Optimisation of Demand-Side Bidding

Jason Alfons Undan
Department of Engineering Science
University of Auckland
New Zealand
jund013@aucklanduni.ac.nz

Abstract

We consider the problem of formulating an optimal bid curve (demand function) for a large consumer in an electricity market aiming to maximise its profits. A dynamic programming approach was suggested in this paper to construct the optimal bid strategy.

The aim of this approach is to use the consumer's profit function to approximate the demand needed in order to maximise its returns. This method could potentially try to reduce consumption of electricity and increase the efficiency of the market by having more active participations from the demand-side. We present numerical results using artificial offer-curves and real-offer curves (obtained from New Zealand Electricity Market for a large consumer).

Key words: Dynamic programming, bidding strategy, revenue function, electricity market.

1 Introduction

Electricity is an important element for large consumers because it is an input for many production and manufacturing processes. As significant amount of electricity are required by these users, it normally becomes the major cost after raw materials. Let us consider an aluminium smelter example where q units of electricity are needed to produce 1 tonne of aluminium. The 1 tonne of aluminium is then sold at market price and the company collects \$ t per tonne of aluminium. Total gross profit gained by the company in producing and selling 1 tonne of aluminium is the revenue of selling 1 tonne of aluminium less the cost of purchasing q units of electricity to produce the 1 tonne of aluminium. This suggests that the revenue generated by the aluminium smelter depends on the amount of electricity units, q purchased to produce a certain amount of aluminium.

Thus, the cost of electricity has direct impact on the revenue of these electricity users. If the price of electricity at a particular trading period is too high that it becomes not profitable to operate, the large consumer has the option to shut down or continue with their production. In our problem, we consider large consumers to have the flexibility to reduce their consumption depending on the spot prices. These consumers can remain shut down, with a cost associated to that decision, if the price of electricity is high that to operate is not a feasible option.

The aim of this project is to develop a methodology for a purchaser to bid for electricity optimally (i.e. submit an optimal demand stack or bid curve to the centralised system) as to maximize the expected profit from purchasing and using certain amount of electricity units, q. We present a dynamic programming approach to solve the optimal bidding problem for the purchaser based on the revenue generated from the purchase of electricity units. In other words, the revenue of the purchaser applied into our method is a function of electricity units purchased. A previous work had been done by Pritchard (2007) to formulate an optimal offer curve for an electric power generator with market power. What this project is trying to do is to develop the opposite formulation for a large consumer to formulate their optimal bid curve.

We developed the method using simple, one-offer curve and three-offer curves examples which can easily be solved analytically. We then used the methodology to solve problems using realistic offer curves obtained from the New Zealand Electricity Market (NZEM) to evaluate its potential of solving real-world problems.

The NZEM has been chosen to explore our methodology in this project. However, other electricity markets may as well be an appropriate market for our methodology. In New Zealand, the term “electricity market” typically refers to the wholesale market where demand every half hour at every node in the network is satisfied at the lowest possible price based on the offers of the generators. In the wholesale market, electricity is purchased by retailers and large consumers.

We describe the dynamic programming formulation and algorithm in Chapter 2 and present the one-offer curve and three-offer curves examples and real-offer curves from NZEM for consideration. For the real-offer curves problem, we assume to be one of the large consumers in New Zealand, i.e. The New Zealand Aluminium Smelter. We analyse and discuss the results in Chapter 3 and finally in Chapter 4 we draw conclusions from this project.

To illustrate an option that a large consumer can consider rather than shutting down completely, supposed in a single node market, where all other demands have been satisfied. There is a residual generation curve and as a large consumer who has a demand that has not been dispatched, the question faced is how to bid optimally in order to maximise the purchaser’s return.

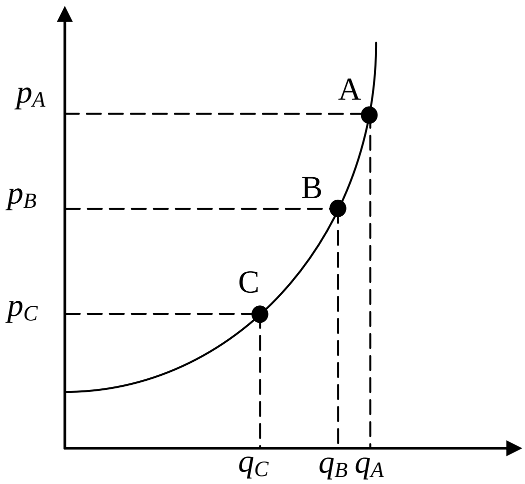

Figure 1.1: An offer stack that has three possible points of dispatch that can be considered by a purchaser

In the diagram above, there are three possible points of dispatch that the large consumer can consider. The initial demand is point A. A natural way of bidding is to bid at A as to get the desired demand of electricity. However, if the consumer can be flexible enough to reduce consumption, bidding at pB is better considering that the price at B is less than A for a relatively small reduction in consumption. Consequently, the cost of electricity can be reduced quite significantly with the decision to be dispatched at reduced consumption.

An optimal point of dispatch cannot be guaranteed unless the actual revenue function of the company is known. Perhaps being dispatched at C may be a good decision too, however it requires greater amounts of consumption to be reduced. The large consumer needs to consider the economic benefits of being dispatched at a reduced consumption, or not being dispatched at all.

Demand-side management proves to provide benefits in electricity markets especially in making the whole system more efficient and effective (Electricity Commission, 2009). However, different market structures and norms pose themselves as challenges for the demand-side management potentials. It will require many consultations with stakeholders, investment capitals and time before a decent demand-side policy will take off effectively in an electricity market. Nevertheless, it is hoped that our method will be able to contribute to more active demand-side participation into the electricity market taking into consideration its unique structures and constraints.

2 Methodology

2.1 Dynamic Programming

Dynamic Programming is an optimisation methodology developed in the 1940s by Richard Bellman to describe the process of solving problems where one needs to find the best decision one after another. The main concept of this technique lies in the principle of optimality which can be stated as follows:

An optimal policy has the property that whatever the initial state and the initial decision are; the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision (Luus, 2000).

Generally, today, it is known as an approach to solving complex problems by breaking them into simpler steps or structures. To put the method into context, we have a complex electricity bidding problem faced by a consumer. Given an offer curve consisting of few tranches, we are interested in knowing the best way to bid for electricity, possibly submitting a bid curve that will maximise our expected profits as consumers. This approach can be done by first subdividing the quantity-price (q, p) plane to form smaller structures or sub problems of the main problem.

2.1.1 Discretisation of (q, p) plane

We begin by subdividing the (q, p) plane with a finite grid of M by N rectangular cells. Any class of admissible bid curves are restricted to follow only the edges of this grid (i.e. a bid curve must intersect each cell of the grid only in its boundary), and are monotone decreasing. Consequently, there are only a finite amount of admissible bid curves, each consisting of a finite sequence of exactly M + N horizontal and vertical line segments (grid edges).

A key observation is that for a fixed grid edge e, the occurrence of a point of dispatch on e is independent of what other grid edges are included in the bid curve. Therefore, for each grid edge that intersects with the offer curves (point of dispatch occurring at e), there is an expected profit V (e) associated, which will be realised if and only if e is part of the bid curve. The expected profit of an admissible bid curve r is then

$$f(r) = \sum_{e \in r} V(e)$$

2.1.2 Edge Value and Revenue Function, $\boldsymbol{\rho}$

Each grid edge, e in the discretised (q, p) plane has an edge value. It is the expected profit that a consumer will receive by purchasing q units of electricity at the price of p, if and only if the grid edge, e is included in the bid curve and point of dispatch occurs at e. To calculate the expected profit for each grid edge in the plane, we need to have a revenue function, ρ for the problem. This can be done via an equation such as $f(e) = \rho(q) - C(q)$; expected profit at any given grid edge is the revenue from using q units of electricity less the payoff (cost of electricity) to the generators when purchasing q units of electricity. Our problem is to find the maximum of f over its (finite) domain.

2.1.3 Problem Formulation

Suppose our grid covers the region $q_{min} \leq q \leq q_{max}$, $p_{min} \leq p \leq p_{max}$, and that all possible points of dispatch lie within this region. For each vertex x of the grid, let W(x) denote the maximal expected profit, due to points of dispatch above x, of any bid curve which passes through x.

Then

$$W(x) = \max\ \big(W_l(x), W_u(x)\big),$$

where

$$W_l(x) = \begin{cases} -\infty, & \text{if } x \text{ is on } q = q_{min} \\ V\big(e_l(x)\big) + W\big(v_l(x)\big), & \text{otherwise,} \end{cases}$$

and $e_l(x)$, $v_l(x)$ are respectively, the vertex adjacent to x to the left, and the edge linking x to that neighbour.

Similarly,

$$W_u(x) = \begin{cases} -\infty, & \text{if } x \text{ is on } p = p_{max} \\ V\big(e_u(x)\big) + W\big(v_u(x)\big), & \text{otherwise,} \end{cases}$$

where $e_u(x)$, $v_u(x)$ are respectively, the vertex adjacent to x above, and the edge linking x to that neighbour.

We have $W(x) = \big((q_{min},\ p_{max})\big) = 0$. It is thus straightforward to successfully evaluate W(x) for each vertex x of the grid, and hence determine the optimal admissible bid curve. This procedure clearly yields a global optimum for the discretised problem which can also be taken as an approximation of the optimal bid curve for the original problem.

2.1.4 Estimation of Optimal Expected Profit

There are two main factors that affect the optimal expected profit of an admissible bid curve for the discretised problem, i.e. size of the grid and the revenue function that is put into the problem.

The size of the grid refers to the vertical height and horizontal width of a particular rectangular cell in the (q, p) plane after being subdivided into M by N cells. In theory, we will expect finer grids to give better optimal expected profits. However as the plane can be divided into any arbitrary values, it can be expected that the optimal expected profit will have an upper bound that defines the maximal limit the profit can be, provided the revenue function is kept constant.

Revenue function is important in computing the expected profit in the discretised problem as the calculation is done at each intersection of the offer curves with the grid to compute the edge values. An accurate revenue function will likely to give an acceptable optimal bid curve for the consumer because our problem is to maximise the sum of edge values in the grid so as to form a monotone decreasing bid curve. A poor revenue function will most probably give an ineffective bid curve and will not likely make any meaningful decision for the consumer.

2.2 One-Offer and Three-Offer Curves Examples

To verify the effectiveness of the dynamic programming algorithm, we use it to solve simple one-offer and three-offer curves examples. Given a (q, p) plane with quantity and prices ranging from $0 \le q \le 40$ and $0.0 \le p \le 4.0$ respectively. We let the revenue function at any point on the grids be $\rho(q) = 5q$. We subdivide the plane into four by four rectangular cells. Table 3.1 and table 3.2 shows the offer stacks and aggregated offer stack submitted for the one-offer curve and three-offer curves examples respectively.

Offer Curve A		
Price ($)	MWh	Cumulative MWh
0.00	5	5
0.04	12	17
1.50	8	25
2.80	3	28

Table 2.1: One-offer curve example

Offer Curve A			Offer Curve B			Offer Curve C		
Price ($)	MWh	Cumulative MWh	Price ($)	MWh	Cumulative MWh	Price ($)	MWh	Cumulative MWh
0.00	5	5	0.00	5	5	0.00	2	2
0.04	12	17	0.05	18	23	0.06	11	13
1.50	8	25	0.07	3	26	1.70	11	24
2.80	3	28	3.20	3	29	2.40	5	29

Table 2.2: Three-offer curves example

We repeat the example using a smaller grid size plane of fifteen by fifteen rectangular grid cells and a different revenue function, $\rho(q) = 6\sqrt{9q}$ is also considered using both grid sizes to provide comparisons between different scenarios.

2.4 Real-Offer Curves from NZEM

After looking into the two examples above, we proceed to use our method to solve a realistic problem using real historic offer curves from the NZEM. As an example to illustrate the potential application of our method specifically for large consumers, we considered ourselves as being the New Zealand Aluminium Smelter at Tiwai Point with our grid exit point (GXP) as TWI2201. The data are obtained from M-Co Ltd website.

The period that was taken as samples for our method is a 5-day week both in summer and winter of 2008, i.e. 11-15 February 2008 and 7-11 July 2008 respectively. However, not all of the 48 periods in a day were considered. We focused on a part of the day (weekday daytime) from 6.30 a.m. to 5.00 p.m. (i.e. period 13 to 34) which are treated as “equivalent periods” which was similarly done by Pritchard and Zakeri (2003). The multiple offer stacks obtained from these periods provide us with a distribution of offer curves that we believe to be similar to each other. These offer stacks present a general pattern that suggests a distribution of an offer curve. For our project, the estimation of the offer stacks distribution is not explored. However with our data, we can strongly believe, it represents an accurate variation of the offer stacks investigated.

2.4.1 Aggregated Residual Generation Curves

We have taken all the offer curves submitted by all generators in the NZEM and combined them to form aggregated generation offer. We also obtained all the demands except the demand at TWI2201 for the periods and aggregated them to form aggregate demand curves. In theory, the NZEM works in such a way that the intersection between the aggregate generation offer and aggregate demand curve is the point of dispatch at a specific period of the day. Having satisfied all the demands except the demand at TWI2201, we have a residual generation curve which is the quantity that the market supplies that is not consumed by other demanders at any given price.

With the residual generation curve, as a large consumer, our method aims to find the best way to bid in for electricity in a way to maximise our expected profit. We have used a linear revenue function to evaluate our optimal expected profit which is $\rho(q) = 250q$. Variations on the grid sizes were also done to analyse the effectiveness of the methodology on a larger scale.

3 Results and Discussion

3.1 One-offer and Three-offer Curves Examples

For the one curve example, we created a four by four rectangular grid on the (q, p) plane. The solution is shown in Figure 3.1. The optimal expected profit from submitting the bid curve is $75.00 (based on a revenue function, $\rho(q) = 5q$). The same one curve example is solved using a finer set of grids of fifteen by fifteen rectangular cells. The optimal expected profit increases to $85.00. In both solutions, the offer stack intersects with the bid stack at the same quantity but different prices. As the bid curves are only allowed to follow the grid edges, a finer grid size will result in more meaningful and useful insights.

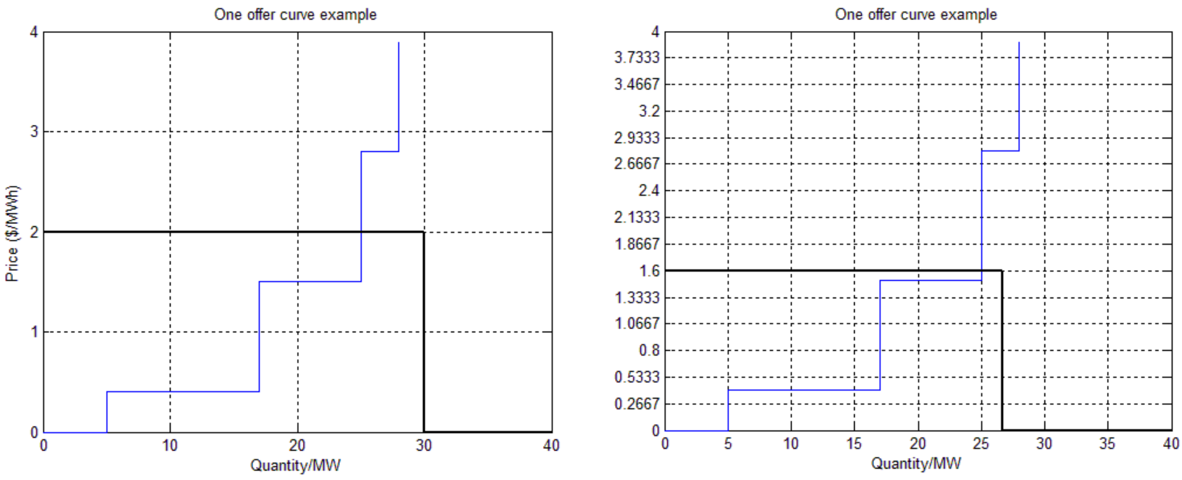

(a) (b)

Figure 3.1: Result for one offer stack example (a) 4 by 4 grid (b) 15 by 15 grid using a linear revenue function, $\rho(q) = 5q$

We also try solving the one curve example using a different revenue function, $\rho(q) = 6\sqrt{9q}$ on a four by four rectangular grid cells. The optimal expected profit from bidding the demand curve is \$57.22. We repeat the example with a fifteen by fifteen rectangular grid and the optimal expected profit increases to \$65.00.

For the three curves example, using a four by four rectangular grid yields an optimal expected profit of \$184.00, while a fifteen by fifteen rectangular grid on the same example gives an optimal expected profit of \$ 272.40. Notice the differences between both solutions in Figure 3.2 are in the setting of the bid price and quantity of demand.

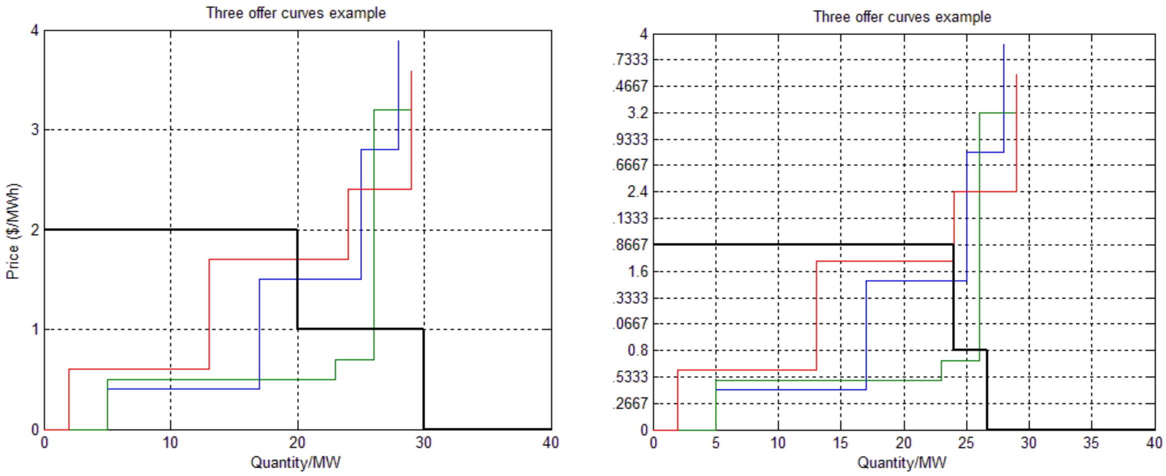

(a) (b)

Figure 3.2: Result for three offer stacks example (a) 4 by 4 grid (b) 15 by 15 grid using a linear revenue function, $\rho(q) = 5q$

The three-offer curves example is solved using a different revenue function, $\rho(q) = 6\sqrt{9q}$ to yield optimal expected profits of \$97.78 and \$193.71 on a four by four and a fifteen by fifteen rectangular grid respectively.

It can be observed that every time we decrease the size of the grid cells, the optimal expected profit tends to increase and the bid curve shows a slight variation. It is either the consumption that is bid becoming less, or the price that we bid electricity is lesser than the price at using the initial four by four rectangular grid.

The maximum optimal expected profit gained by bidding a particular bid curve is the accumulated profit instead of actual profit. This is because in the methodology, every intersection between the offer curves and grids are considered, but in reality only one point of intersection will be realized. Nevertheless, the curve constructed using backward recursion based on maximum optimal expected profit gives a clear indication that any point of dispatch that occurs on the curve will give the best expected profit compared to others.

3.2 Real-Offer Curves from NZEM

The following table show the result of the methodology applied using real offer curves obtained from NZEM. For each season (summer and winter), two refinements of the grid sizes was done in order to evaluate the effects of grid sizes on the maximal optimal expected profits a large consumer can get from a particular bid strategy. The revenue function that is used is $\rho(q) = 250q$.

Season	Initial Result	1st Refinement	2nd Refinement
Summer	\$ 2,469,100	\$ 3,298,800	\$ 3,617,300
Winter	\$ 1,522,000	\$ 1,537,800	\$ 1,543,600

Table 3.1: Summary of result for real-offer curves from NZEM

The initial bid curve gives a maximum expected profit of $ 2,469,100. As it can be observed the curve starts off at a high price of $180/MWh taking into account any generation below that price. However, the price per unit of electricity that the consumer should pay starts to decrease if 200MW or more is required. A further refinement of the grid size is done and this refinement yielded a maximum expected profit of $ 3,298,800; a 33.6% increase in expected profit.

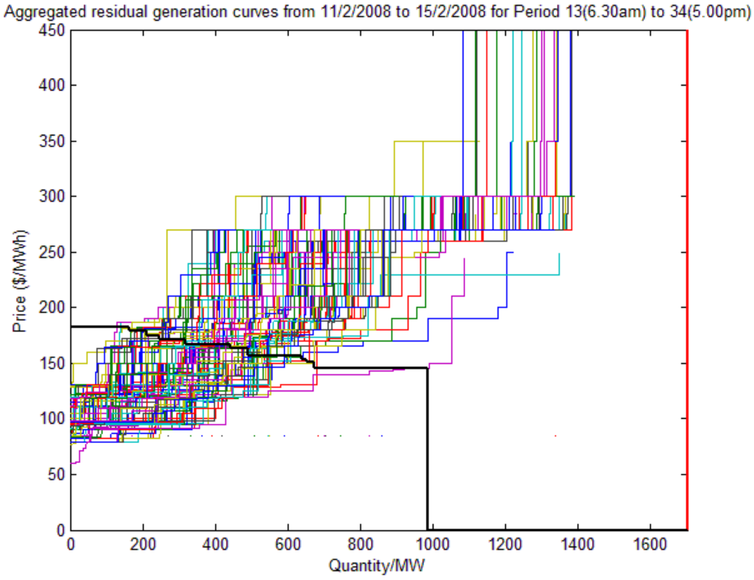

Figure 3.3: Result for real offer curves from NZEM using a linear revenue function

From the result, the difference that can be observed between the refinements is the starting price for the consumer in its bid curve is lower compared to the previous result. Instead of demanding up to approximately 1000MW, a considerable reduction in quantity of electricity demanded is noted (approximately 700MW). The price starts to decrease in a greater amount well before reaching 200MW. A further refinement of the size grids is done and the maximum expected profit found to be $ 3,617,300, which is a 9.6% increase from the previous result. There are not many differences between these two results other than the curve following different edge lines between 150MW and 600MW.

For the periods in winter of 2008, a similar approach is done to the data. The following Figure 3.5 shows the bid curve that a large consumer can bid in to maximise its return. The bidding strategy yields a maximal expected profit of $ 1,522,000. A refinement of the grid size gives a slight increase of 1% in the expected profit, i.e. $ 1,537,800.

A further refinement improves the expected profit to $ 1,543,600. It can be noted that very small improvement can be done onto the result due to the less variability in the winter curves compared to the summer curves; following particular edge lines are necessary in order to ensure maximum optimal expected profit is achieved.

The expected profits during the winter of 2008 are much less compared to the summer of the same year because the electricity generations become more expensive in winter with increasing demands, especially for heating purposes. Note that another reason for the increasing values of optimal expected profits is due to the dynamic programming methodology to bid much less consumption rather than expected actual demand. In particular, less demand means more profit can be made. Due to the finer grids in which the bid curves are confined to follow, more offer stacks are likely to intersect with these grids which lead to calculation of expected profit and accumulated values at each vertex. The dynamic programming method do not take demand into account but rather finds the best way (most optimal and profitable) way in which a purchase, in our case, the large consumer to bid in.

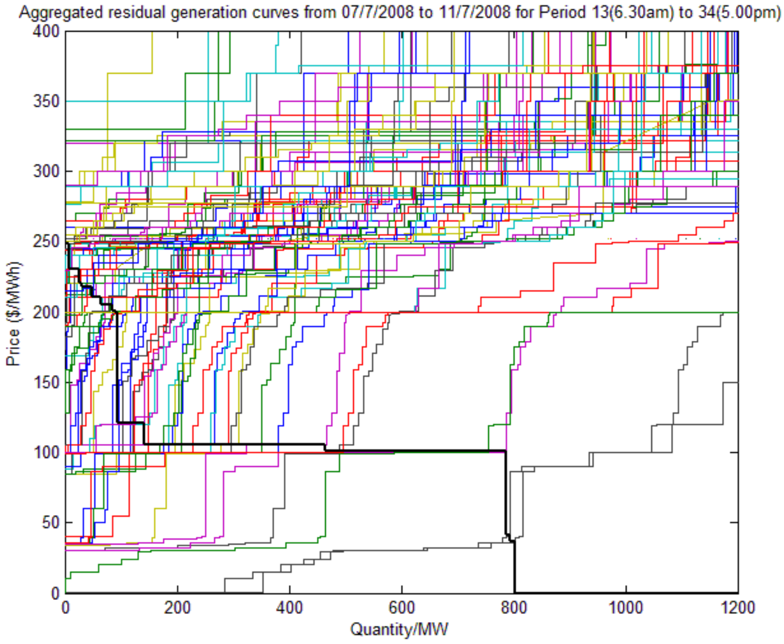

Figure 3.4: Result for real offer curves from NZEM using a linear revenue function with second improvement on grid sizes

This can become quite a significant limitation to the method as certain consumer, although we assume to be flexible, still requires some lower bound of electricity based on its demand to operate efficiently to generate profits. To operate below the identified lower bound, it may be comparable with shutting down for a particular trading period.

In reality, a large consumer may likely have a step-wise revenue function or even demand-related functions. However, as revenue functions differ significantly going from one consumer to another, the dynamic programming method is flexible in taking various classes of revenue functions to evaluate the expected profit at any point in the grid of the (q, p) plane. The consideration in submitting a revenue function into our method is that it should be scaled and forecasted accurately in order to obtain a meaningful and useful bidding strategy for the user.

4 Conclusions and Future Work

In this project, we have managed to develop a dynamic programming method for a large consumer to bid optimally into a given offer stack. This was done by discretising the (q, p) plane into M by N rectangular cells and only allow the bid curve to follow by the grid edges as to form a monotone decreasing bid stack. Each vertex in the grid is the maximal expected profit that can be obtained by a consumer if it was to bid a horizontal edge to the right or vertically downwards. The dynamic programming method then finds the best way of bidding into the offering problem by tracking back the maximal values of each vertex in the grid forming an optimal bid curve.

This method especially allows large consumers to access the potential or possibility to reduce consumption when the price of electricity is very high. A slight reduction in consumption can lead to significant cost savings for the consumer, thus impacting positively onto the grid system as a whole. It will also encourage an active participation on the demand-side of the market as consumers can manipulate the prices and bid quantity depending on the conditions of the market.

One limitation of the method is that it does not take the demand of the consumer into account. We assume consumers have the flexibility to reduce consumption, however in reality, the ability to do so is not easily achieved due to the regulations of the system and types of business a consumer is operating.

A business that has non-disruptable processes does not have as high a flexibility to shut down but it may consider reducing consumptions to avoid high spot prices. Our method gives the alternative to shutting down by allowing consumers to analyse what the total economic benefits are by reducing consumption by how much it can do so legitimately.

Another limitation of the method is that the grid is constructed over a huge (q, p) plane (i.e. maximum range of between 3000 - 6000MW) where in practice, the possibility of being dispatched for a large consumer lie at the maximum range between 450 – 700 MW. To construct a large grid where most of the spaces are not used presents itself as inefficiency in the computational technique. Multiple grid sizes in the (q, p) plane where focus is given in more likely dispatched zones is thus a possible area of future development.

More empirical studies can improve the implementation of the methodology. The data that was used to represent the offer curves distribution was quite a small sample. Perhaps using more data and applying the method using data from other electricity markets to see the variation in offer patterns is also a possible extension.

As presented in this project, the methodology was developed using single-node market examples. Another future work that can be considered is to extend the problem into a multiple-node electricity market setting with nodal pricing scheme implemented. By doing so, the effectiveness of the methodology to find an optimal bid curve, taking consideration of congestion constraints within a system network can be evaluated.

Acknowledgements

I would like to thank my project supervisor, Dr Golbon Zakeri from the Department of Engineering Science, University of Auckland for her guidance, patience and support in completing this project. Thanks to Tony Downward for his opinions and ideas that he had willingly contributed.

5 References

Electricity Commission. 2009. Dispatchable demand: options – Consultation Paper. September 2009.

Luus, Rein.2000. *Iterative Dynamic Programming*. Chapman and Hall/CRC.

M-Co Ltd. *Historic Orders*. [Online] Available: http://www.electricityinfo.co.nz/comitFta/Ongoing_bidoffer.ongoing [18 August 2010]

M-Co Ltd. *Demand*. [Online] Available: http://www.electricityinfo.co.nz/comitFta/ftaPage.demand [23 August 2010]

Pritchard G. 2007. Optimal offering in electric power networks. *Pacific Journal of Optimisation*, 3(3):425-438.

Pritchard G. and Zakeri G. 2003. Market offering strategies for hydroelectric generators. *Operations Research*, 51(4):601-612.