/energy-unit if it consumes power less than $D_k^{max,1}$, given that the market price of electricity is less than $P_k^{max,1}$; and r_k^2 \$/energy-unit if it consumes power less than $D_k^{max,2}$ within one period *k*, given that the market price of electricity is less than P_k^2.

c) Customer j will pay r_k^3 \$/energy-unit if it consumes power less than $D_k^{max,3}$ within one period at any price of electricity.

Suppose that an LSE purchases power from a spot market. If the market price is lower than $P_k^{maz,1}$ and/or $P_k^{max,2}$, customers will benefit from buying power with lower prices than their willingness-to-pay ($P_k^{max,1}$ and/or $P_k^{max,2}$). Therefore, a possible method is to view the contracts a) and c) that an LSE sells a "strip" of "put" options (for all periods) to a customer with a striking price $P_k^{max,i}$ for each option in the strip. Put options with striking-prices equal to $P_k^{max,1}$ and/or $P_k^{max,2}$ are shown in Figure 7.

The payoff (p_k^i) of an option or the price of a put option with striking price $P_k^{max,i}$ is equal to:

$$p_k^1 = \max(P_k^{max,1} - P_k^s, 0) \tag{1}$$

and/or $p_k^2 = \max(P_k^{max,2} - P_k^s, 0)$ (2)

With a well-defined price process (such as described in [16]), one could use a Monte Carlo simulation or a closed-form solution (if it is possible) to determine the value of these options. However, one must keep in mind that a valid price model to value these options correctly must capture the effect of demand adjustment on prices.

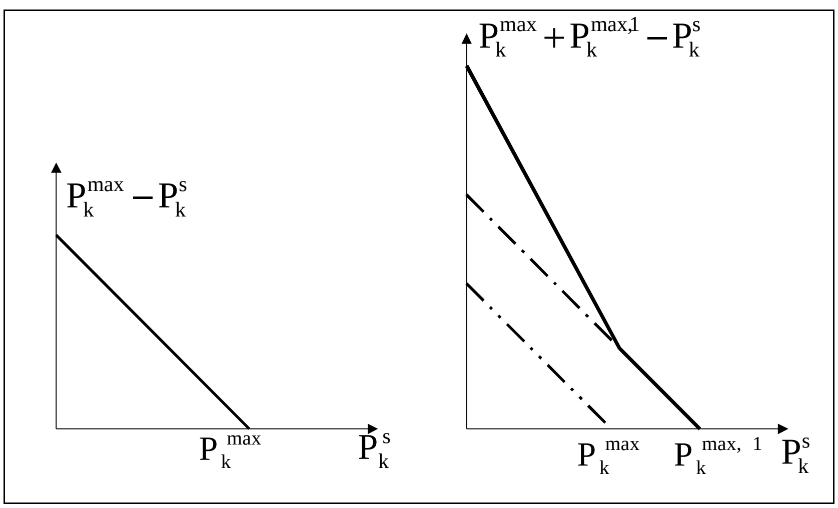

Figure 7 - Example of Put Options Offered by an LSE to Customers

Since each put option is assumed to be for buying one energy-unit (such as MWh or kWh), to fulfill demand customers must buy multiple (but less than $D_k^{max,1}$) of these put options. This indicates that the amount that a customer will consume in real-time varies stochastically with the spot price. Perhaps, a customer should prepay for the amount of electricity it expects to consume.

If it consumes more than anticipated, a penalty-fee might be imposed. Also, if it consumes less than anticipated, a compensation fee might be imposed. Note that if an LSE buys an option, it will not be compensated for under-consumption.

In order to issue these options, an LSE might need to collect a "load profile" of each customer, together with the price profile so that the LSE could roughly determine "applicable" strike prices and a nearly accurate size of contracts to be bought. Without real knowledge of the customers' load profiles and price preferences, the LSE might find that customers violate the contract more often in real-time, because the contracts do not match with customers' characteristics. In this case, the LSE would likely end up losing money because it would need to make up the difference by buying on the spot market and is then subject to greater price volatility.

Suppose that these options are sold to the customers, the remaining question is how an LSE could monitor whether the customers adhere to their contracts. "Smart" equipment to automatically adjust electricity usage in real-time is needed and discussed in the last section of this paper.

3.2.3 How does an LSE create demand elasticity?

Suppose an LSE has N customers to serve. It combines the demand functions (willingness-to-pay) from each customer to form an aggregated marginal utility curve, or an LSE aggregates a series of put options with various strike prices, stacking them from the most expensive willingness-to-pay or the most expensive strike prices. Note that the utility function of an LSE is obtained from piece-wise integration of its aggregate marginal utility function. By using an options pricing formulation, it is implied that under supply scarcity conditions, with high day ahead electricity prices, and where the LSE has to purchase electricity from the spot market, a customer who has the least expensive striking price will be "curtailed" first.

Note, however, that the "put option" contract does not include the insurance for the real-time service. Thus, the LSE faces a different optimization problem if responding to real-time prices and the need to curtail in real-time. However, contracts for the ability to be curtailed or insurance schemes for levels of curtailment need to be a part of the portfolio of products offered by the LSE to hedge its own price and service contract risks **Error! Reference source not found.**.

3.2.4 Other LSE Strategic bidding and Modeling Issues

In order to be able to hedge risk and maximize profits, the LSE would need some mechanism for either paying or penalizing customers who under or over consume. In the case of under consumption (i.e, several customers reduce load beyond expected value), the LSE could sell this power back to the retail market. However, if customers over consume, the LSE may be required to purchase additional power and/or services from the spot or ancillary service market to cover the load and system stability. The LSE would need to set specific contractual provisions to account for these perturbations in order to ensure costs are adequately covered. Figure 8 illustrates the idea of determining the appropriate cost for curtailment compensation or over-consumption fees:

Over-consumption fee paid by consumer j to an LSE over-supply at time k:

$$F_k^{O,j} = \max(\alpha_k^{0,j} \cdot (D_k^j - D_{k,\max}^j),0) \tag{3}$$

Compensation payment of an LSE to load j from load curtailment obligation at time k:

$$F_k^{C,j} = \max(\alpha_k^{C,j} \cdot (D_{k,\min}^j - D_k^j),0) \tag{4}$$

Where:

α_k^C Compensation payment per unit power from curtailment at time k

α_k^O Over-consumption fee per unit power at time k

D_k^j Power consumption at time k

$D_{k,\min}^j$ Minimum supplied power by an LSE without compensation payment at time k by customer j

$D_{k,\max}^j$ Maximum supplied power by an LSE without over-consumption fee at time k by customer j

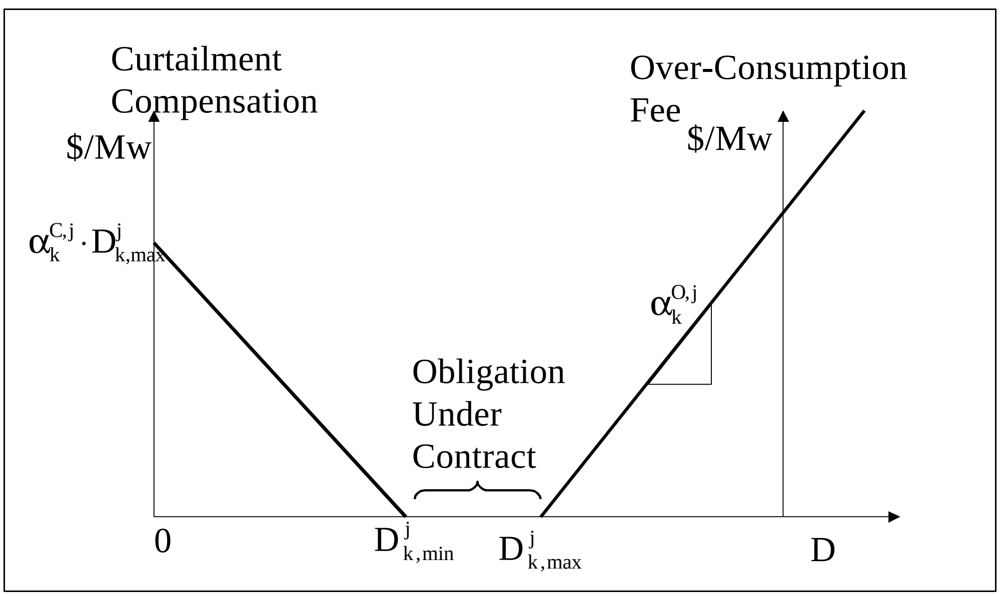

Figure 8 - An Example of Curtailment Compensation and Over-consumption Fee

4. Effects of Demand Response on Load and Market Price

In order to illustrate the effect demand response can have on aggregate load and the market clearing price an example is presented of a simple thermostat control device used to shift peak load demand for residential air conditioning. In summer months, air conditioning likely presents the largest portion of peak load in most of the U.S. In California, it accounts for 29% of the peak demand with residential AC load contributing 14% and commercial AC load contributing 15%. Because AC is the primary load, it is useful to examine the amount of load shifting or shedding required to have an effect on the market price. A simple control scheme for residential AC is used as an example to demonstrate the ability of a smart consumer (or AC unit) to respond to rising prices in anticipation of a high peak load price.

The example is presented in two parts. The first part evaluates the level of control required to maintain indoor air temperature within a comfortable range between 68°F and 72° F when the outside temperature is 88°F over the typical 6-hour afternoon peak demand period and calculates the cost savings compared to maintaining an AC unit at a constant temperature of 70°F throughout the period[6]. The amount of peak shifting calculated is then extrapolated to estimate the effect on total peak load and wholesale market price in CA on a typical summer day. The second part of the example examines other elements of the residential load and simple behavioral changes that can also contribute to shifting peak load and reduce consumer cost using relatively inexpensive and commercially available control technologies to manage load.

4.1 POTENTIAL FOR PEAK LOAD SHIFTING AMONG RESIDENTIAL CONSUMERS

The following example illustrates the potential peak load savings (both economic and power) resulting using from a load control scheme to the air conditioning of a typical household. The control scheme, based on a control model outlined in [4], uses real pricing data selected from the CA ISO for a summer day last year (2000) to control the output of the air conditioning system for optimal cost savings. The result provides the potential economic savings a consumer could have enjoyed from employing such a control scheme, as well as the resultant reductions in peak load power usage.

The Consumer's objective is to minimize the cost of air conditioning while maintaining the indoor air temperature within a certain range.

$$Min_e C_{ac} = \sum_i P_i * q_i \qquad (5)$$

s.t.

$$0 \le q_i \le q^{\max}$$

$$T^{\min} \le T_i \le T^{\max}$$

where:

$T^{min} = T^{ideal} - d$

$T^{max} = T^{ideal} + d$

d = Acceptable temperature deviation

q_i - energy (kWh) consumed by air conditioning in hour i.

P_i - price of electricity (\$/kWh) in hour i

The hourly household temperature is determined by (parameters and definitions are given in Table 1):

$$T_{i+1} = \varepsilon T_i + (1- \varepsilon)(T^o - \eta * q_i / A) \qquad (6)$$

[6]Normal cycling of the thermostat will cause the temperature to "naturally" vary around the ideal, however this cycling will occur such that the amount of power consumed will be evenly distributed among each hour.

To maintain a constant indoor temperature at 70° with and outside temperature of 88°F, i.e., (T^o_i - η*e_i /A) = T_i = 70°, the power required is calculated to be q_i = 1.008 KW continuous. Similarly, for an outside temperature of 75°, the power required is 0.28 KW continuous and for an outside temperature of 95°, the required power is 1.4 KW continuous.

Table 2 compares the results of applying a load control scheme to the AC versus the base case of allowing the AC to run on a single thermostat setting. It is assumed that the consumer is indifferent to indoor temperature fluctuations between 68 and 72° (T^{ideal} = 70°, d = 2°). During the 6 highest price hours of the day (from 1pm to 6pm), an 83% reduction in peak demand for air conditioning can be achieved by load shifting without moving outside the indifferent temperature range.

There is, however, a large increase in consumption in the hour immediately before and immediately after the 6 peak hours. The overall power consumption is nearly equivalent for the eight hour period using the load control scheme. By expanding the control scheme beyond the 8 hours used in the example, additional [cost] savings may be realized.

Table 1 – Parameters and Values of Residential AC control model

Variable	**Value**	**Description**
T_0	70	(°F) initial temperature
η	2.5*	Efficiency of AC
q_i		Energy use by AC in hour i (KWh)
q_{max}	3.5	(KWh) maximum energy used by AC per hour
ε	0.96*	System inertia = exp(-τ /TC)
TC	25*	(hr) Time Constant for home
τ	1*	(hr) duration of control period
A	0.14*	(KW/°F) Thermal Conductivity
T^o	88°	Outside Temp (°F)
T_d	70°	Desired household Temp
T_{max}	72°	Highest acceptable household Temp
T_{min}	68°	lowest acceptable household Temp
T_i	70° (initial value)	Current household Temp

*Parameters from [13]

Table 2 – Comparison of control scheme to base case

		Constant Temperature Case			**Load Control Case**		
Hr	Price $/MWh	Temp	Output (KWh)	Cost (mils)	Temp	Output (KWh)	Cost (mils)
12	231.96	70	1.008	234	68.2	3.5	812
13	442.61	70	1.008	446	69	0	0
14	750	70	1.008	756	69.8	0	0
15	750	70	1.008	756	70.5	0	0
16	724.75	70	1.008	731	71.2	0	0
17	749.95	70	1.008	756	71.9	0	0
18	353.46	70	1.008	356	71.8	1.041	241
19	192.94	70	1.008	194	70	3.5	812
		Total	**8.064**	**4,229**		**8.041**	**1,865**

The price data from a single summer day (June 29, 2000) in Burbank, California where the temperature reached 88 degrees, is shown in Table 2. It is observed that a single residential consumer could have saved 55% off their cost of power for air conditioning alone. This does not include any effects of reducing the peak prices due to lowering the overall peak load discussed later (many consumers would have to employ such a scheme for this effect to become significant).

In addition, the power consumption would be lower in a more detailed example where daily fluctuations in actual outdoor temperature are accounted for. For simplicity, this example assumes a constant outdoor temperature of 88° over the entire 8-hour period. The amount of potential shifting can actually be enhanced, depending upon the time it takes for the temperature to change (i.e., it will take less power at the end of the cycle to bring the temperature down from the upper thermostat bound if the outside temp has also come down during that time and less initial energy to achieve the low temperature bound if it is not yet at the hottest time of the day.)

4.2 Extrapolated Effect of Controlling Residential AC on Market Price

Small adjustments to end-use devices enabled by simple control schemes and real-time pricing can have a tremendous effect on electricity market prices and efficiency. Using the simple residential AC load shifting example and applying it to the entire CA market demonstrates this point. The summer 2000 peak load in CA was approximately 51,000 MW. The residential AC portion of the peak CA load is 14%. Extending the calculated 83% load shift to the entire residential load over a 6-hour peak, we calculate a system wide peak load reduction of 12% [15]. Given publicly available data on the operating costs of CA power suppliers, an estimate of the perfectly competitive market supply curve is shown in Figure 9. Note, that the aggregate supply in this figure falls short of the 51,000MW peak demand; this implies that CA needs to import power to meet its daily peak demand.[7] When demand outstrips supply, the cost for energy is off the chart. What is important to take notice of is the relative trade-off between quantity and price between point A where the slope of the curve begins to become quite steep (i.e., currently representing expensive peak load capacity) and point B where the supply is completely inelastic. A responsive demand can shift the demand curve downward from point B to A and reduce price dramatically with only a small reduction in load. Thus, a 12% load reduction realized from simply shifting the residential AC load, results in a 47% price reduction. In a very competitive market where suppliers are bidding in actual cost, this reduction would result in a price difference from $71.58/MW to $41.36/MW for a total savings of $11.4 million for the six-hour period.

[7] Increased demand from Northwest power routinely imported to CA left CA supply short and was one of the major reasons for high wholesale power prices in the CA market .

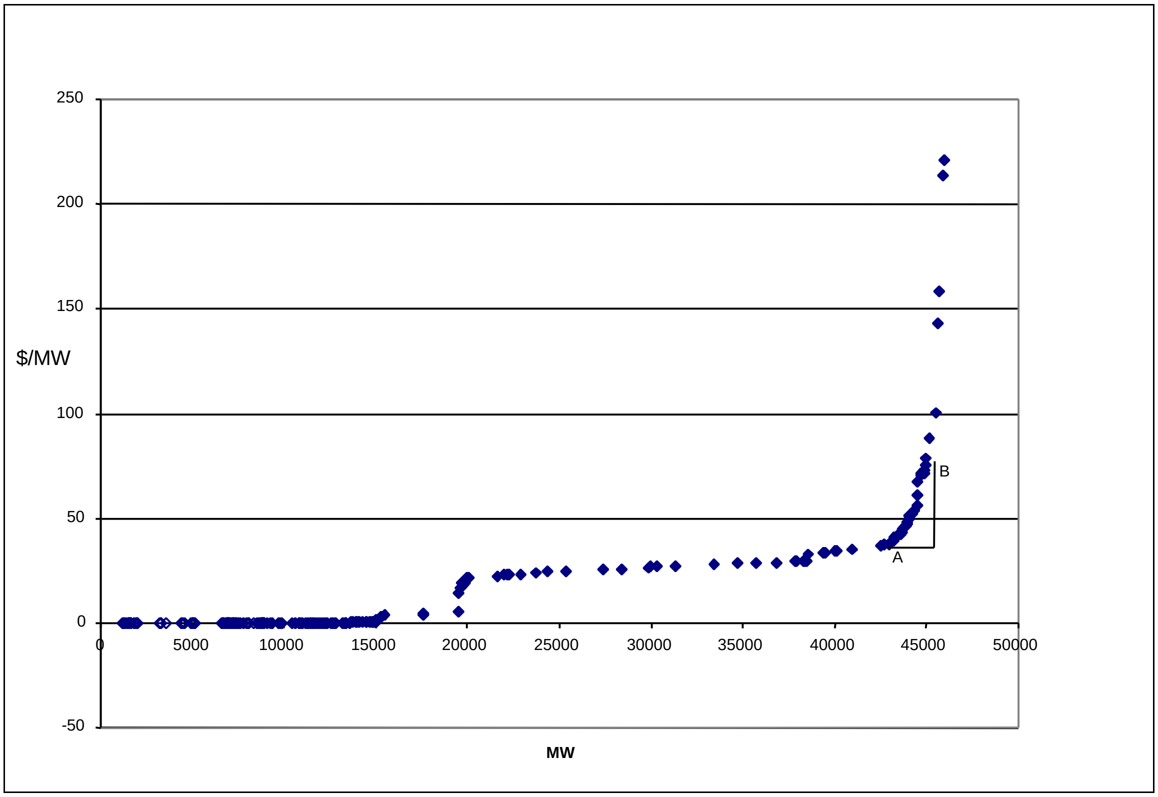

Figure 9 – California Generator Market Cost Curve

To further illustrate how a responsive demand could have mitigated the market crisis in CA, Figure 10 shows the unconstrained market clearing price and system load for June 29, 2000, a day where the wholesale market hit the $750/MW price cap from 2:00 PM through 5:00 PM. The difference between the peak demand for the day of 41606 MW at 4:00 PM and the demand at 11:00 PM of approximately 37536 MW is 5%. This 5% difference in demand accounts for a 41% difference in spot market price. If 35% of the residential load is shifted according to the air conditioning control scheme discussed above, cost savings during the 6-hour peak demand would be approximately $51.44M. However, this would lead to a very high demand spike in the hours proceeding and following the 6-hour peak causing very high prices and could completely offset the majority of cost savings. A more optimal situation would be to have just enough residential load shifted in order to flatten out (and expand) the peak period without causing large spikes on either end. This situation requires an approximately 15% shift of the residential load which leads to an overall peak reduction of 2.1% and an overall cost savings of $29M [15],[17]. Expanding the control scheme beyond the 8 hours simulated above would further flatten the overall load profile and lower the peak load even more.

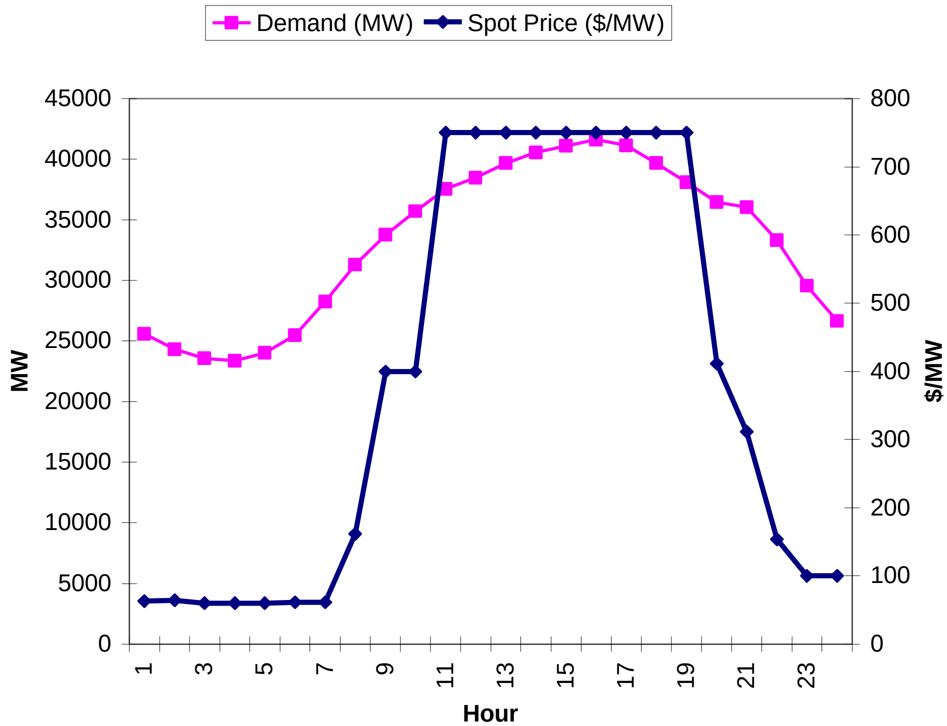

Figure 10 – Hourly Unconstrained market clearing Price – June 29, 2000

4.3 OTHER OPTIONS FOR REDUCING THE OVERALL PEAK RESIDENTIAL LOAD

The following scenario provides additional examples of the effect of load shifting various types of residential end-use loads on peak shaving. Table 3 provides a list of common residential appliances and the concomitant contribution to peak and off-peak loads[8]. For these examples, calculations are based on the following assumptions:

- 6 Peak hours per day
- 1/3 of the total average household consumption is during the 6 peak hours
- Average Household consumption is 2100 KWh per month during the summer[9].
- Assumes the following breakdown of peak (6hrs/day) to off peak (18hrs/day) usage for appliances[10]:

---

[8] based on load profile data from Niagara Mowhawk, 2001
[9] Assumes family of four with air conditioning and electric stove and water heater.
[10] The assumed breakdown is based on the aggregate load profile from Niagara Mohawk, so that 1/3 of all power is consumed during peak hours. A 25% to 75% split indicates uniform usage.

Table 3 – Load Composition – Residential Appliances[11]

Appliance:	Monthly Use (kW)	% Peak	% Off-Peak
Air Conditioner	642	40%	60%
Water Heater	513	25%	75%
Refrigerator	112	25%	75%
Range	375	50%	50%
Freezer	100	25%	75%
Dryer	85	60%	40%
Washer	12	60%	40%
Dishwasher	12	60%	40%

Based on the above assumptions, the following reductions in peak load are possible:

1) Shifting 50% of peak use of washer, dryer, and dishwasher to off peak times leads to a *4% reduction in peak household usage*

2) Using Load control for Air Conditioning leads to a *28% reduction in peak household usage*

3) Using similar load control devices for water heater, refrigerator, and freezer leads to a *20% reduction in peak household usage*

4) Incorporating all of the above leads to a *52% reduction in peak household usage*

5) 100% shift of washer, dryer, and dishwasher leads to a *8% reduction in peak household usage*

6) 100% shift of washer, dryer, and dishwasher + load controls (2 and 3) leads to *56% reduction in peak household usage*.

4.4 Control Technology Implementation

The potential for reducing peak household consumption during the summer by implementing these few simple steps appears to be quite large. By simply modifying their behavior in washing clothes and dishes, consumers can reduce their peak demand by between 4 to 8%. This would require no additional investment whatsoever, and have a relatively minor impact on consumer utility (quality of life). Manufacturers could easily add timers to washers, dryers, and dishwashers for little additional cost, but there is currently little demand for such features. Current pricing structures (i.e. flat rate) provide little or no incentives for consumers to modify their behavior, however. Most consumers are not even aware that electricity costs more during certain times of the day.

[11] Appliance data from: Otter Tail Power Company Website, May 2001. http://www.otpco.com/home/edu_usage.asp

Switching to a real time pricing scheme would require the installation of automated metering and a communications infrastructure [18], as well as regulatory changes. The potential savings from peak reduction, however, is large enough to pay for the installation of such infrastructure.

Implementing a load control scheme for AC or other energy (vs. power) consuming appliances is slightly more difficult, but also feasible. The potential for shifting load from peak hours using such devices is quite large. In addition to the communications infrastructure necessary to implement real time pricing, the controllers themselves will need to be developed, as well as the software to determine to optimal usage profile for the appliances. Programmable thermostats currently exist for AC and heating, but it would currently be rather difficult for the average consumer to install such a thermostat and link it to his/her computer where a control algorithm would automatically program the usage according to the day's hourly electricity prices. Standardization of this equipment is necessary so that consumers can "plug and play" with little or no hassle.

Incorporation of controls on refrigerators, freezers, and water heaters would be relatively simple and low cost. Once again, the lack of a demand for such controllers is the main obstacle to manufacturers including them as a standard feature with these appliances.

5. Conclusions

This paper asserts that load-side management is an important missing piece in the current US electricity markets. Problems with existing deregulation policy and practice limiting demand responsiveness were identified and discussed. The benefits of demand responsiveness were discussed with respect to the effect on market electricity price, which creates potential profit for an LSE, and service choices and lower costs for the end-use consumers. The ability to expose customers to real-time pricing provides the needed incentive to create demand elasticity. LSEs, through better understanding of load profiles, customer's demand elasticity and willingness to reduce or shift load in exchange for compensation, can more effectively bid demand into the wholesale electricity markets and reduce overall market price (moving the demand curve from point B towards point A in Figure 9).

As shown for an anticipated price profile, a residential customer can determine optimal (least cost) consumption patterns without forsaking comfort (in this case: a cool house on a hot day). When this type of peak shifting behavior is conducted on a more aggregated level, e.g., by the LSE through the use of appropriate appliance controllers, the aggregate reduction can result in a lower hourly wholesale electricity price. The examples provided in this paper present a simple case illustrating the important effects of load response to anticipated electricity price. However, the dynamics of a real-time market are not accounted for and are beyond the scope of this paper.

The assertion that responsive demand can correct some of the observed and anticipated market problems in the future is at least empirically supported by the arguments set forth in this paper. However, more research is needed on a variety of levels to provide a sound basis for future policies and investments. Of particular interest is the role of state and local government in developing and implementing deregulation policies. In the short-term, government has the ability to conduct pilot studies on the effect of demand responsiveness, load aggregation and consumer elasticity on market price. Well-designed experiments could be very helpful in fully characterizing consumer demand elasticity, the effects of load aggregation, overall benefits of load response, the infrastructure technologies and costs, and affect on market behavior. These studies would be done in collaboration with LSEs and technology companies interested in analyzing the performance of this system.

Broader policy issues regarding the role of government in the market, beyond pilot programs and policy makers are also important. Can or should the government participate as a market player either a seller or buyer of power on behalf of its constituents? Should municipalities take on the role of an LSE on behalf of residential and small business customers and government services in order to negotiate better contracts and perhaps, affect the market price. What would the economic and political ramifications of governments role? What policies would need to be developed to enable this action? Is it the role of government to serve as a catalyst to foster demand responsive activity?

Other important questions that need resolution include: understanding the economic trade-off between investments in centralized generation, and transmission and investments in well-designed control schemes that enable distribution-side demand responsiveness and distributed generation is essential to optimizing future investments in power system infrastructure.

6. REFERENCES:

[1] Bertsekas, D., *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, New Jersey. 1987.

[2] Black, J. Ilic, M., *Survey Of Technologies And Cost Estimates For Residential Electricity Services*, Accepted for Publication, IEEE Proceedings, Summer 2001.

[3] California Public Utility Commission - Office of the Ratepayer Advocate, *Monthly summary choosing direct access*, April, 2001 http://www.cpuc.ora.gov.

[4] Constantopoulos, P., Schweppe, F. and Larson, R., 'ESTIA: A Real-time Consumer control Scheme for Space Conditioning Usage Under Spot Electricity Pricing' C*omputers Operations Research*, Vol 19, no. 8, pp 751-765, 1991.

[5] Cox, J., Rubinstein, M., *Options Markets*, Prentice Hall, New Jersey, 1985.

[6] Fumagalli, E., Black, J., Ilic, M., Vogelsang, I., *A Reliability Insurance Scheme for the Electricity Distribution Grid*, MIT Energy Lab Working Paper EL 01-001 WP, January 2001 and IEEE Proceedings, PES Summer 2001.

[7] Fumagalli, E, *Quality of Service: Comparing the Italian Regulatory scheme with an Insurance-Based Approach*, Gas and Electricity Forum, Scuola Superiore E. Mattei, Milano, Italy, 21-22 June 2001

[8] Ho Y. C., Chiu D. M., *A model for indirect load-management through optimal pricing*, NSF Research Paper.

[9] Hull, J.C., *Options, Futures, and Other Derivatives, 4th ed*. Prentice-Hall, 1999.

[10] Ilic M., Skantze P., Yu C.N., Fink L., Cardell J., *Power Exchange for frequency control (PXFC)*, Proceedings of the IEEE PES Winter Meeting, New York, NY, January 1999.

[11] Massachusetts DOER, 2001.

[12] Pennsylvania Public Utility Commission, Office of the Ratepayer Advocate. Monthly Data on customers switching from default utility, April 2001.

[13] Pindyck. R.S., Rubinsfeld, D.L., *Microeconomics 4th ed.*, Prentice-Hall, New Jersey, 1997.
[14] Radar, N. Hempling, S, *Promoting Competitive Electricity Markets Through Community Purchasing: The role of Municipal Aggregation*, American Public Power Association, January 2000.
[15] Rohre R., *Data from California Energy Commission*, Energy Demand and Forecasting Group.
[16] Skantze, P., Gubina, A. F., Ilic, M.D., *Bid-based Stochastic Model for Electricity Prices: the Impact of Fundamental Drivers in Market Dynamics*, The MIT Energy Lab Technical Report EL 00-004, MIT, November 2000.
[17] Supply-Demand bid curve data June 29, 2000 from the California Power Exchange http://www.calpx.com
[18] Unconstrained hourly market clearing price from the California ISO website. http://www.caliso.org
[19] Wagner, M., *Hedging Optimization Algorithms for Deregulated Electricity Markets*, Master of Engineering Thesis, EECS Department, MIT, June 2001.