Scalable, Real-Time Electricity Demand Response Optimization in the Smart Grid

Intel Science Talent Search

November, 2013

Abstract

With the growing smart grid effort, and the fine-grained electricity demand data that comes with it, utility companies have gained the resources to implement Demand Response, through which electricity prices are dynamically changed to incentivize consumers to cut down on electricity consumption. However, current Demand Response programs do not utilize the bulk of this new grid data. These programs, which usually regulate electricity consumption on a pre-determined fixed set of facilities during a set interval, are inefficient and inflexible because of their coarse-grained and static nature. This research introduces an optimization algorithm to implement fine-grained Demand Response in a real-time, scalable fashion, to balance electricity loads continuously over the duration of the day. Through optimization, we prevent utilities from wasting electricity and paying large premiums to combat supply shortage, by fitting electricity demand to the ideal shape: a flat line. This research first develops a simple optimization model that takes into account a basic 2-price Demand Response program for a set of facilities, that optimizes over a single time interval. This optimization is then expanded to continuously optimize over the course of an entire day containing multiple time intervals. To solve this optimization efficiently, we present a relaxation which permits the use of a greedy algorithm. We present results using real-world smart meter data, which shows the efficacy of this algorithm in achieving a desired electricity load profile, while maximizing utility profit.

1 Introduction

As the smart grid effort continues to become pervasive throughout both the United States and the world, utility companies are beginning to look for ways through which they can not only reduce overall energy consumption, but eliminate the expensive and unnecessary generation costs that come with extreme grid demand. New smart grid technology permits a granular view of electricity consumption, which makes it possible to do a better job of electricity demand and price management. Using this newly available data, the government and utility companies have partnered to implement Demand Response [1], through which a dynamic pricing is used to incentivize consumers to cut down their electricity consumption when demand for electricity is extremely high. This works off a simple supply and demand metric: when supply is low, prices are raised to inhibit demand.

At the moment, scheduling Demand Response events ahead of time, over a small area, works as a short-term fix to the problem. More often than not, though, electricity demand significantly lowers during the event, but peaks again both before and after the event occurs. As a result, utility companies are left using a coarse-grained, static process, which is not as effective as it can be. However, as the smart grid and electricity demand grow, this newly available fine-grained data (result of the new technology) allows for a better solution, in which Demand Response is optimized in a continuous, dynamic fashion.

As it currently exists, electricity load over an area is most accurately represented by a sinusoid: as the day progresses, electricity consumption peaks around mid-afternoon, and then returns to a period of low demand. At the moment, when electricity demand exceeds capacity, utility companies turn on peaker plants, to compensate for the difference. A peaker plant is essentially a backup power plant that is only turned on during times of peak demand (an extremely small portion of the year). For the most part, peaker plants lay unused, but unfortunately, utilities spend hundreds of million dollars a year for construction and maintenence [2]. By finding a real-time, continuous optimization of Demand Response, to flatten the demand sinusoid, we can effectively cap electricity demand, saving money for both the utilities and the consumers, as well as save millions of gigawatt-hours a year. In addition, we will be able to eliminate the need for these peaker plants.

Other researchers, to try to solve this problem have proposed control algorithms to dictate how utilities should control consumer demand through direct load management [3, 4]. Others have used direct load management to optimize utility expenses for generation [5, 6, 7, 8], propose theoretical models to improve utility profit margins [9], or develop pricing models to maximize social welfare [10]. In general, the nature of the past work either focuses on using direct load control to manipulate consumer demand, or uses theoretical models to simulate solutions on small datasets.

This research introduces a greedy optimization algorithm to prevent utility companies from overpaying on electricity generation costs by effectively shaping electricity demand in such a way that grid consumption always remains below a certain threshold. This creates the desired flat-line load shape. Unlike past research, this algorithm takes into account the stochastic nature of grid dynamics; that is to say, that the optimization takes into account the stochastic nature of consumer choice, rather than rely on methods of direct load management like papers past. Furthermore, the adaptation of this optimization to solve for both Demand Response optimization in the United States, and outage prevention in developing nations reflects the flexibility and effectiveness of the solution algorithm.

2 Demand Response

A key part of Demand Response and the smart grid is the ability to monitor electricity on a granular level. To do this, utility companies and the government have invested in Advanced Metering Infrastructure (AMI). AMI differ from region to region, or based on provider, but its key function remains the same; to provide regular interval data regarding electricity usage of the facility at hand [11].

In the last few years, especially as the US Smart Grid program has taken off, the number of AMI has increased, at an almost exponential rate. Figure 1 shows the number of AMI in the US by type per year, from 2007 2011 [12].

Table 10.6. Advanced Metering Count by Technology Type, 2007 through 2011

Year	Residential	Commercial
Automated Meter Reading (AMR)		
2007	25,785,782	2,322,329
2008	36,425,943	3,529,985
2009	41,462,111	4,239,531
2010	43,913,225	4,611,877
2011	41,451,888	4,341,105
Advanced Metering Infrastructure (AMI)		
2007	2,202,222	262,159
2008	4,190,244	444,003
2009	8,712,297	876,419
2010	18,369,908	1,904,983
2011	33,453,548	3,682,159

Prior to 2010, the count was the number of customers, not number of meters.
Source: U.S. Energy Information Administration, Form EIA-861, "Annual Electric Power Industry Report."

Figure 1: Table of Advanced Metering Infrastructure (AMI) in the United States from 2007 2011. Notice the growth rate of AMI from year to year, indicating a growth in the US Smart Grid.

AMI data, for the most part, comes in on a 15-minute interval. This means that for each AMI meter that exists, there is a reading in Kilowatt-Hours (KWH) every 15 minutes, creating a 96-interval time series for each meter over the course of a day.

Because this data exists in such a fine granularity (facility to facility), utility providers can use it to establish trends, find discrepancies, and build regional pricing models for electricity. This gives utilities a better view of electricity demand, and pricing as a whole. Demand Response uses this AMI data in two key ways: Forecasting, and Response.

2.1 Forecasting

As the main purpose of Demand Response is to reduce electricity demand, it is extremely important that utility companies have a general idea of when demand for electricity will spike. Currently, the way they do this is by creating forecasting algorithms, built off of the data received from AMI in the region at hand.

Current forecasting models are generally built using regression models like GLM. However, the underlying principle is the same. All forecasting algorithms are trained off of past electricity interval data received from AMI, and factor in a series of other elements, including weather data, time of day, and past pricing schemas [10]. As a result, most forecasting algorithms are able to predict future demand to a certain degree of accuracy.

In some small regions, forecasting is done on the region as a whole, where data from all AMI in the region are used to build a single forecast. For other regions encompassing larger areas, where some factors vary from place to place (i.e. Weather), forecasting is done on a much smaller level, often meter to meter [13]. For the purposes of this research, the assumption was made that the forecasting was done on a meter to meter granularity, over a large area.

2.2 Response

The second part of Demand Response is how utility companies respond to spikes in electricity demand. There are two main components that are considered when it comes to responding to peak load events: the utility companies, and the consumers.

As mentioned before, the bulk of past research on Demand Response optimization is based on the idea of direct load control. With direct load control, utility companies directly control the electricity of demand for each facility. Examples of this may include utility managers shutting off wireless air conditioning systems or turning down smart thermostats by directly connecting to the facilitys wireless network [7]. This is extremely effective in facilities where wireless appliances exist, but otherwise, it gives utility companies little to no control over a facilitys demand. The second and most important factor to consider is the consumer. When a Demand Response event is called, it is ultimately up to the consumer to decide how much they wish to reduce their electricity consumption.

In the same way that Demand Response uses forecasting algorithms to predict future demand, other models are in effect to predict how much a certain consumer will reduce electricity consumption by during a certain event. Like demand forecasting, there are a multitude of factors that are taken into account, including previous Demand Response event data, pricing data, and the type of facility the event is being called on.

With both of these forecasting options, utility companies have an accurate representation of Demand Response load information per facility. Using these options, it becomes possible to figure out how to optimize Demand Response in real-time, allowing us to make Demand Response a sustainable option.

3 Methods

Optimization problems such as this one require an incremental approach in finding a solution algorithm. The problem is built in the following manner: First, a simple 2-price Demand Response model over a single interval is created. Then, the optimization is expanded to consider a continuous interval (the course of a day). Finally, the optimization statement is relaxed from an Integer Quadratic Program, and solved with an efficient greedy algorithm.

3.1 Parameters and Givens

In order to find a solution to this optimization problem, the following assumptions are made.

Facilities

The key variable in this optimization problem is the set of facilities on which the optimization will take effect. We are given a set of i facilities with Advanced Metering Infrastructure, in a specific power region. The assumption is made that all facilities are supervised by the same power utility.

Demand Forecasts

It is also important that we have forecasts of baseline demand for each facility. This allows us to effectively predict both when Demand Response should be used to obtain a load reduction, as well as providing us with baseline reduction information. These forecasts are built off of a current industry-tested Generalized Linear Model, provided by the facility where this research is being done. These forecasts report expected value.

Reduction Forecasts

In addition to the demand forecasts, it is also necessary that we have forecasts of the predicted reduction, when Demand Response is in effect. These forecasts are also built off of a current industry-tested Generalized Linear Model, provided by the same research facility. These forecasts also report expected value.

Pricing Model

It is also imperative that we have a cost model to compute the maximization clause of the optimization. The following pricing model is modeled after existing Critical Peak Pricing models [12] which are already used in Demand Response programs across the country.

3.1.1 Pricing Model

A critical-peak pricing model (CPP) for electricity demand works by charging the consumer a higher(peak) price for electricity, in order to incentivize a reduction in demand. CPP models are used by utility programs across the country for Demand Response and time of use (TOU) electricity programs [14].

When we optimize Demand Response, we will be lookig for a subset of the total facilities to enact critical-peak pricing (enact Demand Response) on, in order to achive the required load reduction. For the purposes of this paper, enacting Demand Response will be synonymous with enacting critical peak-pricing.

To do this, we introduce two pricing variables, P_{DR} and P_0, where P_{DR} is the critical price of electricity per Kilowatt-Hour (KWH) during Demand Response, and P_0 is the baseline price of electricity per KWH. We also introduce two consumption variables, L_i, and L_i^{DR}, where L_i is the forecasted electricity consumption without CPP, and L_i^{DR} is the forecasted consumption with Demand Response, at facility i.

An example pricing model for three consumers is as follows:

Consumer	KWH without CPP	KWH with CPP	Cost without CPP	Cost with CPP
C_1	L_1	L_1^{DR}	$L_1 P_0$	$L_1^{DR} P_{DR}$
C_2	L_2	L_2^{DR}	$L_2 P_0$	$L_2^{DR} P_{DR}$
C_3	L_3	L_3^{DR}	$L_3 P_0$	$L_3^{DR} P_{DR}$

This brings us the basic problem at hand; how do we optimize Demand Response in a continuous fashion such that the desired load cap is always met, and the most money is saved (profit is maximized).

3.2 Simple Discrete 2-Price Demand Response Optimization

3.2.1 Problem Statement

Let us take a set of i facilities with Advanced Metering Infrastructure. Each consumer i has a forecasted baseline electricity consumption L_i (in KWH), and a forecasted critical-peak pricing consumption L_i^{DR}. There is a load cap L_{MAX} (in KWH). Note that this optimization is only true for a certain interval in time, rather than a continuous period of time.

To set up this problem, we need to establish the constraints, as well as the function we are trying to maximize. So far, the only constraint to this optimization is that our total load consumption after enacting Demand Response should meet the required load cap, L_{MAX}.

With this load cap, the optimization is as follows:

$$\begin{aligned} \max_i \quad & \sum_{i_{DR}} P_{DR} L_i^{DR} + \sum_{i_{BASE}} P_0 L_i \\ \text{such that} \quad & \sum_{i_{DR}} L_i^{DR} + \sum_{i_{BASE}} L_i < L_{MAX} \end{aligned}$$

Or: Find the subset of facilities out of i on which to enact critical-peak pricing (Demand Response) in order to both maximize profit, and meet the required load cap L_{MAX}.

3.2.2 The Model

At its most general, this optimization is a simple profit maximization. As the problem is stated, we only need to find a set of facilities on which to enact Demand Response so that the reduction is met, and profit is maximized. Because P_{DR} is a higher price than P_0, it logically follows that Demand Response should be enacted on facilities where the difference $L^{DR} - L$ is small, or where the resulting load is still fairly high. In doing so, the utility companies will be maximizing their profit.

Therefore, the solution is as follows:

1. Compute the difference $L^{DR} - L$ for each of the facilities
2. Sort this list from lowest difference to highest difference
3. Pick the first n items such that $\sum_n L_n^{DR} < L_{MAX}$

The solution, even on scale, is fairly efficient, and can be easily modified to deal with large amounts of data. However, while there may not be problems with the solution, there are problems with how the model is structured.

The first, and perhaps the largest, flaw with the current model is that it makes the assumption that utilities can randomly pick facilities on which to enact a different price of electricity. The key problem with this is that this quickly becomes unfair, as some people are charged more for the same goods as others.

The second flaw with the model is that it only encompasses a single point in time. As mentioned previously, in order to truly optimize Demand Response, a continuous model that takes all time into account is necessary, in order to account for all cases. Therefore, it seems necessary that we create a model that takes into account the continuous, real-time nature of the optimization, and find a solution that fairly treats all members of the Demand Response program.

3.3 Continuous 2-Price Demand Response Optimization

Whereas in the previous problem, we dealt with a single interval in time, we will now be dealing with the electricity demand over an interval. Instead of a single electricity demand value in KWH, we will now be looking at a time-series load curve, or load profile.

Figure 2 shows the load profile of a power region over the course of a single day. Notice the high peak value on the graph. In order to prepare for demand at any given time, utility companies supply the peak amount of electricity throughout the course of the day. Not only is this inefficient, but it also results in a massive waste of electricity, as much of the supply goes unused for most of the day.

Instead, what I propose is the following: Instead of having electricity demand spike at a certain point in the day, there should exist a maximum load cap, L_{MAX} such that at any point in the day, the load profile will not exceed the maximum value. This will result in a much flatter load profile, with both less wasted electricity, and less electricity generated

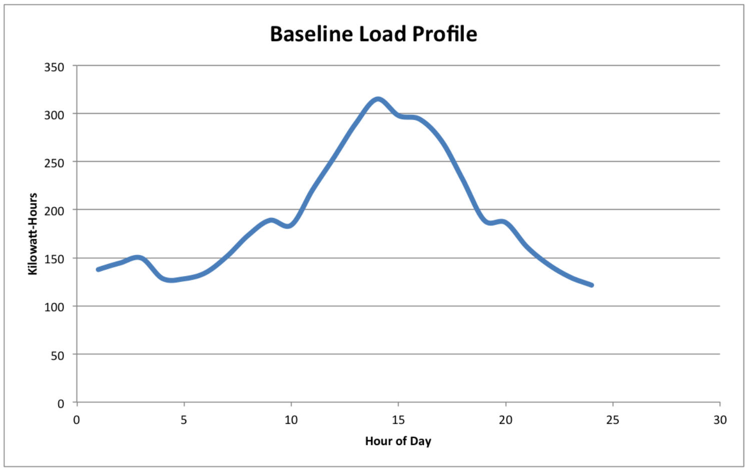

Figure 2: Typical baseline load profile for a power region. Notice the peak during the mid-day hours, and the sinusoidal nature of the graph.

overall, eliminating the need for peaker plants.

In order to do this however, it is necessary to find a way to use Demand Response in a manner that will both bring down overall electricity demand at the peak, but also maintain a relatively low electricity demand over the course of a day. To do this, I propose the following solution:

$$\max_{a_{i,j}} \quad \sum_i R_i$$

$$\text{where} \quad R_i = \sum_j p_{i,j} L_{i,j}$$

$$L_{i,j} = a_{i,10:00} L^{DR}_{i,10:00} + (1 - a_{i,10:00}) L_{i,10:00} \cdots + a_{i,17:00} L^{DR}_{i,17:00} + (1 - a_{i,17:00}) L_{i,17:00}$$

$$p_{i,j} = a_{i,10:00} p^{DR}_{i,10:00} + (1 - a_{i,10:00}) p_{i,10:00} \cdots + a_{i,17:00} p^{DR}_{i,17:00} + (1 - a_{i,17:00}) p_{i,17:00}$$

$$\text{such that} \quad \sum_j L_{i,j} < L_{MAX} \text{ for all } j$$

$$\sum_j a_{i,j} = 1$$

$$a_{i,j} \in 0, 1$$

The overall premise is very simple. Each household/facility i will have Demand Response called on it for a single hour, over the course of a day. Not only does this ensure fairness across all participants, but it provides the optimization with a certain degree of flexibility. However, as Demand Response is only effective when electricity consumption is fairly high, the optimization is limited to the hours of 10 A.M. (10:00) to 5 P.M. (17:00), when demand is highest. The forecasted load (dependent on the price) is represented by $L_{i,j}$, where j signifies the time in question. The term $a_{i,j}$ and its respective constraints allow the optimization to choose only one of the possible loads for each time interval (i.e. for 10 A.M., there will only be one load $L_{i,j}$). The constraint equation $\sum_j a_{i,j} = 1$ ensures that only one of the loads for each facility i will be the Demand Response load.

Alternately, the hourly price at time j is represented by $p_{i,j}$. In much the same way, the $a_{i,j}$ term and its constraints guarantees that only one of the interval's will be charged at the Demand Response price, and the rest at the baseline price.

Therefore, we are left with what seems to be a profit maximization problem, with the constraint that $L_{i,j} < L_{MAX}$ for all j. However, there is a flaw with this expression of the algorithm. With the constraint $a_{i,j} \in 0, 1$, this optimization turns into a integer quadratic problem, which is by definition NP-Hard (exponential Order of Growth). Especially when dealing with an average power region, with $n > 100$ independently controlled facilities, an optimal solution does not seem feasible, scalable, or time-sensitive. Therefore, I propose the following relaxation:

$$
\begin{aligned}
\max_{a_{i,j}} \quad & \sum_i R_i \\
\text{where} \quad & R_i = \sum_j p_{i,j} L_{i,j} \\
\text{such that} \quad & \sum_j L_{i,j} < L_{MAX} \text{ for all } j \\
& \sum_j a_{i,j} \leq 1 \\
& a_{i,j} \in 0, 1
\end{aligned}
$$

Instead of finding the optimal solution to this problem, this new algorithm relaxes the problem into an integer linear program. In this relaxed optimization, the equation $a_{i,j} = 1$ has been replaced with the equation $a_{i,j} \leq 1$. In the previous statement of the optimization, the purpose of this equation was simply to ensure that Demand Response was called on each facility once per day. As a result of changing this constraint, there exists the possibility that there are facilities on which Demand Response is not necessary. However, as the Demand Response load is always less than the predicted baseline load, utility companies can arbitrarily assign the remaining set of facilities to different demand response periods, without compromising the results of the optimization.

In order to solve this new, relaxed optimization, one of the more efficient options is with a greedy approximation. A greedy algorithm, while suboptimal, is both extremely efficient and readily scalable, both of which are crucial when dealing with a problem such as this. Furthermore, the adaptation of a greedy algorithm gets us fairly close to an optimal answer, with minimal effort. On the contrary, a brute force approach for a problem of this scale would take eons to find the optimal solution, and slightly less time to find a single solution that meets the given constraints.

The greedy algorithm is structured as follows:

1. The first step is to sort the time intervals by maximum demand. The premise behind this greedy algorithm is to drop electricity demand below the supply threshold over all time intervals j by first invoking Demand Response during times of peak demand, then slowly spreading out through the remaining intervals.

2. Now, the algorithm becomes a simple single interval optimization problem, with the same exact solving technique as before. This ensures that the overall load remains below the supply threshold, while maximizing profit.

3. Compute the difference $L_{i,j}^{DR} - L_{i,j}$ for each of the facilities separately. This value

represents the difference between the load during Demand Response, and the Baseline load. A larger difference signifies a larger drop in demand.

4. Sort this list from lowest difference to highest difference.

5. Pick the first x facilites such that $\sum_i L_{i,j}^{DR} < L_{MAX}$ for the given interval j. As a result of this step, Demand Response is called on enough facilities to drop demand below the threshold, as well as to cut down on electricity demand as a whole, while ensuring profit is at its maximum.

6. Continue to iterate through the remaining time intervals, until either there are no longer facilities available on which Demand Response can be called, or there are no longer any intervals for which Demand Response is necessary.

Because the algorithm prioritizes the intervals during which a shortage is most likely to occur, any deficit in available facilities will have a minute impact on the overall profile, especially with larger power regions. In the case that Demand Response is no longer necessary, utilities can pick random intervals on which to trigger events for the remaining households, as at any given time, $L_{i,j}^{DR} \leq L_{i,j}$.

By ensuring that the forecasted electricity demand (load) never exceeds a certain threshold (L_{MAX}), and Demand Response is called enough to substantially decrease overall electricity demand, this solution proves effective. Not only are eliminating the need for peaker plants, but we are also saving electricity utility companies millions each year, without adversely affecting thousands of people.

4 Results and Discussion

In order to test the results, the data used is from an actual United States power region, under the following conditions:

1. The raw data encompasses 100 facilities ($n = 100$) over the course of a year.

2. This data was used to train the Forecasting GLM, which was used to produce both the baseline and demand response forecasts for each facility. For the purposes of this research, each forecast is limited to a specific day (same day for each facility forecast).

3. The timeseries data for each facility exists, and is forecasted on, a 15-minute granularity. This means that for each facility, there will be a forecast of 96 data points for the day ($24 * 4 = 96$).

4. Demand Response can only be called during the hours of 10 A.M. to 5 P.M. (the peak hours).

5. When Demand Response is called on a facility, the interval will be 1 hour (i.e. a Demand Response event called at 11 A.M. will last until 12 P.M.).

6. The two prices used will be the Standard price of electricity per Kilowatt-Hour, and the Critical (Demand Response) price of electricity per Kilowatt-Hour, both of which are set by the utility company.

7. The maximum peak load (L_{MAX}) is equivalent to utility peak supply. For our purposes, L_{MAX} will be 75 KWH at any given 15-minute interval.

Before taking a look at the results of the algorithm, let us first create a basis for comparison. Figure 3 shows the baseline forecast for the test dataset, in which data is represented on a 15-minute granularity.

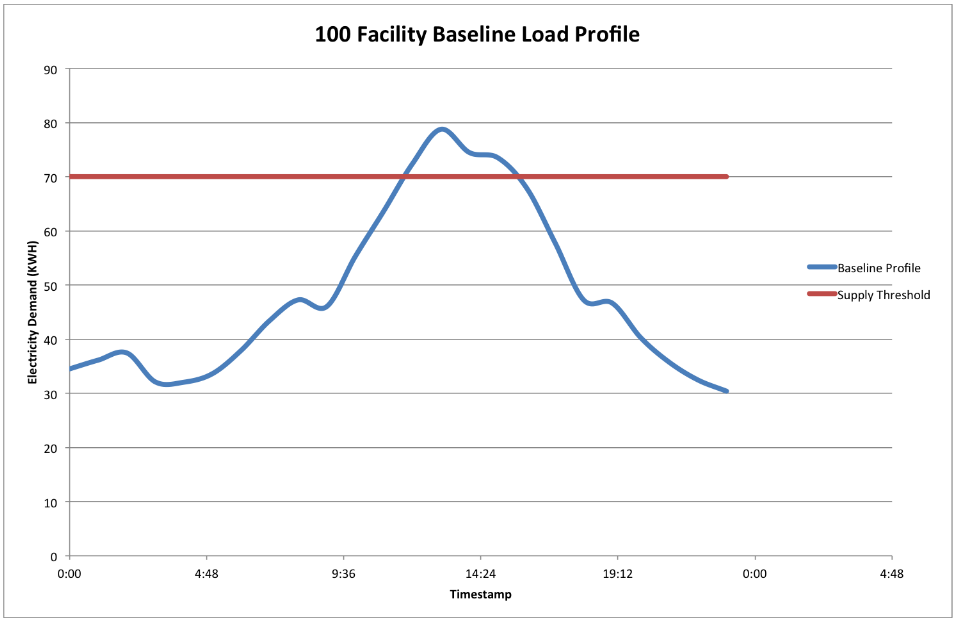

Figure 3: Baseline load profile (15-min granularity) for the test dataset. Notice how the Baseline Profile exceeds the Supply Threshold during the mid-day peak.

Not only does this baseline profile show the typical sinusoidal nature of electricity demand, but it also has a mid-day peak that exceeds the supply threshold (represented by the red line). Because the demand curve exceeds the supply threshold of 70 KWH, a shortage of electricity is created, a problem that can only be resolved by turning on a peaker-plant, and generating additional electricity on the spot. However, as our goal is to eliminate the necessity for peaker-plants, this is the perfect dataset to test the optimization algorithm.

In order to be a true solution to this problem, the end result (after the optimization) should consist of a profile somewhat similar to the baseline forecast, but meeting the supply threshold at any point along the line. In essence, the peak and overall electricity consumption of the optimized profile should be slightly lower than the baseline: however, since there are only 100 facilities in question, there should only be a minute difference.

Figure 4 shows the post-optimization profile, in comparison to the baseline profile, and the supply threshold:

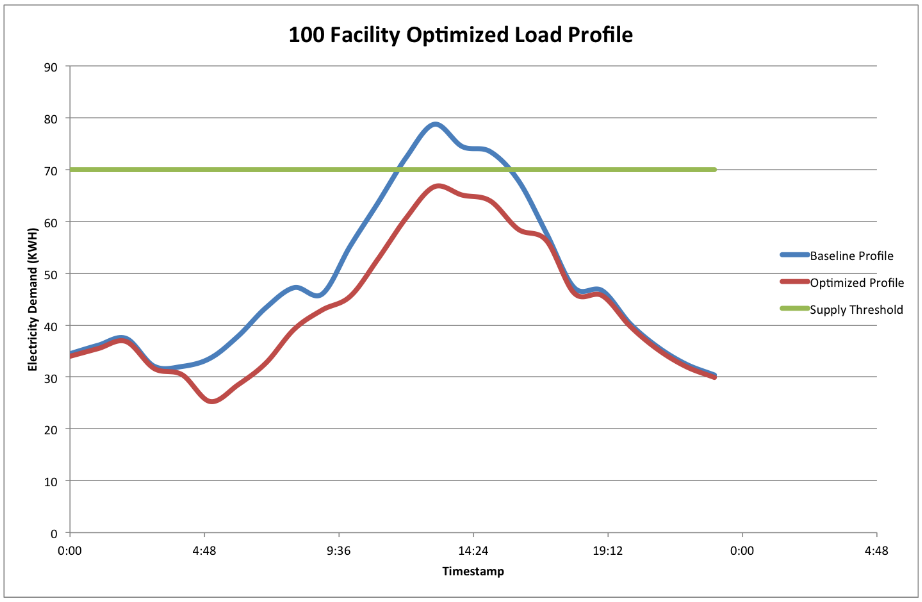

Figure 4: Comparison of the optimized profile to the baseline profile. Note the flatter nature of the optimized demand curve, as well as the adherence to the supply threshold. Also note the amount of reduction in total demand, between the optimized profile and the baseline profile.

In addition to the peak load reduction, the optimization profile also seems to be reducing demand as a whole, across the entire time series. By taking the numerical integral across the entire time series, the following breakdown of overall and peak demand is produced, providing a way in which to numerically compare the data from the optimization with the baseline.

Load Profile	Total Demand (Consumption)	Peak (Maximum) Demand (Consumption)
Baseline Profile	4627.18 KWH	78.75 KWH
Optimized Profile	4142.42 KWH	66.70 KWH
Percent Change	10.47%	15.3%

While 10.47% total reduction, and 15.3% peak reduction seem like small numbers, in truth, they are actually slightly misleading in nature. Recall that the current data is limited to $n = 100$ facilities over the course of a single day, in a small power region. The average power region has $n > 1000$ facilities, with certain larger regions reaching into the $10,000s$.

While the reduction may scale slightly sublinearly as the power region becomes larger, a 10% reduction in a power region with $n = 5000$ facilities contributes to gigawatts of electricity saved each day. Furthermore, if the peak reduction is also scaled upwards to a power region of such size, there will generally be much less electricity generated at any given time. As a result, the overall supply threshold can possibly be decreased, therefore both saving utility companies generation costs, and saving the total amount of electricity generated total.

One of the benefits of using a greedy algorithm to solve this problem is that the overall time to find a solution is order of magnitudes smaller. This means the algorithm can be used in a real-time manner, to quickly optimize over a set of facilities, even minutes before the start of the first Demand Response interval. In this way, this solution is much more flexible than existing methods, and as a result, more effective. Furthermore, in addition to providing utilities with a real-time algorithm for optimization, this greedy algorithm also allows for a high degree of scalability. As most power regions contain thousands of independent facilities, a scalable solution such as this one is not only highly beneficial, but necessary.

4.0.1 Extension to Outage Prevention

There is not much that changes with the solution when applied to outage prevention in developing nations. On a macro level, the cost of failure is higher, as a large-scale power outage is more detrimental in the long-term than slightly expensive generation costs.

However, one of the more common problems with developing nations regards the stability of their grid systems. Data received from the smart grid in developing countries may not be as reliable, or as accurate as is necessary (due to problems like electricity theft, malfunctioning AMI, etc.). Instead of ensuring that load at each interval is less than the supply threshold L_{MAX}, the new optimization should use a new value $L_{MAX} - \epsilon$. ϵ is an independent constant, and it provides the utility company with a certain amount of wiggle room, such that the supply threshold will definitely be met.

5 Conclusions and Future Work

While this solution is an optimistic start to a relatively untouched problem, there still exist ways to improve upon it. The next logical step after this solution would be to actually implement the optimization algorithm on a large scale (i.e. $n > 10000$), and see if there exist reasonable results. Furthermore, there are a multitude of different methods through which one could find a more optimal solution to this problem, rather than one that just meets the constraints. One method would be to use a sparse approximation (subset selection) approach to solving this problem: that is to say, find a way to relax this optimization problem into a convex optimization, to find a close-to-optimal solution in little time. A method such as this one would result in an efficiency rate much higher than the current greedy algorithm, and would also ensure a high degree of scalability.

This research began with an explanation of the current problem with Demand Response, and the necessary steps needed to take to find an optimal solution. Current Demand Response is too static and coarse-grained to provide an effective, flexible solution for controlling electricity demand. To develop a fine-grained, continuous model for Demand Response, this research first looked at a simpler subproblem, the optimization of Demand Response over a single interval. Then, the optimization statement evolved into an Integer Quadratic (NP-Hard) optimization problem. Next, through justified relaxations, the problem was reduced to an optimization that can be solved with an efficient, real-time, and scalable greedy algorithm.

As a result of this greedy algorithm, utility companies in the United States are able to eliminate the need for these expensive and unnecessary peaker-plants. Both the utilities and the general public can save millions of dollars each year, at a low, almost non-existent cost. Furthermore, the application of this solution to outage prevention in developing nations proves its flexibility and effectiveness, placing utilities in a position of preventing large-scale blackouts altogether.

References

[1] Katherine Tweed. *Green Tech Media*, July 2013.

[2] David Silva. Digging deeper: 'peaker' plant under construction in city of industry. *ENR California*, April 2012.

[3] Iordanis Koutsopoulos and Leandros Tassiulas. Control and Optimization Meet the Smart Power Grid: Scheduling of Power Demands for Optimal Energy Management. *CoRR*, abs/1008.3614, August 2010.

[4] Haitham Hindi, Daniel Greene, and Caitlin Laventall. Coordinating Regulation and Demand Response in Electric Power Grids using Multirate Model Predictive Control. *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*, pages 1–8, January 2011.

[5] A. I. Cohen and C. C. Wang. An optimization method for load management scheduling. *IEEE Transactions on Power Systems*, 3(2):612–618, 1998.

[6] Y. Y. Hsu and C. C. Su. Dispatch of direct load control using dynamic programming. *IEEE Transactions on Power Systems*, 6(3):1056–1061, 1991.

[7] D. C. Wei and N. Chen. Air conditioner direct load control by multi-pass dynamic programming. *IEEE Transactions on Power Systems*, 10(1):307–313, 1995.

[8] J. Chen, F. N. Lee, A. M. Breipohl, and R. Adapa. Scheduling direct load control to minimize system operation cost. *IEEE Transactions on Power Systems*, 10(4):1994–2001, 1995.

[9] K. H. Ng and G. B. Sheble. Direct load control a profit-based load management using linear programming. *IEEE Transactions on Power Systems*, 13(2):688–695, 1998.

[10] Lijun Chen, Na Li, Libin Jiang, and Steven H. Low. Optimal Demand Response: Problem Formulation and Deterministic Case. *Power Electronics and Power Systems*, 3:63–85, 2012.

[11] Advanced metering infrastructure (ami). Website. `http://www.ferc.gov/eventcalendar/Files/20070423091846-EPRI%20-%20Advanced%20Metering.pdf`.

[12] Electric power annual. Website. `http://www.eia.gov/electricity/annual/`.

[13] Cruz E. Borges, Yoseba K. Penya, and Ivan Fernandez. Evaluating combined load forecasting in large power systems and smart grids. *IEEE Transactions on Industrial Informatics*, PP(99), September 2012.

[14] Karen Herter and Seth Wayland. Residential response to critical-peak pricing of electricity: California evidence. *Energy*, 2009.