Developing bidding strategies for electric utilities in a competitive environment

by

Charles William Richter, Jr.

A thesis submitted to the graduate faculty

in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Department: Electrical and Computer Engineering

Major: Electrical Engineering

Major Professor: Gerald B. Sheblé

Iowa State University

Ames, Iowa

1996

Graduate College
Iowa State University

This is to certify that the Master's thesis of

Charles William Richter, Jr.

has met the thesis requirements of Iowa State University

Signatures have been redacted for privacy

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

Currently deregulation is changing the interaction of buyers and sellers in the electric power marketplace to promote a more efficient usage of the power system. In their push toward a more competitive marketplace, many proponents of the deregulation have proposed auction methods implemented via a regional commodity exchange as the new standard for buying and selling power. The author is building on a proposed environment in which generation companies (gencos) are separated from distribution companies (discos), and must communicate through the independent system operator (ISO) via the bidding process. In a competitive bidding environment such as the one proposed, an electric utility's profit will require a successful bidding strategy. This research describes evolving adaptive agents which model generation companies. As the agents evolve, those with more successful bidding strategies, i.e. those yielding more profit, survive over the generations. New agents are created using offshoots of the successful strategies (according to the rules of genetic algorithms) which are sometimes better than the original. The results obtained could be useful to electric utilities operating in a competitive environment. By including real generation cost curves and company specific information to the models used in this software, the method used this research could be used to simulate companies participating in an auction. By studying how bidding practices affect agent's profit in the computer simulation, those trade via auctions should be able to develop bidding strategies which will help them maximize their profit.

CHAPTER I: INTRODUCTION

Recent changes to regulations governing the electric utility industry are being implemented to make the electric system in the United States more efficient. Increasing competition through open access to the transmission network is expected to increase efficiencies and benefit electric consumers.

Other industries in the United States have recently been successfully deregulated, or re-regulated, to promote competition. Among these, the communication industry, the airline industry, and the natural gas industry have made the shift to more competitive operations, and many argue that the customers have benefited. Those companies operating efficiently enough to make a profit have continued to operate, while those companies that relied on over-built and inefficient systems either changed their operating procedures, or succumbed to the pressures of competition.

In order to promote competition in the electric marketplace, the Energy Policy Act (EPAct) was signed into law in 1992 by President Bush. This action expanded the power that the Federal Energy Regulatory Commission (FERC) has over the electric energy marketplace. The FERC now has the power to order utilities to provide transmission services for wholesale customers. In theory, transmission-line owners no longer have any priority over other would-be transmission line users to use transmission lines. This was termed *open access*. Under open access, independent power producers, cogenerators and large retail customers have the same access to transmission facilities as the traditional electric utility generators, even though they may not own any of those facilities. Prior to passage of EPAct, the traditional vertically integrated electric utilities were accustomed to having a guaranteed amount of customers and transmission privileges. After EPAct was signed into law, these utilities were theoretically required to compete with other types of power producers with access to inexpensive energy. Supposedly, power producers with inexpensive power were given a foothold into the electric market. In some areas, traditional electric utilities found themselves with competitors after monopolizing the electricity market for years.

In addition to affecting the utilities, the signing of the EPAct potentially held implications for electric customers. Large customers gained the opportunity to deal directly with power suppliers. One would assume that these customers are able to purchase power less expensively when they have a choice, instead of being forced to buy from their present local distribution company. In order to qualify for these benefits of EPAct, smaller residential customers could form groups representing themselves to negotiate better rates with the power providers in the surrounding regions.

EPAct was intended to promote competition in the electric marketplace, but many of the utilities have decided that they weren't ready for such dramatic change without more encouragement. In some cases, utilities have refused to allow others access to their transmission lines, or charged "unfair" rates for the usage of their transmission equipment. True open access is not here yet in all areas. With the release of its NOPRs (notice of proposed regulation) the FERC is beginning to convince the electric industry that it is serious about making an environment with open access and competition, a reality.

When open access is fully implemented, the next step will be to take competition even further, so-called retail competition. When this happens, all electric customers will have a choice of electricity providers. This will be a vast change from the electric marketplace of today. Customers won't necessarily make out a monthly check to the same electric utility. Power suppliers and customers could arrange their own transactions via contracts for market-priced electric energy. The everyday residential customer might not be interested in negotiating with power suppliers, but many service provider companies would be willing to represent groups of residential customers. If the electric industry follows the communications industry under its deregulation, there could be electricity service providers requesting that people use their company to provide electric service much the same way as the communications companies compete for customers.

With so many customers, distribution companies, and service providers wanting to negotiate with power producers for their electricity, it might be difficult to determine what the going price for electricity should be at any given moment. It could be difficult to learn which supplier is sold out, or has a surplus of electricity. Many have proposed that it would be convenient to establish regional energy exchanges where buyers and sellers participate in an electric power auction. Under a scenario where there are regional energy exchanges, interested parties could discover what the market price is in their region without contacting all individual power producers. This system has worked well at the Chicago Mercantile, New York Stock Exchange and many other exchanges where commodities (other than electricity) have been traded for many years. Recently the New York Stock Exchange has added electricity futures to their offerings.

In the fully deregulated, or re-regulated environment, electric energy market participants could arrange their own transaction directly, or make transactions through a broker at an electric commodity exchange. The commodity exchange approach would allow large numbers of power generators and purchasers to buy and sell power at prices obtained via the auction process. The auction process that is proposed in this thesis involves taking bids from buyers and sellers, and matching them so that marginal utility is maximized. As part of operating the electric system under this new system, a central coordinator/broker would evaluate transactions to ensure operation within stability limits and ensure that transmission lines were not overloaded. Proposed transactions that would compromise system integrity or force operation outside of limits would not be allowed.

The main emphasis of this thesis is the development of bidding strategies for the agents participating in a energy exchange type auction. In particular, strategies to increase the profit of those agents that produce power and offer it for sale in an auction are investigated. Genetic algorithm (GA) based agents repeatedly play out the auction scenario. Successful agents are allowed to reproduce according to the rules of GAs, and their strategies are evolved resulting in better bidding strategies.

An integral part of developing a successful bidding strategy is successful price forecasting. Agents which are able to correctly predict the equilibrium price, are better equipped to place an effective bid which will result in high profits. The adaptive agents described in this research are given their choice of different statistical prediction methods to project prices. These price projections will be used to make bids via expert system rules. The bids are then submitted to an auctioneer who matches bids between buyers and sellers. If sufficient transmission capacity exists, and an agreeable price is discovered, the buyers and sellers have a contract. Knowing the selling or buying price and a cost of production, it is possible to calculate the profit made by each agent. The profit becomes an indicator of how successfully each agent's price prediction and bid making rules are functioning.

Chapter II describes recently published work related to this thesis. It discusses research that has been performed dealing with evolving economic agents, developing successful bidding strategies, deregulation, and the auction environments. Chapter III describes the methods investigated and/or used for the research presented in this thesis, including genetic algorithms, price forecasting, and neural networks. Chapter IV presents the results of this research. Finally, Chapter V presents the conclusions made as a result of the research and lists several directions in which this research may be extended.

CHAPTER II. LITERATURE REVIEW

This chapter presents recent research related to developing competitive bidding, evolving trading agents, and deregulating the electric marketplace. It introduces some of the basic auction methods developed by others, used in this research to develop results.

The move toward more competitive reliable electric marketplaces is not confined to North America. The England and Wales Power Pool is presently operating their electric system using a "competitive bidding dispatch mechanism." The country's transmission grid is owned and controlled centrally by the government. Generators bid for the right to produce electricity in a particular time period. They submit price, capacity, and other information crucial to the stable operation of the electric system to the central coordinator. The central coordinator performs a unit commitment analysis, and based on the projected demand and prices, offers to produce power at a price are either accepted or rejected. If a generator's bid is accepted, they have the right to generate electricity for the time period in question.

Finlay [4] analyzes bidding strategies for the English system, and shows mathematically that there exists an optimal bidding strategy for bidders competing in the English system. Finlay does not seek to maximize the profit of the individual generation companies, but uses a mathematical approach to show the new English system is equivalent to the cooperative power pool approach.

Post [19] describes four types of basic auctions, including the English auction, the Dutch auction, first-price sealed-bid auction, and the second-price sealed-bid auction. The English method involves orally announcing bids until the bid announced is considered acceptable by only one buyer. The bidding stops at a level approximately equal to the valuation of the item by the second highest bidder. (A bidder's valuation of the item, is the worth that the bidder places on that item. Through his bids, the second bidder has been indicating that the item in question has not yet reached his valuation. At the time he stops bidding, that bidder's valuation of the item is known to be bounded by the last acceptable bid and the final bid.) The Dutch auction is a descending-bid auction during which an auctioneer calls a decreasing set of prices beginning with an initially high price, until a potential buyer finally accepts a the current price. There is only one true bid, which is by the bidder who accepts the price first. The third method of auction that Post describes is the first-price sealed-bid auction. Potential buyers submit sealed bids for an item. The item is awarded to the bidder who has submitted the highest bid, for the price that he bid. The final method described by Post is the second-price sealed-bid auction. Bidders submit sealed bids, and the item being auctioned is sold to the bidder with the highest price, but for a price which is equal to the second highest bid. Although to the reader, one may seem more logical than others, all of these have been used by various industries at one time or another.

Post lists variations on these auctions include imposing reservation prices, time limits for submitting bids, entrance fees for participation, assigning royalties, and standardizing on a minimum acceptable increment to the last bid accepted.

When more than one unit of a good is to be sold, and each bidder demands only one unit, a multiple auction may be used. The multiple auction allows more that one bidder to be awarded a bid. Post also describes variations on the multiple auction.

In addition to the four basic types of auctions mentioned above, Post also describes the double auction. Under the double auction, several buyers, and several sellers submit bids and offers for one unit of a good. If and when a buyer accepts a seller's offer, or a seller accepts a buyer's bid, a binding contract is made. Post points out that "few theoretical results on double auctions exist since modeling strategic behavior on both sides of the market is difficult."

Recent work by Kumar and Sheblé [12] proposes a framework for an energy brokerage system. Similar to the double auction described by Post, [19], the buyers and sellers interact through a central coordinator who matches the bids. The central coordinator is responsible for ensuring that the energy transactions resulting from the matched bids do not upset the electrical transmission system.

Under the framework described in Kumar and Sheblé, vertically integrated monopolies are not allowed. Companies presently having both generation and distribution facilities would be divided. In the framework developed by Kumar and Sheblé, all power is generated by generation companies (gencos), and all power is sold to distribution companies (discos). See Figure 2.1. Each of the gencos and discos is a completely separate entity that does not communicate directly with any other concerning projected demand, capacity or cost. Companies previously having combined coordination over their generation and distribution systems, will now only coordinate via prices transmitted to a central coordinator or Independent System Operator (ISO). It is necessary to have a central coordinator to monitor and respond to the power system limits and transmission capacities. Gencos and discos are required to cooperate with the ISO in maintaining system reliability by supplying crucial generator parameters which will not available to the other gencos or discos.

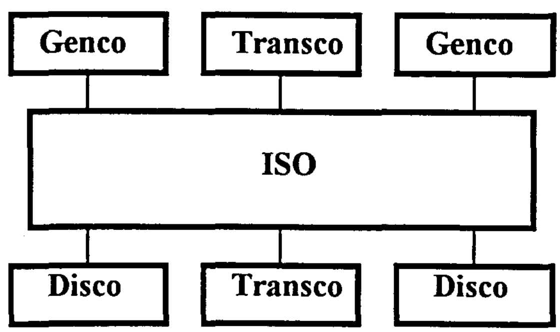

Disco: Distribution Company
Genco: Generation Company
ISO: Independent System Operator
Transco: Transmission Company

Figure 2.1 Brokerage system model

The research described in this thesis uses a double auction model with seller's being gencos, and the buyer's action being simulated by one large disco. The gencos are represented by selling agents and are allowed to participate in the auction simulations. They evolve based on their performance. Those whose bidding results in greater profits than the others are rewarded with another opportunity to participate in the auction. Those whose profit lower than the others, will not be allowed to participate further in the auction, and will be replaced by new agents. Those agents which still remain after hundreds or thousands of trials, may demonstrate qualities which are useful in developing bidding strategies for real generators. The rules that they are using to develop their bids can be decoded and duplicated for a real genco.

Developing bidding strategies in this manner for the deregulated electrical utility industry is a new field of research. Outside of the electrical utility industry, interest has grown in recent years for using adaptive agents to simulate trading behavior. Experimenting with adaptive agents has provided a useful means of exploring different aspects of trading markets.

LeBaron, in [13] describes work on evolution and learning as it has been applied to financial markets. The types of evolutionary learning used in computational modeling ranges from traditional least squares learning to more advanced or complicated artificial intelligence techniques like neural networks, classifier systems, and genetic algorithms. LeBaron states that one of the most crucial points of economic modeling and simulating financial markets is the actual trading mechanism used in the simulation. According to LeBaron, it can be difficult to get the interacting agents to clear the market, and how one goes about implementing this task can greatly affect the results.

In [13], LeBaron outlines experiments from the ongoing project of Santa Fe researchers, called the Santa Fe Stock Market. The environment in which the experiments take place has agents selecting between a risk free bond and a risky stock paying a stochastic dividend for their investment portfolio. The bonds pay a constant interest rate, and there is an unlimited supply of bonds.

Agents in this environment make "intelligent" bidding decisions using a classifier system. The classifier system chooses one (or more) set of forecasting parameters, a, b to be used in a general formula for determining the bid amount. Each set of parameters will be used in a general formula (see equation 2.1) to develop a conditional linear forecast.

$$E_{t,i} (p_{t+1} + d_{t+1}) = a_{i,j} (p_t + d_t) + b_{i,j} \quad \text{(equation 2.1)}$$

where:

p_t = price at time t

d_t = dividend at time t

i,j = jth rule in the ith agent's rule set

Agents are allowed to select up to twelve discrete indicators to invoke the different sets of forecasting parameters. In addition to the twelve indicators defining different levels of return on investment and price, LeBaron also includes three additional inputs. LeBaron's classification parameters are as follows:

1. [price*interest / dividend] > 1/2
2. [price*interest / dividend] > 3/4
3. [price*interest / dividend] > 7/8
4. [price*interest / dividend] > 1
5. [price*interest / dividend] > 9/8
6. [price*interest / dividend] > 5/4
7. [price*interest / dividend] > 3/2
8. price > 5 - period moving average
9. price > 10 - period moving average
10. price > 100 - period moving average
11. price > 500 - period moving average
12. always on
13. always off
14. random (on off)

The Agents have twelve evolving bits in their gene which decide what criteria shall be used to invoke a particular set of forecasting parameters. The forecasting parameter sets are part of the genetic algorithm gene, as shown in Figure 2.2.

The classification antecedent (the first fourteen bits) and forecasting parameter sets (the last part of the gene) that perform the worst are replaced each generation and new sets are developed using standard GA crossover and mutation. LeBaron defines a symbol (#) which has neither a 0 nor 1 value, but which is used like a wildcard. In the described environment the use of the don't cares makes it possible to select more than one strategy initially and then use the one that produces the best result in the final analysis.

LeBaron discusses some of the simple experiments that have been performed using the described environment. LeBaron mentions that the population of agents can be unstable,

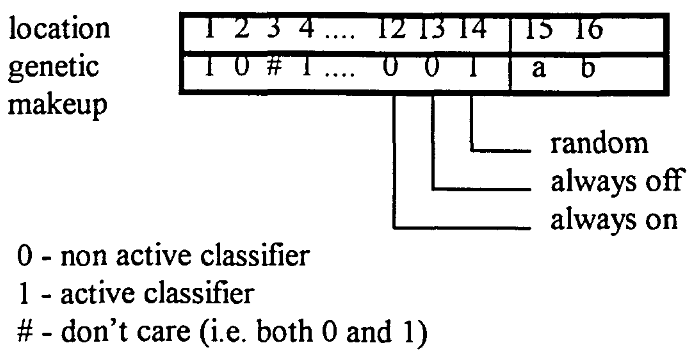

Figure 2.2. GA gene as in LeBaron, [13]

and they can react differently than one would expect based on traditional theoretical models. This observed scenario is commonly exhibited in systems with evolutionary learning. Certain creatures will take advantage of those that do not have adequate defenses or that act stupidly, consequently their population increases, until the players who they are able to take advantage of, become extinct. LeBaron concludes that the financial evolutionary learning field is fairly young.

Other researchers have developed similar platforms for examining financial trading behavior, in addition to other types of behaviors. In [9], Ioannides reviews different trading structures and their associated topologies. In [25], Tesfatsion describes research in which trading agents have an expected payoff for entering into a competition with another agent. They are allowed to compete with those agents that will increase their fitness. In [25], the competition is an iterated prisoner's dilemma game, but could modified to allow the agents to participate in a round of double auctions. Interesting behaviors are exhibited as a result of the interaction between agents. Ashlock, in reference [1], uses genetic programming combined with a finite state automata to play a theoretical game called Divide the Dollar. In Divide the Dollar, two players each submit a bid in the range from \$0 to \$100. If the sum of the two bids is less than or equal to \$100, the players each receive the amount of their bid. If the sum exceeds \$100, neither player receives anything. Interesting behavior patterns emerged when Ashlock used genetic programming and the finite state automata in evolving players for this game. Although Divide the Dollar is not financial trading, the resulting player behaviors might provide insight into how traders cooperate to increase market efficiencies and profits.

CHAPTER III: METHODS

This chapter describes the methods used in developing the bidding strategies presented in this thesis. It presents techniques investigated and/or used to forecast equilibrium prices during the bidding process. Basic genetic algorithms, used in evolving competitive agents and their strategies, are presented in this chapter. Following the discussion of basic GAs, the specific GA used in this research is described in detail. The auction model used to obtain the results presented in chapter IV is outlined.

Price Forecasting Methods

It would be easier to maximize profit if a particular agent knew in advance what the price of electricity would be. It is impossible, without insider information, to know a priori what other agents' bids will be, and consequently impossible to know, in advance, the exact bid that will be accepted. Even though it is not possible to know future prices exactly, it is possible to forecast future prices, and use the forecasted price in developing a successful bid, or bidding strategy.

In the research being described here, several methods were investigated for the purpose of forecasting prices based on historical price data. The following price prediction methods are described in this section:

- linear regression
- moving average
- exponentially weighted average
- weighted moving average
- neural network

Linear regression

Simple linear regression involves using historical data to develop the equation of a line that "fits" the data. Assuming that the equilibrium prices in future time periods will continue to fit this linear equation, the slope of this line can be used to predict future equilibrium prices. In order to make a good prediction, it is necessary to select a window, or number of data points, over which data will be considered for the linear regression. If data points covering a long history are selected, the line equation will not be sufficiently influenced by the most recent data, and the amount of error will be high. To obtain the equation of the line the following formula is used:

$$x(t+1) = x(t)\,(\eta) + x(t-1) \qquad \text{(equation 3.1)}$$

where:

t = time period index

$x(t)$ = price at time period t

$x(t+1)$ = predicted price at time period $t + 1$

$x(t-1)$ = price at time period $t - 1$

η = slope of the line

The linear regression method of statistical forecasting is the least likely of traditional prediction schemes to simply lag the actual value. It tends to project prices that are more volatile than the actual prices. Although it does not simply lag the actual values, it can have a rather high mean squared error due to the volatility in the forecasted prices.

Moving average

The moving average predicts that the price will be the average of the last s observed prices, where the number of historical prices considered in developing the prediction. It gives an estimate that tends not to respond rapidly to changes the input, which is the equilibrium price in this research. It works more as a smoothing filter rather than a true predictor of prices. The amount of smoothing depends on the number of observations, or its width, being used to determine the moving average. The width or number of observations to consider can be increased or decreased to produce a predictions which tend to vary less or more respectively with fluctuations in the observed values.

$$x(t+1) = [\, x(t-s+1) + \ldots + x(t) \,]/s \qquad \text{(equation 3.2)}$$

where:

s = width of the window, or number of prices to be considered

t = time period index

x(t) = price at time period t

x(t+1) = predicted price at time period t + 1

For example, if the width of the data window were 4, and the goal was to predict the next value in the series of numbers [8, 7, 4, 3, 2.5, 5], we would have (5+2.5+3+4)/4 = 3.625, as our prediction using the moving average. As previously stated, moving averages tend provide a smoothed representation of a series. If the prediction is to be used to for stock market speculation where it is desirable to take advantage of short term fluctuations in price, then moving averages would not be the best prediction selection technique. However, if we are attempting to reduce the effect of hourly or daily fluctuations in a price, then moving averages would be very appropriate.

Weighted moving average

If we are interested in predictions that can respond quickly to increases or decreases in the price of a commodity, and be used to influence bidding practices, we might be interested in a prediction technique which doesn't simply smooth the data and give an average. It is possible to modify the normal moving average so that more emphasis is placed on more recent data. Weighting the more recent data more heavily can help produce a better response to quickly changing markets. This might be a good trade off between the smoothing of the plain moving average and the linear regression in our application.

With the weighted moving average, as in the moving average, we focus on a certain window of observations when making our predictions. More recently observed values are weighted more heavily by multiplying them by some constant. Observed values further into the past are multiplied by a smaller factor, making their contribution to the prediction smaller. If information is available that correlates elapsed time and a how much weight an observed value deserves, it could be used in determining the weighting factors. The observed values, adjusted by the weighting factor, are summed and the final prediction is made by dividing the summation of the weights. Gilchrist in [5] illustrates this with the following example which considers four observations:

$$x(t+1) = 1/10\,[\,x(t-3) + 2\,x(t-2) + 3\,x(t-1) + 4\,x(t)\,], \qquad \text{(equation 3.3)}$$

The ten is a summation of the weightings associated with the data values. i.e., 1 + 2 + 3 + 4 = 10. Here, the most recently observed value is weighted four times as heavily as the first, the next most recent is weighted three times as heavily as the first and so on.

Exponentially weighted moving average

The exponentially weighted moving average (EWMA) weights the observed values in as described weighted moving average. It is a specific case of weighted moving average that makes it easy to consider many past values. This advantage of the EWMA gives it a very simple recurrence form. The general form is given by the following:

$$x\,(t+1) = \frac{\sum_{r=1}^{t} a^r \bullet x(r)}{\sum_{r=1}^{t} a^r} \qquad \text{(equation 3.4)}$$

Gilchrist in [5] uses the following example to clarify. Assume that we are examining some weekly data where the observed values are 7, 5, and 6 for the first, second and third weeks respectively. Our prediction for the fourth week's data would be $(0.8^2 \times 7) + (0.8^1 \times 5) + (0.8^0 \times 6) / 2.44$, or 4.46.

It is not necessary to recalculate each part of this equation for every prediction, as the predicted value can be obtained using numbers that were multiplied and added to get the last prediction. In the example above, we can use the prediction for the present week to help get next week's prediction.

Neural networks

More correctly termed artificial neural networks (ANNs), ANNs were investigated during this research for predicting prices. They have been used in power systems for predicting demand which is directly involved in predicting the cost of power. See Papalexopoulos [18] and Mohammed [16]. Until now, electric power has always been sold at a rate based on a cost, not a price, therefore neural networks have not been used specifically for predicting equilibrium prices for electric power. ANNs have been applied to predict prices of stocks and the like.

According to Hakin [7] "a neural network is a massively parallel distributed processor that has a natural propensity for storing experiential knowledge and making it available for use. It resembles the brain in that its knowledge is acquired by the network through a learning process, and that inter-neuron connection strengths known as synaptic weights are used to store the knowledge." For the lay person, a neural network can be thought of as a black box that when presented with new inputs can respond with an appropriate output based on correlations drawn via observed input-output pair relationships. See Figure 3.1. Prior to using the neural network, it must be trained using a learning algorithm.

The neural network architecture used in this research is the multilayer perceptron (MLP). See Figure 3.2. The MLP is constructed of interconnected perceptrons or neurons. See Figure 3.3. The perceptron is a simple device having multiple inputs and a single output.

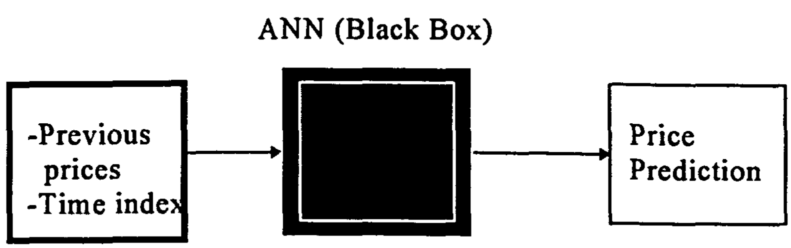

Figure 3.1. Lay person's diagram of an ANN

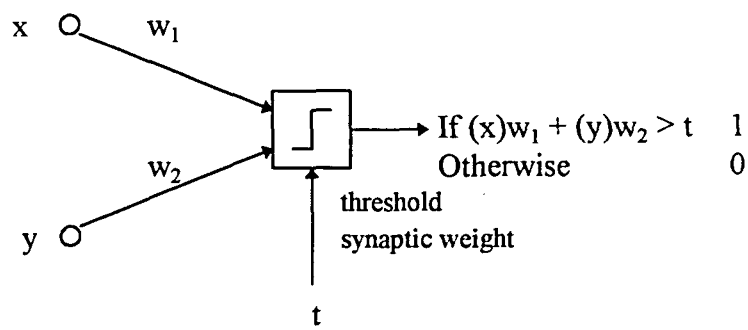

If the summation of the product of the inputs with their synaptic weights is greater than the threshold, the output is turned on and vice versa.

Figure 3.2. Perceptron model

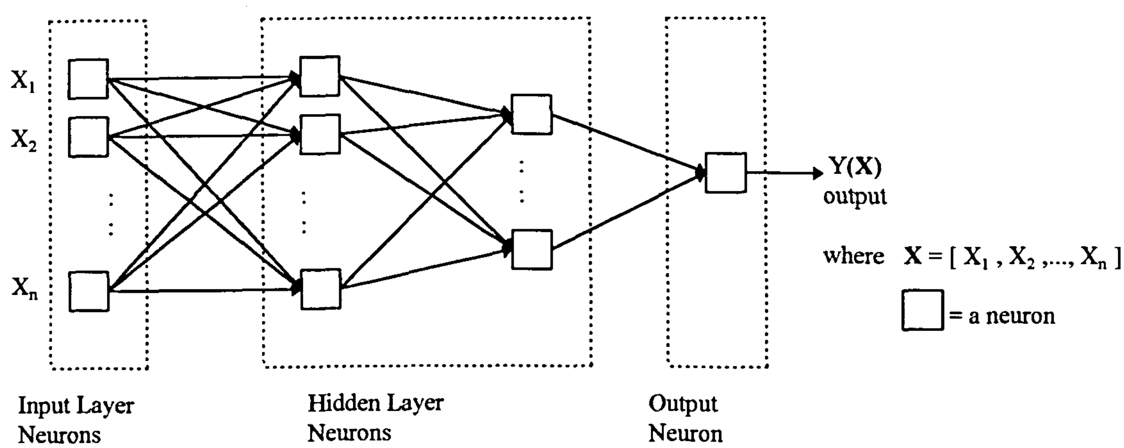

Figure 3.3. Architectural graph of multilayer perceptron with two hidden layers

It produces an output that is either on or off, based on the weighted sum of the inputs that it observes. If the summation of the weighted inputs are above a threshold value, the neuron outputs one. If the summation of the weighted inputs is below the threshold value, the neuron outputs zero.

If the output of the ANN is incorrect, the weights may be modified. During a training period, data (with a known output response) is presented to the ANN. When the output of the ANN does not agree with the answer, the weights are modified. Weights can be modified after presentation of all of the training data (batch training), or after the presentation of each set.

The training method used to set the synaptic connection weights in this research is the back propagation learning algorithm discussed in [7]. See Hakin for the weight updating equations. Training is done using supervised (each training input has a known output) learning. As mentioned above, inputs are presented to the network, and the difference between the corresponding network output and the known correct output is used to adjust the connection weights.

During training, the mean squared error (MSE), which is based on difference between desired and observed outputs, is monitored. Training continues until the MSE is below some tolerance limit. The weights are then considered tuned or set, and the network is then ready respond with output when presented with inputs. For a more complete description of network training, see Hakin [7].

For the research discussed in this thesis, ANNs would seem to work well for predicting equilibrium prices. Application of ANNs to the price prediction portion of this research is still under development, and for the results described in chapter IV, the other methods discussed in this section are used.

Basic Genetic Algorithms

In the research described here, agents are evolved with a genetic algorithm. Genetic Algorithms (GAs) are general-purpose search techniques derived from the biological model of evolution. They are often able to find solutions to problems which can not be handled by normal optimization techniques. Genetic algorithms have been shown to work on problems like the unit commitment [10] and the economic dispatch [27] problems in the electrical power systems field. Genetic algorithms can search solution spaces in parallel. Entire populations of candidate solutions termed creatures, are initialized and allowed to evolve. These creatures are often represented by binary string representations, sometimes referred to as the creature's DNA, which can be mapped into the solution space for evaluation. Each creature is assigned a fitness by a heuristic measure of its quality. Then during the evolutionary process, those creatures which exhibit a higher fitness are favored and allowed to procreate. During each generation of the evolutionary process, creatures are randomly selected with a fitness bias for reproduction. The type of fitness bias determines the *parent selection method.* Through the processes of *crossover* and *mutation*, new creatures are developed which explore a different area of the solution space. These new creatures replace lesser fit creatures from the existing population. See Figure 3.4 for a block diagram representation. The basic algorithm can be written as follows:

1. Randomly initialize a population and set the generation counter to zero.
2. Until done, do the following:
    - Calculate the fitness of each member of the population.
    - Select parents using some fitness bias.
    - Create offspring from the selected parents via crossover.
    - Mutate these new offspring.
    - Replace the lesser fit members of the population with the newly created offspring.
    - Increment the generation counter and go to the beginning of step 2.

Parent selection methods

Similar to Darwin's theory of survival of the fittest for evolving biological species, members of the GA population are allowed to reproduce based on the fitness they exhibit.

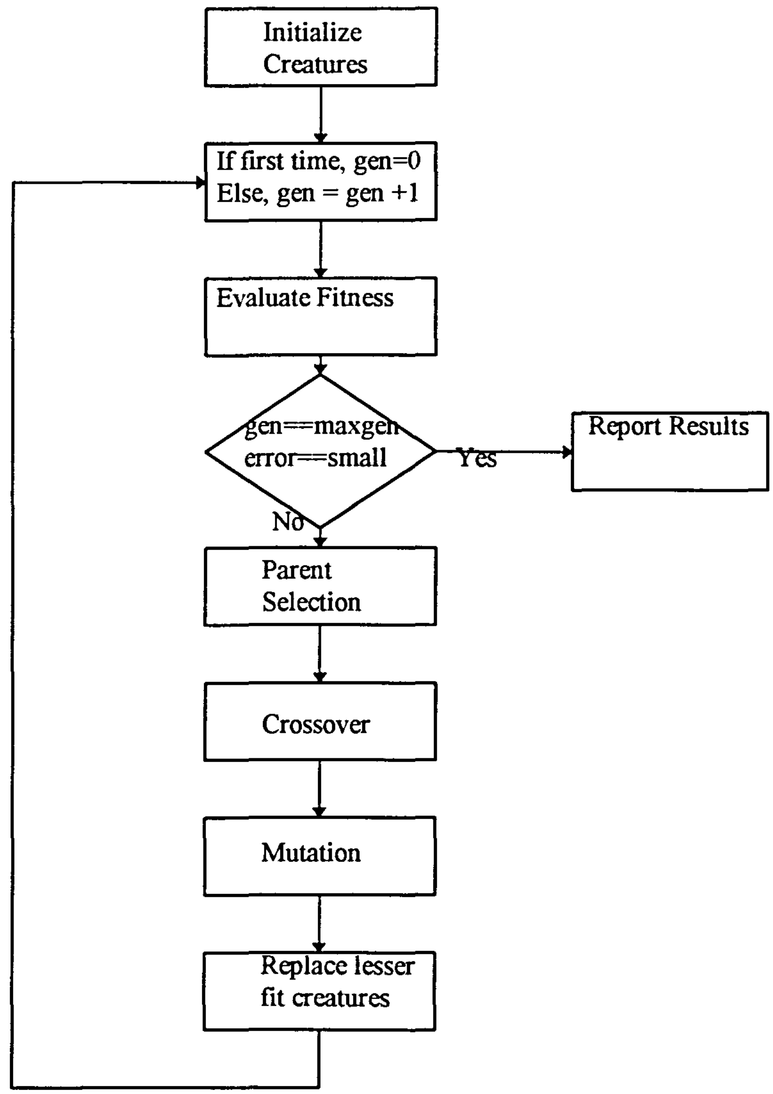

Figure 3.4. Genetic algorithm block diagram

Creatures selected for reproduction are called parents. Parents are selected randomly from the population with a fitness bias which tends to select more highly fit creatures for reproduction. There are several successful means of selecting parents, each with different selection fitness biases which can move the collective fitness of the population in different directions. Some of the more commonly used methods are described in the following paragraphs.

Roulette selection

Roulette selection chooses creatures to be parents with some probability. The probability that a creature will be chosen is proportional to that creatures fitness among the entire population. Many people prefer to visualize a wheel or circle whose area is 1. The fitnesses of each of the individual members of the population are be mapped onto this probability "wheel". Each member's assigned space on the wheel corresponds to its proportion of the entire population's fitness. This probability wheel is similar to the roulette wheel used for gambling, except that the areas on the wheel which are not equally sized. The "wheel" is spun. When it stops, it points to the creature selected to be a parent. See Figure 3.5.

Rank selection

Rank selection is very similar to roulette selection. First the members of the population are sorted by their fitness value. Next, they are mapped onto a probability wheel as in roulette selection. However, under rank selection, the sections of the wheel do not correspond to the individual's fitness proportion, but are proportional to the creatures rank

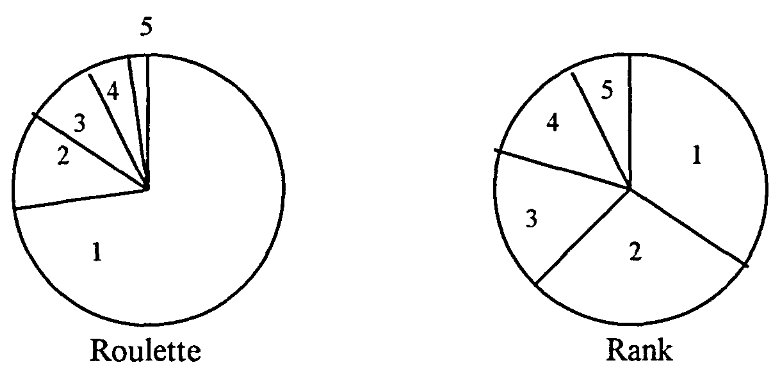

Note: Roulette selection can produce excess fitness bias for parent selection. Rank selection tends to avoid rapid convergence to a local optima better than roulette selection.

Figure 3.5. Roulette and rank probability wheels

within the population. Rank selection does not bias the selection to the most highly fit creature as strongly as roulette selection. This is sometimes a desirable trait, as it can help prevent the GA from getting caught in a local optima. See Figure 3.5. After a species reproduces for many generations, if there is not a large gene pool, members of the species could start to look identical. In a genetic algorithm, these members of the population that look the same are searching the same solution space. A population that is not searching the same solution space, but has a large pool of genetic material at its disposal is said to be biodiverse. In many cases, especially in those cases with a rough fitness landscape (see Figure 3.6), it is beneficial to maintain biodiversity for as long a time as is possible.

Tournament selection

Tournament selection is substantially different from the roulette and rank selection methods. It maintains biodiversity and works well for those problems which have a rough fitness landscape characterized by tall narrow peaks. Figure 3.6 shows both a smooth and a rough fitness landscape. This method does not involve the use of a probability wheel. Instead, each creature is shuffled into a group of four (tournament) based on a number randomly selected from a uniform distribution.

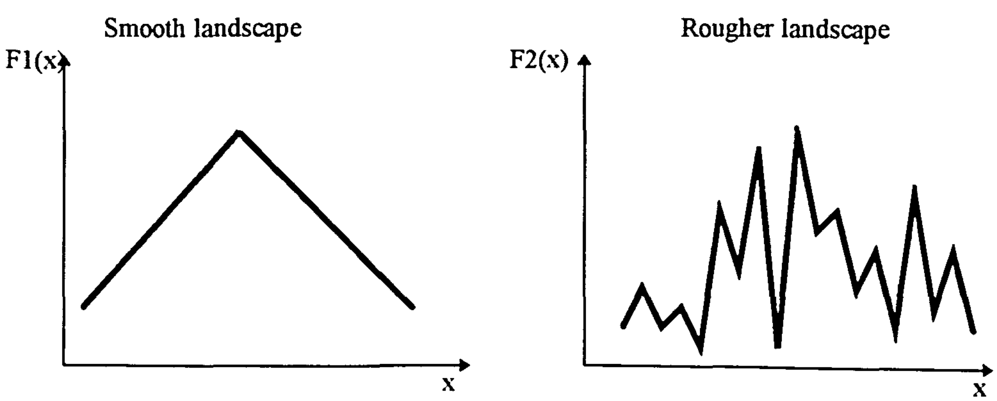

Note: Maintaining biodiversity can aid in finding the global optima of more complex functions.

Figure 3.6. Fitness landscapes

The two creatures with the highest fitness in each tournament are selected as parents, and their offspring replace the two creatures with the lowest fitness in their group. The first and second most fit creatures in this method are safe from replacement. Conversely, the first and second most unfit creatures are guaranteed to be replaced, but it is not possible to say what will happen deterministically to the other creatures at each generation.

Crossover

Once the parents have been selected (and matched in a pair), they have children. New creatures, commonly termed children, are created by recombining the parents' binary strings. Parents are required to be in pairs for reproduction, and the result is two children for pair. Children are created by copying the bits from parent 1 into child 1 and from parent 2 into child 2 until a randomly selected crossover location is reached. At this point, bits are copied from parent 1 into child 2, and from parent 2 into child 1. If there is one crossover point separating the areas to be swapped, it is known as single point crossover, and is presented graphically in Figure 3.7. Multiple point crossover operates on the same principle

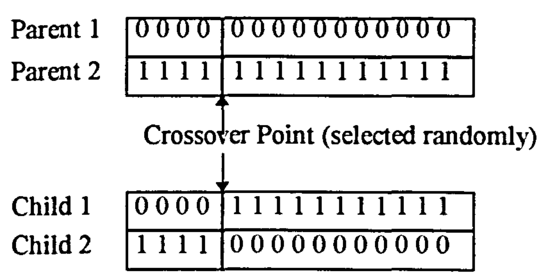

Figure 3.7. Single point crossover

but uses more crossover locations. Experimentation can help decide how many crossover points to use for a particular problem.

Mutation

Following the crossover process, the children are mutated. Biological mutation results from errors during DNA replication. These errors may or may not have a large effect on the resulting creature. Mutation in genetic algorithms is derived from biological mutation, and serves to prevent the GA from getting stuck in local optima. It does this by introducing new genetic material into the gene at some low rate. If the gene to be mutated in the child is represented by a binary string, mutation involves flipping the bit (0 goes to 1, 1 goes to 0) at each location in the string with some probability. See Figure 3.8.

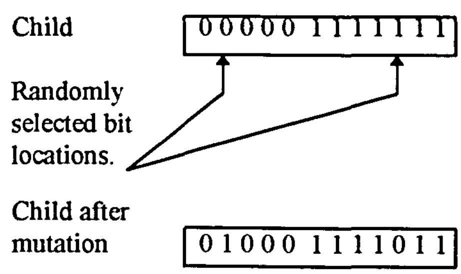

Figure 3.8. GA mutation

Fitness evaluation

A creature's fitness is a heuristic measure of the creatures performance or goodness. There are many ways to evaluate a creature's fitness. The GA designer must use engineering judgment in determining how to award fitness. For example, if a genetic algorithm was being used to solve the unit commitment problem, a creatures fitness should be inversely proportional to the cost of the solution that it finds. Costs are first evaluated for the corresponding solution using our cost equations, and those creatures with higher associated costs are awarded a lower fitness. If the goal of the GA is to develop bidding strategies for auctions which maximize profit, creatures which find a solution resulting in a higher profit are awarded higher fitness than those agents making a smaller profit. In the genetic algorithm described in this research, the fitness is proportional to the genco's profit. It is possible to have the fitness depend on more than one parameter. A GA can have a primary and a secondary fitness function. In such a case a solution represented by a member of the population might receive fitness for profit, but if no profit existed, it might award some fitness points for not having violated any environmental or operational constraints.

Auction Mechanism

This section presents the auction used to simulate transactions between the gencos and discos through the independent system operator (ISO). It describes the rules, according to which, the auctioneer or ISO will match the bids. It also describes the process by which the agents develop their bids prior to submitting them to the ISO.

The research presented here uses a double auction similar to that described in Post [19]. Distribution companies, or discos, are always the buyers, and the generation companies, or gencos, are sellers. Buyers and sellers submit their bids and offers respectively to a central auctioneer, which in our case, is the ISO. This is based on the auction market framework proposed by Kumar and Sheblé [12]. This thesis is primarily concerned with developing bidding strategies for the gencos and pays less attention to strategies of the discos. See Figure 3.9. In future simulations, discos and gencos will both

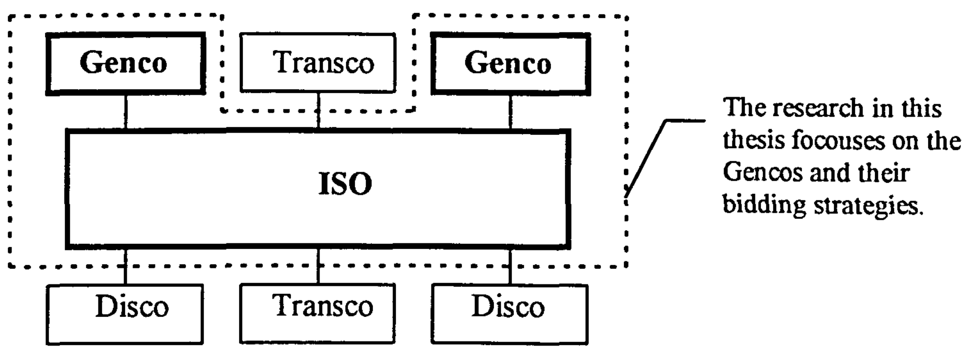

Figure 3.9. Auction framework used for simulations

be allowed to buy or sell electric power for the purposes of reselling, provided this scenario is profitable.

The bids submitted by the discos and gencos to the ISO are sorted so that the disco bids are arranged in descending order, and the genco offers are arranged in ascending order. Each bid is specified in dollars per megawatt. Contracts are standardized and are for one

MW of electrical power. The contract also specifies delivery points, power quality, amount of voltage regulation to be provided, etc. The disco with the highest bid for a contract is matched with the genco offering power for sale at the lowest price. The difference in prices is split, and the equilibrium price becomes the midpoint of the bid and offer being matched. If these two auction participants have exhausted their demand or supply, the ISO then would match up the second highest bid with the second lowest offer, and so on. See Table 3.1 for an example. This is done until there are no valid matches left. A buyer or seller may bid on more than one contract. Each bid is entered and sorted as in the example shown in Table 3.1.

Table 3.1. Example of auction bid matching

Buyer bid ($/MW)	Seller offer ($/MW)	Contract	Equilibrium price ($/MW)
12.50	8.50	Yes	10.50
12.00	9.00	Yes	10.50
11.80	10.00	Yes	10.90
10.00	10.50	No	NA
9.50	11.00	No	NA

Note: In this example, buyers and sellers are bidding on contracts of one MW each.

The bidding takes place in rounds or cycles. Discos and gencos submit bids and offers to the ISO, and the ISO, after verifying that the system can sustain the transaction, matches the bids. If the bids are such that there are no valid matches, or there are many players not matched, true price discovery has not occurred. This is reported to the market players and another round of bids is requested by the ISO. Prior to the close of trade (final round of bidding) determined and known by the ISO, the bidders may request another round of bidding if they want to attempt to buy or sell more power. Each time the discos and gencos submit their bids and offers, they are not be certain that this cycle of bidding will be discovery has occurred. At this point participants whose bids have been matched are said to have a binding contract. See Figure 3.10.

Following a round, buyers or sellers can request additional rounds if they have additional power to buy or sell. During the auction process, if a round is unsuccessful at determining the equilibrium price, buyers are not allowed to decrease, and sellers are not allowed to increase the price that is bid during subsequent rounds. For more details of the auction framework see Kumar [12].

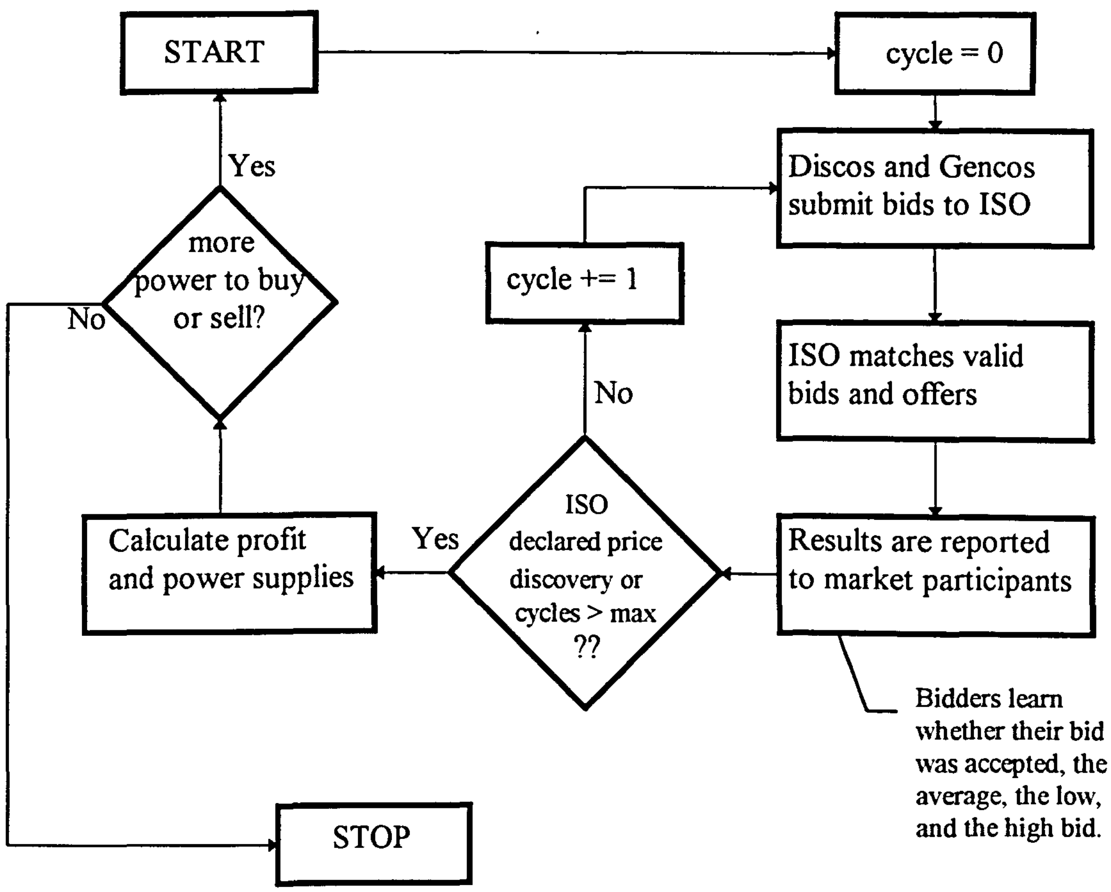

Figure 3.10. Bid matching process

Evolving Price Prediction Parameters with a Genetic Algorithm

Gencos submitting bids that are accepted, and that result in a profit will outperform those who submit bids that are never accepted. This is one reason why accurate prediction of the equilibrium price might become an important issue. If gencos could know a priori those bids which would be accepted and rejected, they could increase the profit that can be made. To aid the gencos in submitting good bids, the algorithm used in this research contains four separate genetic algorithms that evolve prediction parameters used in predicting equilibrium prices. There is one GA population for each of the four prediction techniques that are being used, including linear regression, moving average, weighted moving average, and exponentially weighted moving average.

For the weighted moving average the genetic algorithm evolves the weights of each of the inputs, the spacing between the inputs, and the number of inputs to consider i.e. the width of the input window. For the moving average the genetic algorithm evolves the width of the input window and the spacing between the inputs to be considered. For the linear regression the GA is evolving the window of values and the spacing. For the exponentially weighted moving average, the only value that is evolving is the exponent. These parameters are then used with the respective prediction techniques along with the equilibrium price history to predict the future equilibrium price. See Figure 3.11.

Evolving Agents with the Genetic Algorithm

The genetic algorithm used for this research is fairly complicated and can be subdivided into three areas. The first area consists of those functions related to the agents. This includes agent initialization, cost calculation, developing bids, and the GA processes like parent selection, crossover and mutation associated with the evolution of the bidding agents. The second area or subdivision consists of the equilibrium price prediction mechanism. Separate populations of prediction parameter sets are evolved to optimize the prediction schemes. In this research, four different prediction methods were used to predict price, and each of these four has a population of prediction parameter sets that evolves

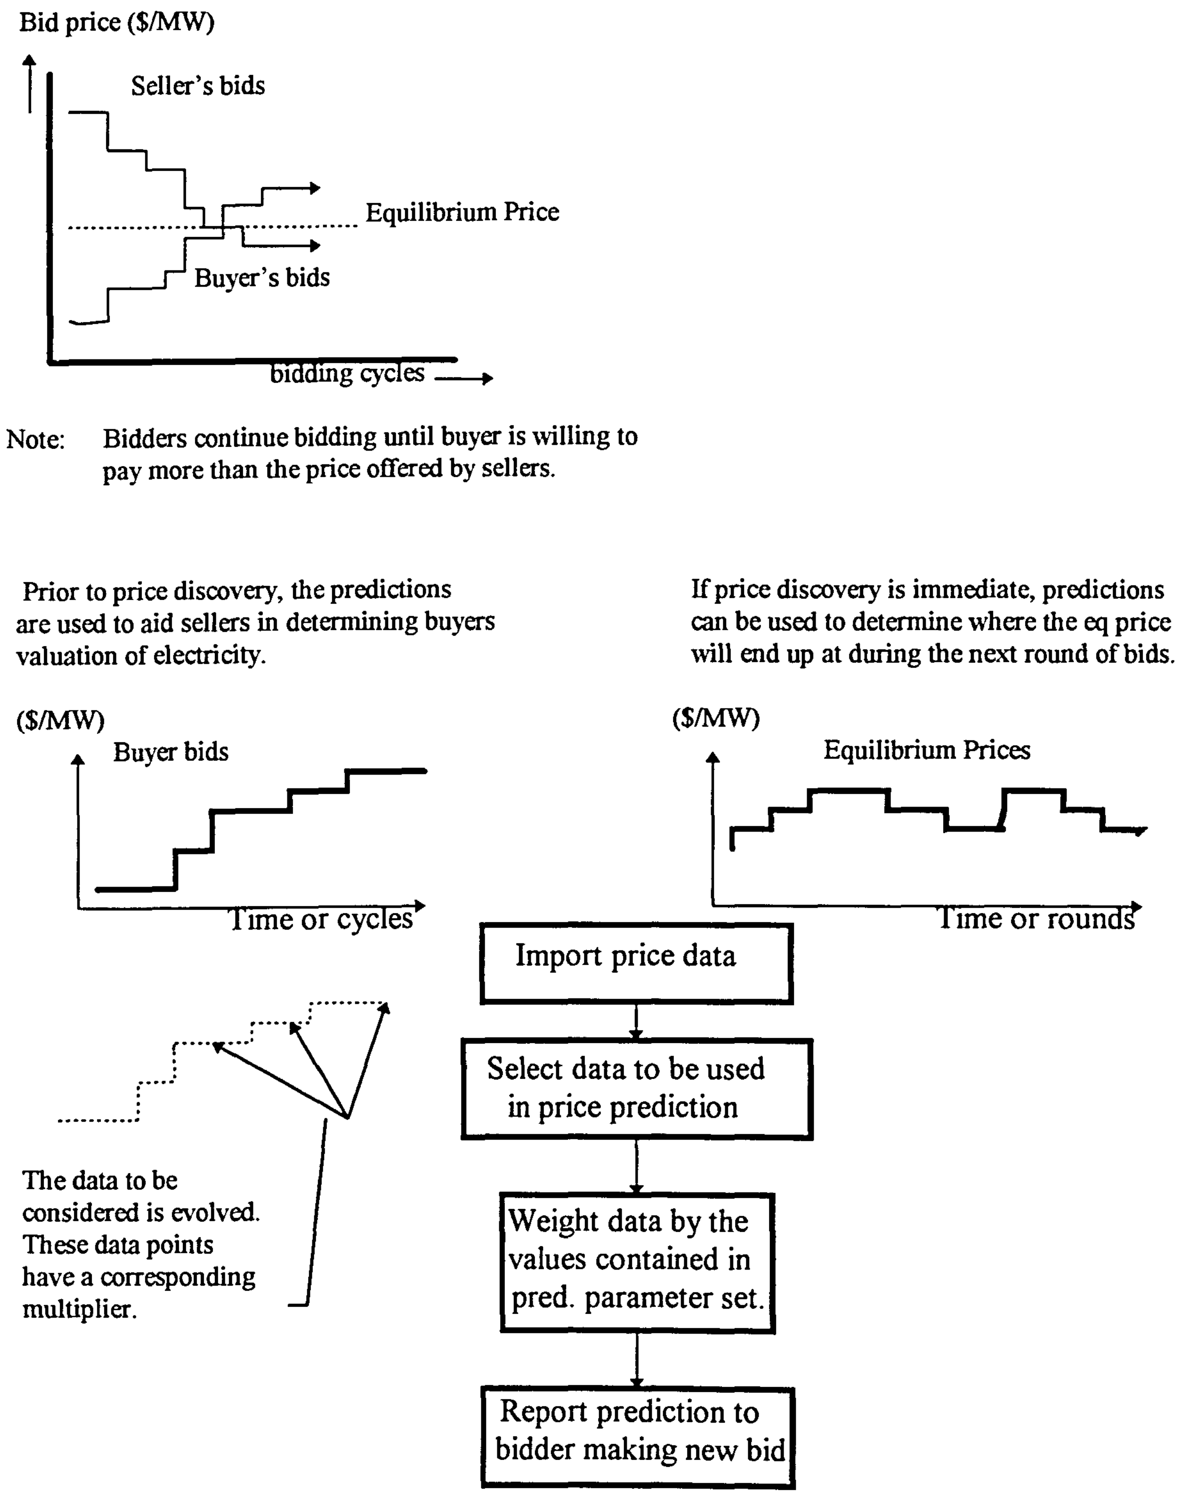

Figure 3.11. Using forecasting techniques to predict equilibrium prices

separately. The third area consists of those functions used to simulate the auction (See Figure 3.10). A block diagram of the evolving agent GA and its interaction with the price prediction GA is shown in Figure 3.12. Figure 3.13 shows the agent GA and its connection to the auction mechanism.

Initialization of gencos

Each of the agents participating in the auction, represents a genco which desires to make a profit via the sale of power. For simplicity, each genco in this research consists of one generating unit that has a minimum and a maximum amount of power that may be produced. These restrictions are due to the physical limitations of generators, turbines and other mechanical equipment that drives the unit. Each generator has a cost associated with producing electricity. This cost at any feasible power level is modeled by the following quadratic curve:

$$F_i = \{A + (B)(P_i) + (C)(P_i)^2\} R_f \qquad \text{(equation 3.4)}$$

where,

F_i = cost of generating unit i

P_i = generation level in MW of generating unit i

R_f = price of the fuel used to power the generating unit

A, B, and C are determined by physical characteristics of the generating unit

The parameters of this curve are read in from file during the initialization portion of the algorithm.

In addition to reading in the cost curve parameters and fuel price for each of the gencos, during the initialization phase of the algorithm each genco has a gene made up by a string of integers that is initialized with the number of contracts to offer at each round of bidding. The reason for this gene is that during the bidding process, as power is sold, demand may vary. If demand is low, the price of power offered for sale must decrease in order to be attractive to buyers. Rather than sell too much power at a low price, the successful gencos will evolve the number of contracts to offer at each round of bidding. A

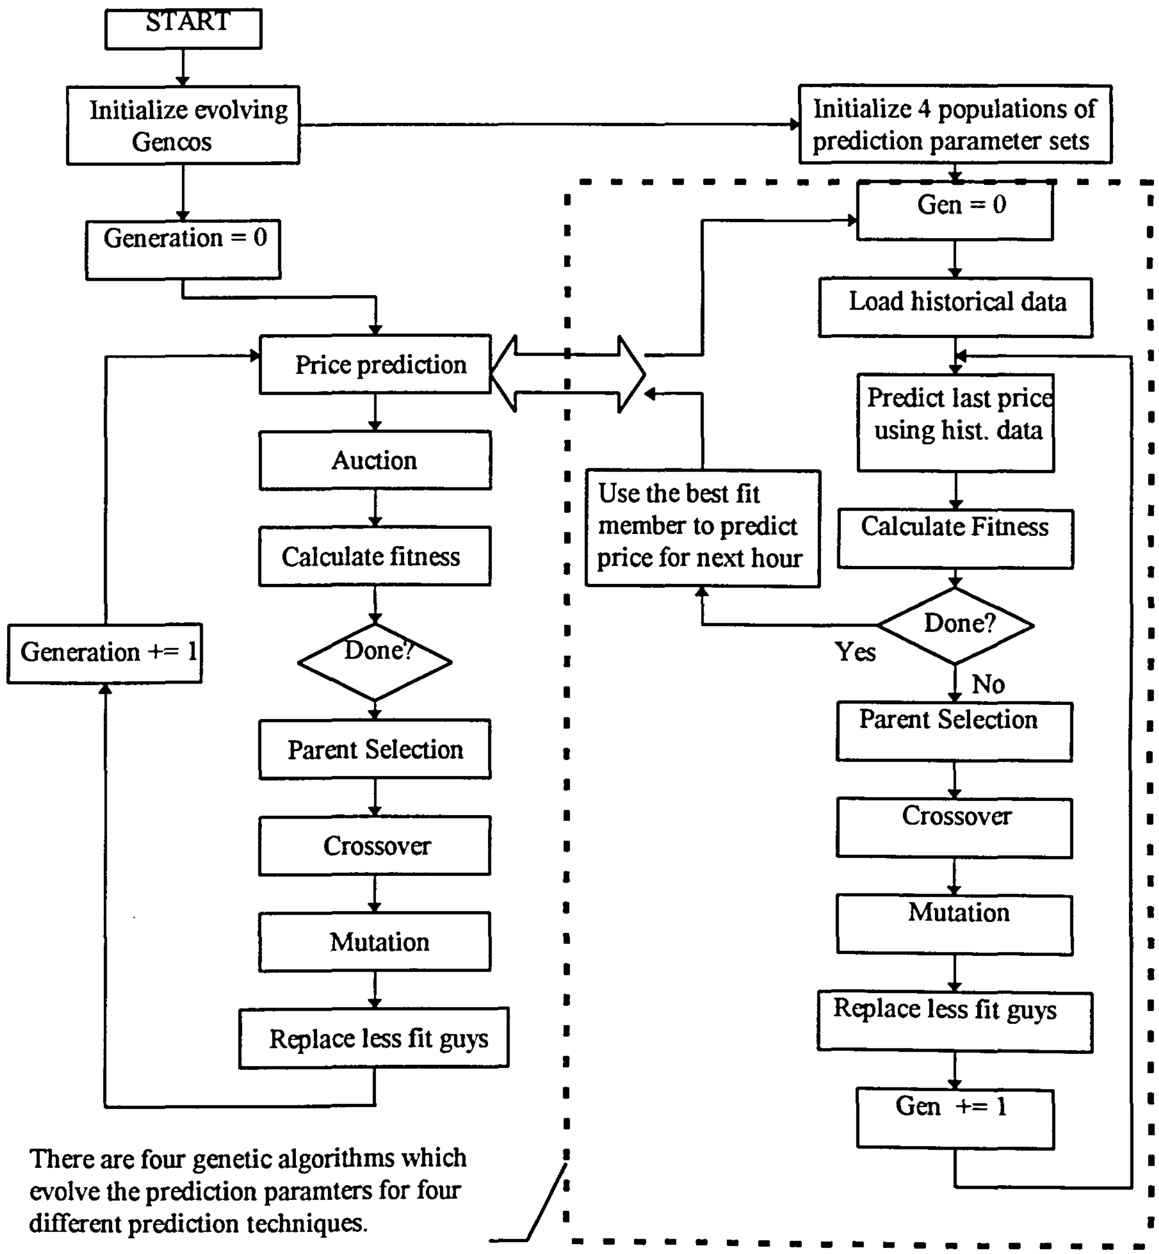

Figure 3.12. Block diagram of the evolving agents GA and interaction with prediction parameter GAs

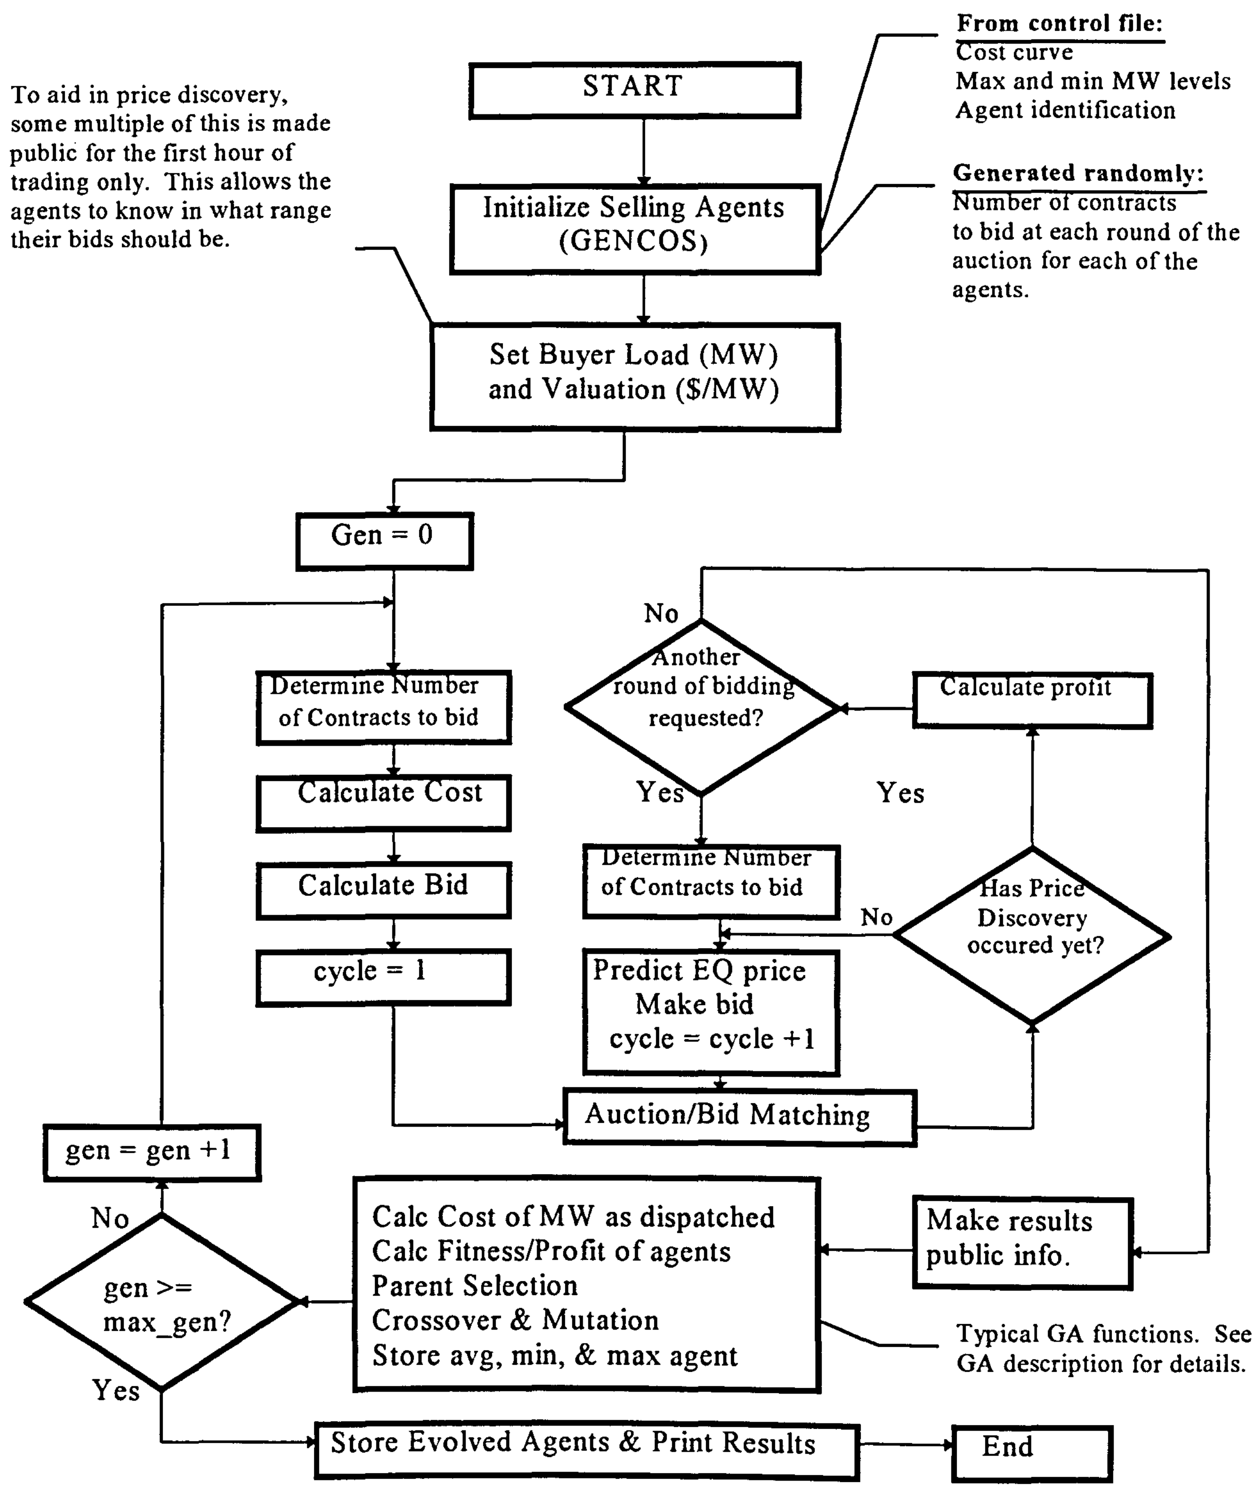

Figure 3.13. Block diagram of evolving agents' GA and auction mechanism

multiple of the cost is used to develop a starting point from which to start bidding. The price prediction method is chosen randomly during initialization. The contracts to offer each round of bidding, the cost multiplier to determine initial bid, and the choice of price prediction method will be evolve while the GA is running. See Figure 3.13 for a graphical representation of these parameters. Figure 3.14 shows how these parameters are used to participate in the auction. Figure 3.15 shows how the agents uses the information in its data structure to develop a bid and participate in the auction.

In the research described in this paper it is assumed that the startup and shutdown costs are high and the generating units have been operating prior to the period of concern. Therefore, a unit commitment analysis is not required at each step to determine whether it is in the best interest of the genco to generate power at that hour. Further research will perform unit commitment analysis including startup and shutdown costs and ramp rates. This will be discussed in the future research section of this thesis.

Buyer valuation and equilibrium pricing

Because this research focuses on the gencos, the buyers are represented by a single disco with a large demand. The disco does not have the same access to the prediction methods that the gencos are using. Instead of a cost curve as the gencos use, the buyer's valuation of electricity at each round of bidding is predetermined using a typical demand curve (in the simulations included in the results section, it is a constant to simplify things). This valuation, which for our purposes can be considered the buyer's bid, is used, in combination with the gencos' bids, to determine the equilibrium price.

For the results presented in Chapter IV, the gencos' bids are developed by one of two methods. The first uses the a bid multiplier ranging from 0 to 1. If the multiplier is 0, the genco bids at its cost. If the multiplier is 1, the genco bids at the expected equilibrium price. The second method uses a bid multiplier ranging from 0.75 to 1.25. The bid multiplier is multiplied directly by the expected equilibrium price. The expected equilibrium was determined via the prediction techniques in some trials, and in the remainder the expected

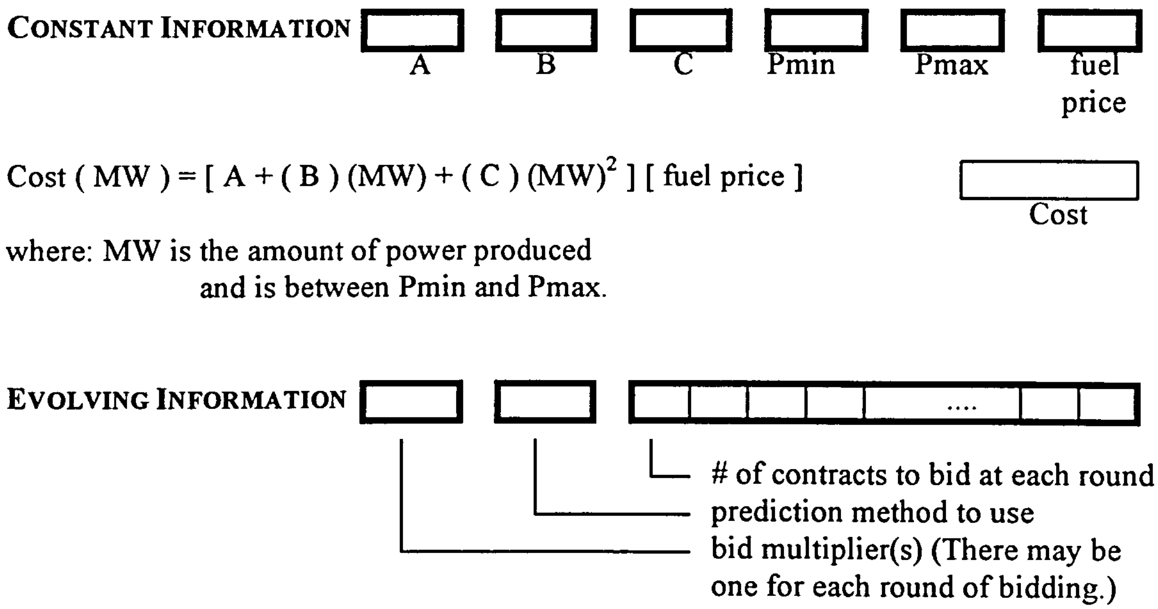

Figure 3.14. Agent's data structure

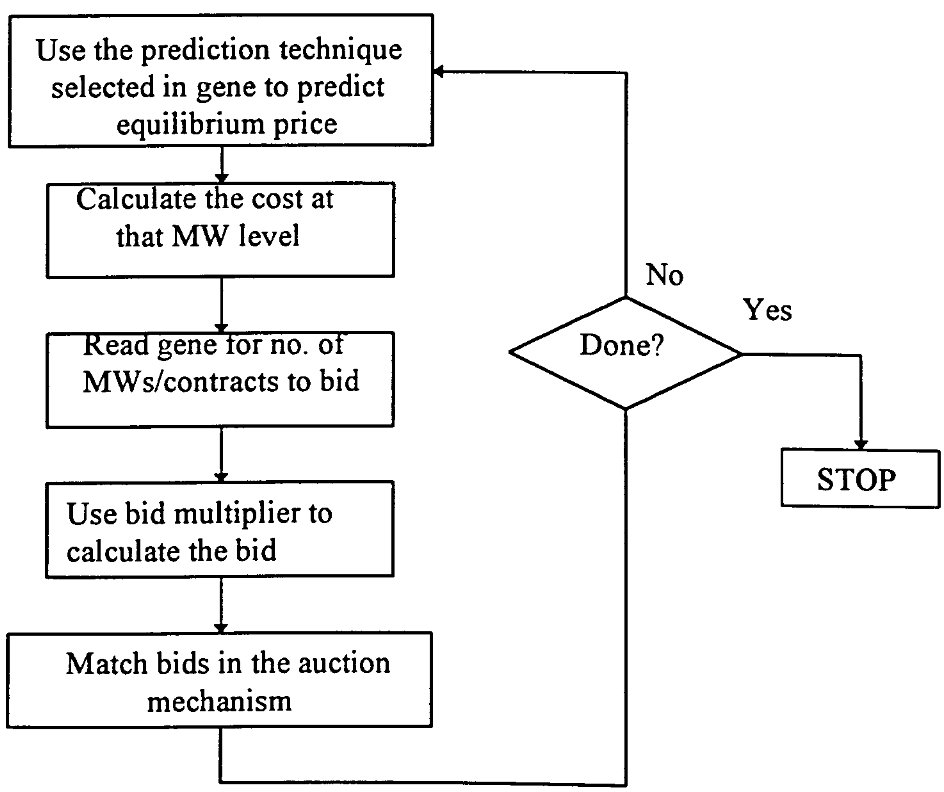

Figure 3.15. Using the information in the agent data structure

equilibrium price was taken to be the previous rounds equilibrium price. The second method involves less coding and would seem to be a fairly accurate assumption in a stable market where prices are not very volatile from round to round. In both cases, the buyer's bid and the seller's bids are sorted and matched, if the buyer's bid is higher than the offer the seller is making, the price difference is split and they have a contract for the amount of power the seller is offering or the amount that the buyer needs, whichever is smaller.

CHAPTER IV. RESULTS

This chapter presents the results of the research obtained using the techniques or methods described in Chapter III. Prior to simulating the auction, the evolution of the prediction parameters developed for this research was tested. Four separate populations were evolved simultaneously. One population each for the parameters necessary to predict prices using linear regression, moving average, weighted moving average, and exponentially weighted moving average. Using the genetic algorithm initialization parameters shown in Table 4.1, the prediction techniques were used to obtain the plots on the following pages. Each prediction technique was given a sinusoidal input with a small amount of noise as an input. The prediction technique predicted the value for the next time step. The results of

Table 4.1. Prediction parameter GA parameters

Parameter	Value	Parameter	Value
population size	16	Xoverpts	3
max_aeta	1	mut_percent	50
number_new	8	parent_selection	tournament
max_placement	5	max_window	5
max_precision	10		

that test confirmed that the prediction parameters were functioning, but also show that the predicted price seemed to lag the actual price by a time period. In Figures 4.1 - 4.2 we can see this. The input curve and the predicted curve are plotted on the same graph. The first graph in Figure 4.1 uses the best moving average prediction parameter set. The second graph in Figure 4.1 uses the best weighted moving average prediction parameter set. Figure 4.2. contains the best prediction parameter sets for the exponentially weighted moving average and the linear regression prediction. In both figures, the dashed line represents the input values, while the solid line represents the predicted values.

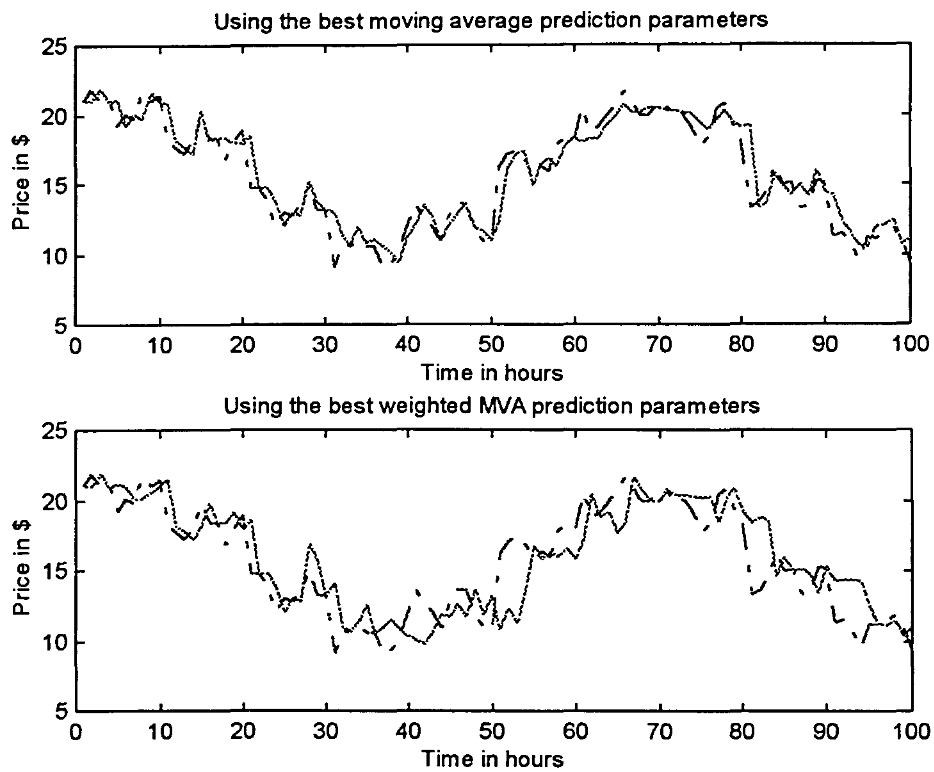

Figure 4.1. MVA and WMVA prediction parameter examples

The best prediction parameter set in each of the four populations performed rather well. The entire population did fairly well, suggesting that perhaps this problem may not require a genetic algorithm to attempt to optimize the prediction parameters. In future research, a fixed prediction parameter set could be used, or perhaps the prediction portion of the algorithm could be eliminated. The population sizes for each of the four prediction parameter populations were all set to sixteen because of memory limitations.

Auction Simulations

This section contains several cases of the evolving agent auction simulations. Each case contains a plot of the gencos' average, maximum and minimum bids, the equilibrium price, the amount of power desired by the buyer, the amount of power actually purchased by

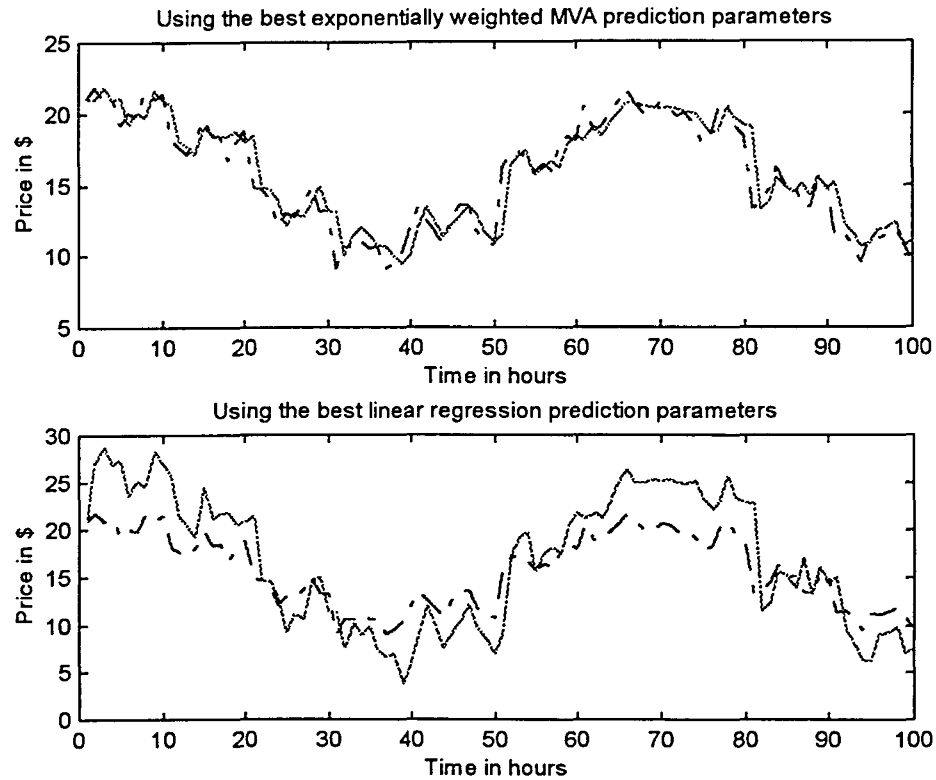

Figure 4.2. EWMVA and linear regression prediction parameter examples

the buyer, the minimum, maximum and average fitness/profitability of the population of agents at each generation, and the amount of power sold by the most profitable generation company in each generation.

For the cases described in this chapter, different combinations of features of the program were used to develop the output. In some of the cases a separate bid multiplier for each round of bidding was used. In other cases a single bid multiplier was used for all of the rounds of bidding. The bid multiplier allows each of the agents to bid an amount greater than or equal to their cost and lower than or equal to the expected equilibrium price for some trials, while for others the bid multiplier allows the agent to bid slightly above or below the genco's expected equilibrium price. The case description describes which bid multiplier scheme is being used, and whether the equilibrium price is being predicted by one of the four techniques.

The expected equilibrium cost was determined in some of the cases using the agents' choice of prediction technique. As mentioned in Chapter III, the agents have their choice of linear regression, moving average, weighted moving average, and exponentially weighted moving average. In other cases, the expected equilibrium price was taken to be the previous rounds equilibrium price, except for the first round of bidding when the expected equilibrium price was taken to be a multiple of the generator's cost. The latter method of determining the expected equilibrium price is easier, and might be a fairly good assumption if the number of players is large, and each player is a small part of the total trade. This assumes that market has stabilized to some steady state value and is not expected to fluctuate a great deal.

For the evolution of the agents, three point crossover was used for both the number of contracts each round gene, and the bid multiplier gene. Two types of mutation were used. Mutation A would, with a certain probability (mutation rate), select each loci of the newly created creatures for mutation. If a loci in the integer portion of the gene (i.e. the number of contracts) was selected to be mutated, a number selected randomly from a uniformly distributed set of valid numbers was added to the value in that loci. If larger than allowed, it was wrapped around to the beginning. If a loci in the bid multiplier section of the gene was selected for mutation, the value in that loci would be flipped (i.e. 1 becomes a 0, and vice versa). Mutation B, with a probability of mutation rate, would shuffle the number of contracts (one number for each round of bidding) within an agent's gene. This would aid in spreading good numbers to all the rounds rather than waiting for mutation A to stumble onto these numbers. Mutation A or mutation B was each called each generation, each with a 50% probability. The number of new creatures each generation was set at one half of the population.

Each of the agents have the same cost of generation. The cost can be defined as a quadratic equation as described in Chapter III. For each of the generating agents, the cost is $200 + (8)(\text{MW level}) + (.00251)(\text{MW level})^2$ [$]. If we take the derivative of this with respect to MW, we get the incremental cost curve, which is $8 + (2)(.00251)(\text{MW level})$ [$/MW].

Case A description

For case A, the bid multipliers for each round of bidding were initialized to 1. The supply was set much greater than the demand. No prediction techniques were used, the expected equilibrium price was taken to be the previous bidding round's equilibrium price. There were independent bidding multipliers for each round of bidding. See Table 4.2 for the GA parameter setpoints. See Figures 4.3 and 4.4 for the results of the simulation.

Table 4.2. Parameter setpoints for case A

Parameter	Value	Parameter	Value
generations	40	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	1:10
popsize	24	Buyer bids ($)	20
number new	12	Case ID	A

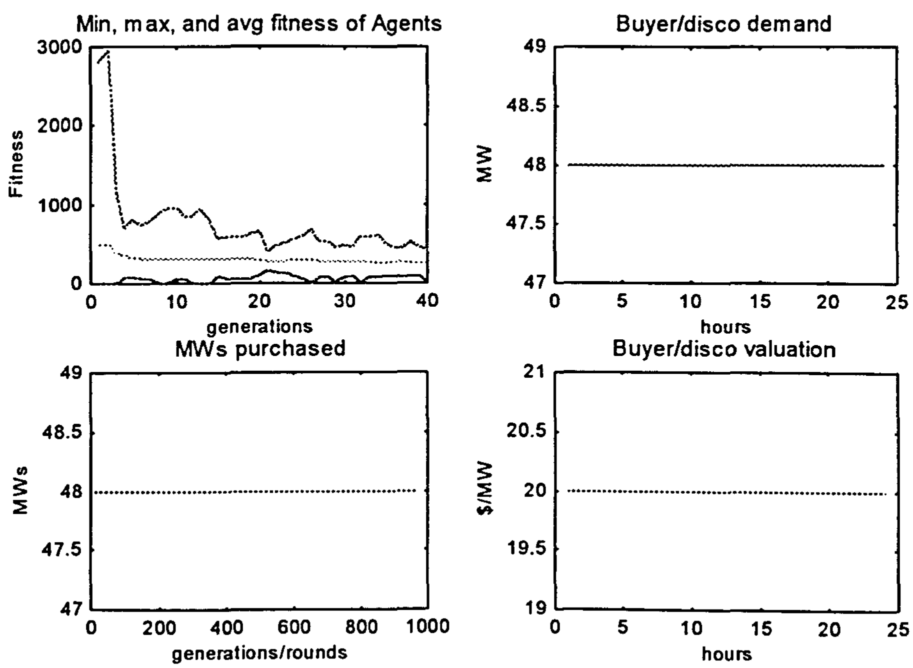

Figure 4.3. Case A plots

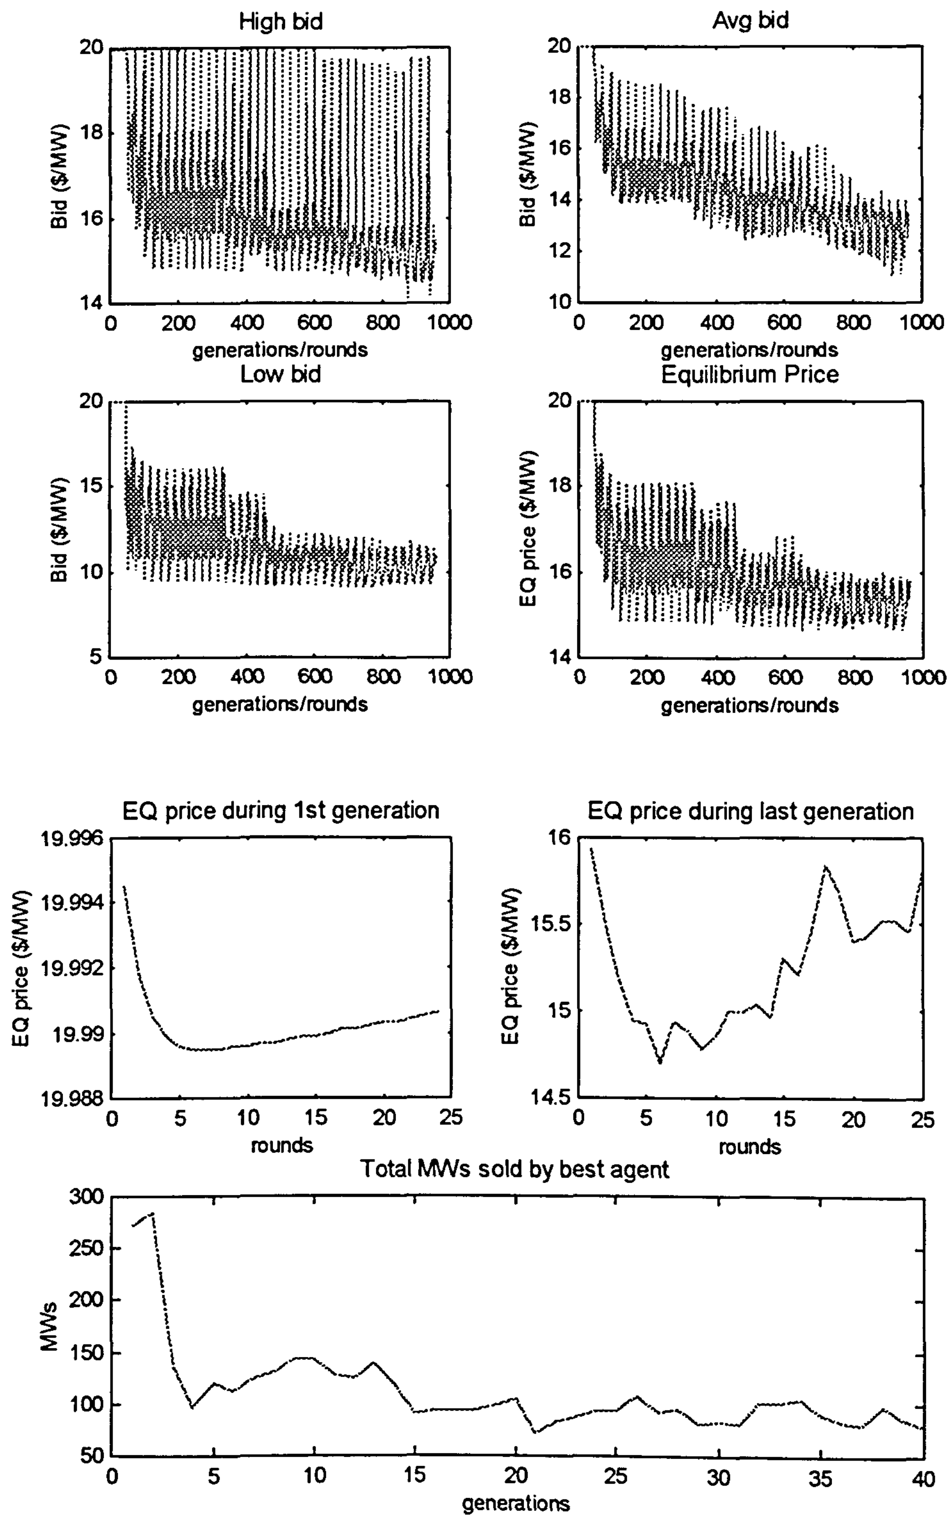

Figure 4.4. Case A plots (continued)

Case A analysis

For case A the bid multipliers started out at one, meaning that the bids were at the expected price. After only a couple of rounds, the bids decreased significantly. This is because when multipliers less than one are introduced via mutation, these agents quickly steal the contracts from those bidding with their bid multiplier of one. The supply is ten times the demand, and there is no room for high prices. The maximum fitness occurs only during the first round of bidding when all sellers are bidding the same inflated price. After that, the sellers' profits drop to a low level and they stay there.

Case B description

For case B the bid multipliers for each round of bidding were changed to allow bidding above the previous equilibrium point. The maximum bid allowed for this case was 1.10 times the difference between the cost and the previous equilibrium price. The supply was again set much greater than the demand. No prediction techniques were used to obtain expected equilibrium price. Independent bid multipliers were used for each round of bidding. The GA was allowed to evolve for 400 generations this time. See Table 4.3 for the parameter setpoints. See Figures 4.5 and 4.6 for the resulting output plots.

Table 4.3. Parameter setpoints for case B

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	1:10
popsize	24	Buyer bids ($)	20
number new	12	Case ID	B

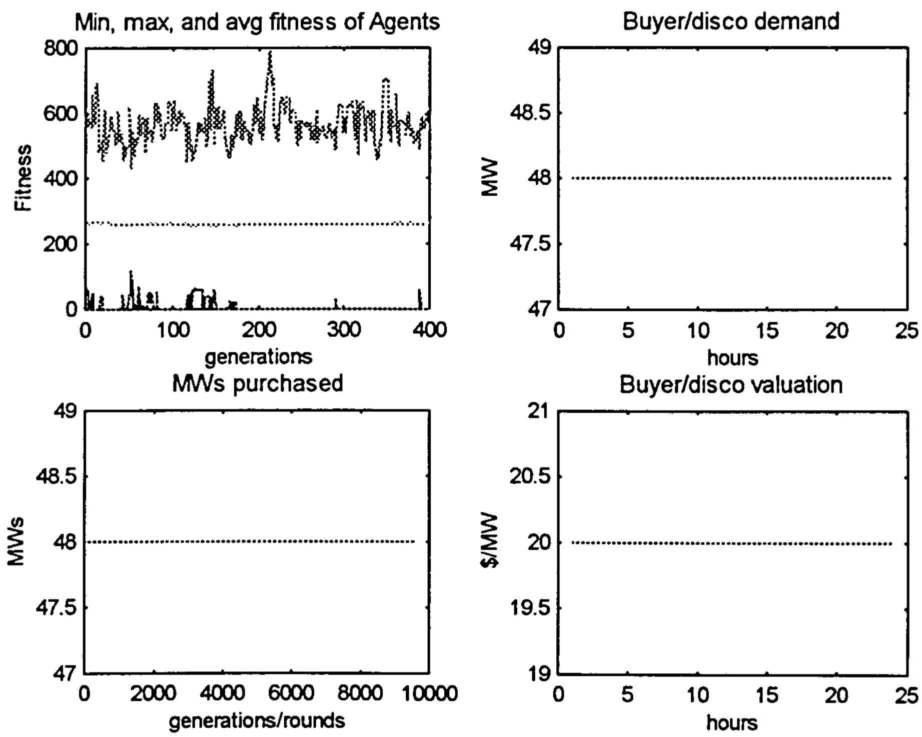

Figure 4.5. Case B plots

Case B analysis

In the figures above we note several things. First of all, the minimum, maximum and average fitnesses do not seem to have any definite trends upward or downward. In effect, there is little that the agents can do to increase their profit. Demand is much smaller than the supply, and if the sellers increase their bids even slightly, they are left with no contract. Since it is possible for sellers to bid over the expected price, some of them take this opportunity to bid more than what the buyer is willing to pay resulting in no contract. This is possible since for the first round of bidding each generation, the forecasted or expected equilibrium price is a multiple (1.5 in this case) of the generation cost.

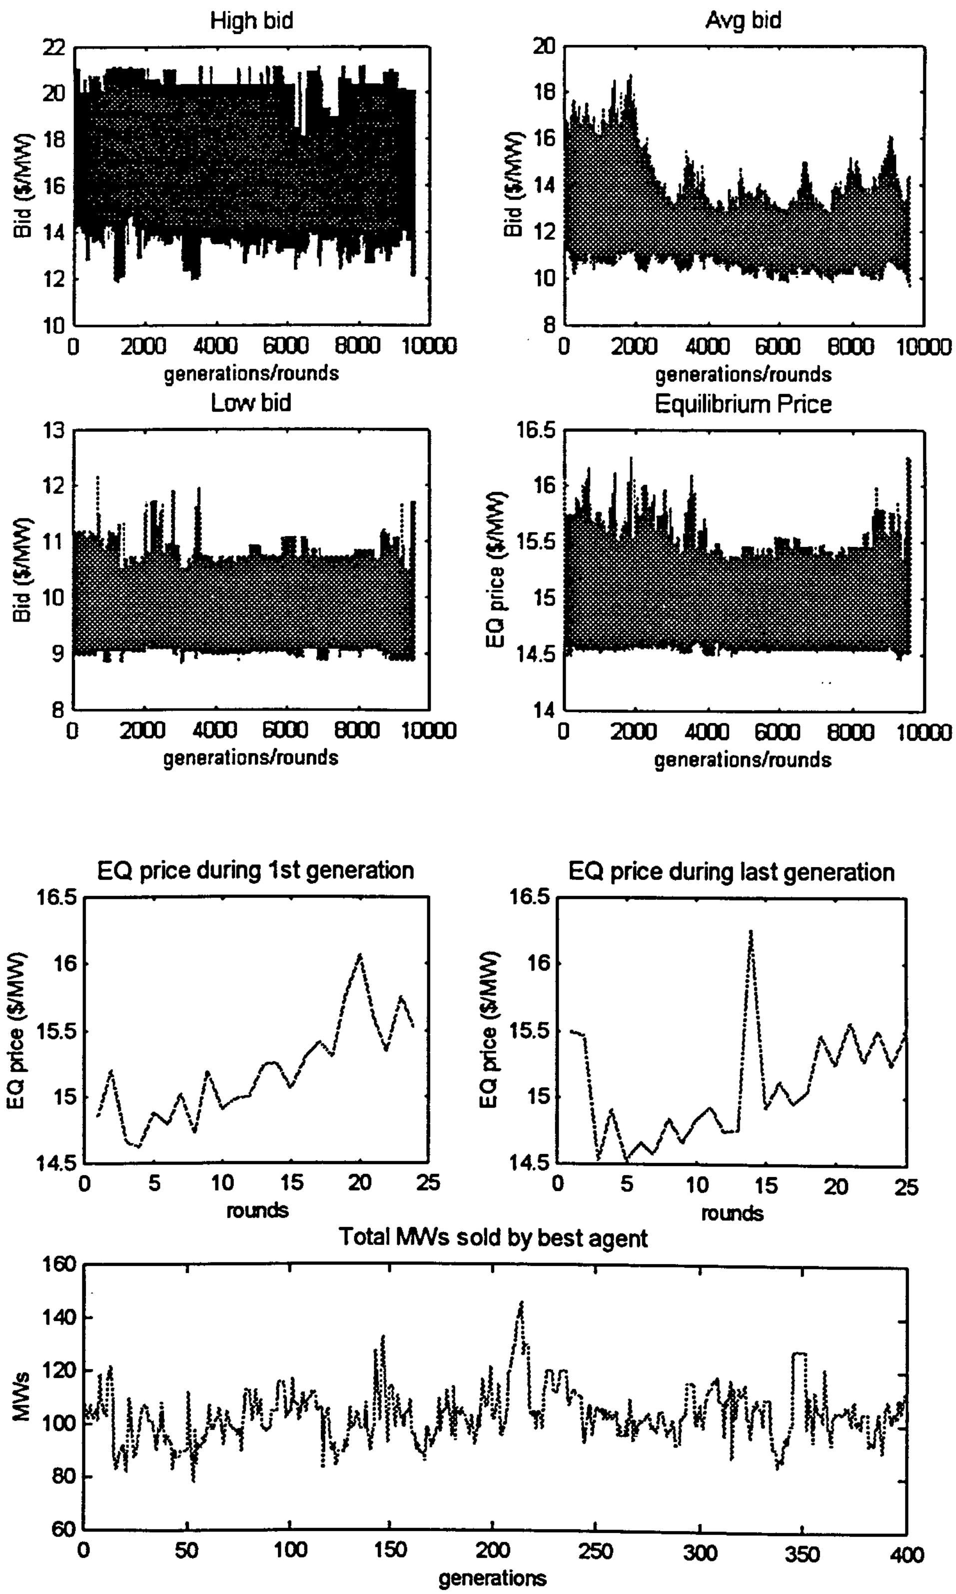

Figure 4.6. Case B plots (continued)

Case C description

For case C the bid multiplier was returned to its original range, so that the previous round's equilibrium point was its maximum bid as in case A. In case C, the demand was set much higher than supply, so that rather than having a surplus of electricity, there was twice as much power demanded as supplied. Since the agents evolve both their bid multipliers and the number of contracts to offer for sale at each round of bidding, they should increase the amount of electricity offered for sale. This case does not use a prediction technique to determine the expected price and has a separate bid multiplier for each of the bidding rounds. See Table 4.4 for the parameter setpoints. See Figures 4.7 and 4.8 for the resulting output plots.

Table 4.4. Parameter setpoints for case C

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	2:1
popsize	24	Buyer bids ($)	20
number new	12	Case ID	C

Case C analysis

As expected, the average and maximum fitness exhibited by the agents increased over the generations. This is because the gene choosing the number of contracts (each contract is equivalent to one MW) offered for sale at each round of play evolves to its maximum. The maximum amount that any genco can sell during any one round is the total number of MWs that their unit is capable of producing divided by the number of rounds. With the maximum bid being the previous round's equilibrium price, it is hard for the gencos to fully exploit the buyer's predicament. The equilibrium price approaches the buyer's bids rather slowly, especially since new agents have a tendency to bid randomly.

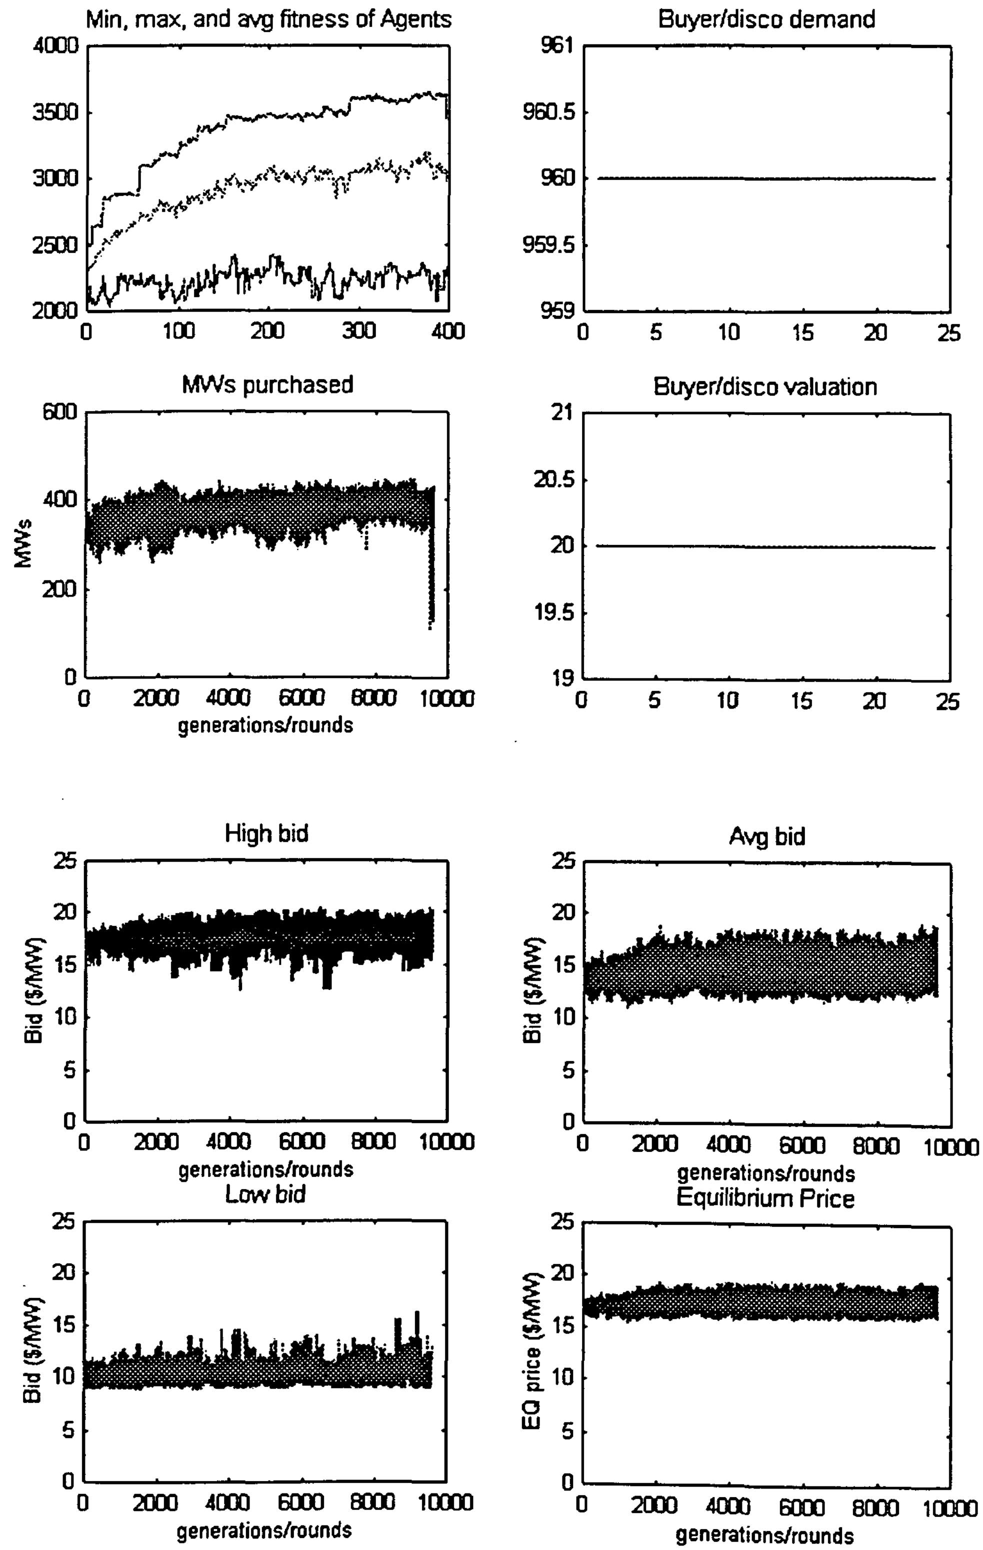

Figure 4.7. Case C plots

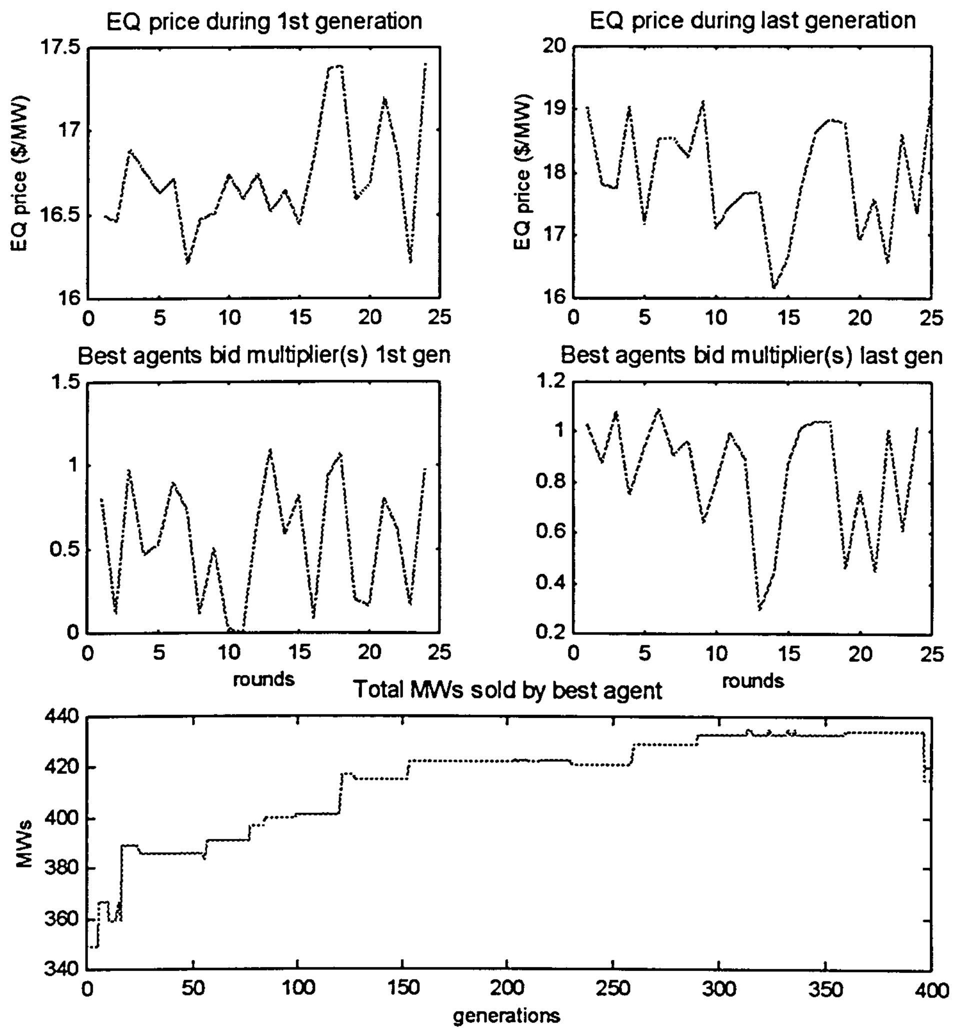

Figure 4.8. Case C plots (continued)

Case D description

For case D, the maximum bid that the gencos are allowed to make is again changed so that it is not limited by the previous rounds equilibrium price. This case should demonstrate whether the gencos more quickly take advantage of the powerless buyer. In case B this was tried, but the supply was much greater than demand. The parameters are set

4.5. Parameter setpoints for case D

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	2:1
popsize	24	Buyer bids ($)	20
number new	12	Case ID	D

the same as in case C and are shown in Table 4.5. No prediction technique is used to **Table** determine the expected equilibrium price for this case. Each round of bidding has a separate bid multiplier. See Figures 4.9 and 4.10 for the resulting plots for this case.

Case D analysis

From the graphs we can see that the sellers are not able to fully exploit the buyer

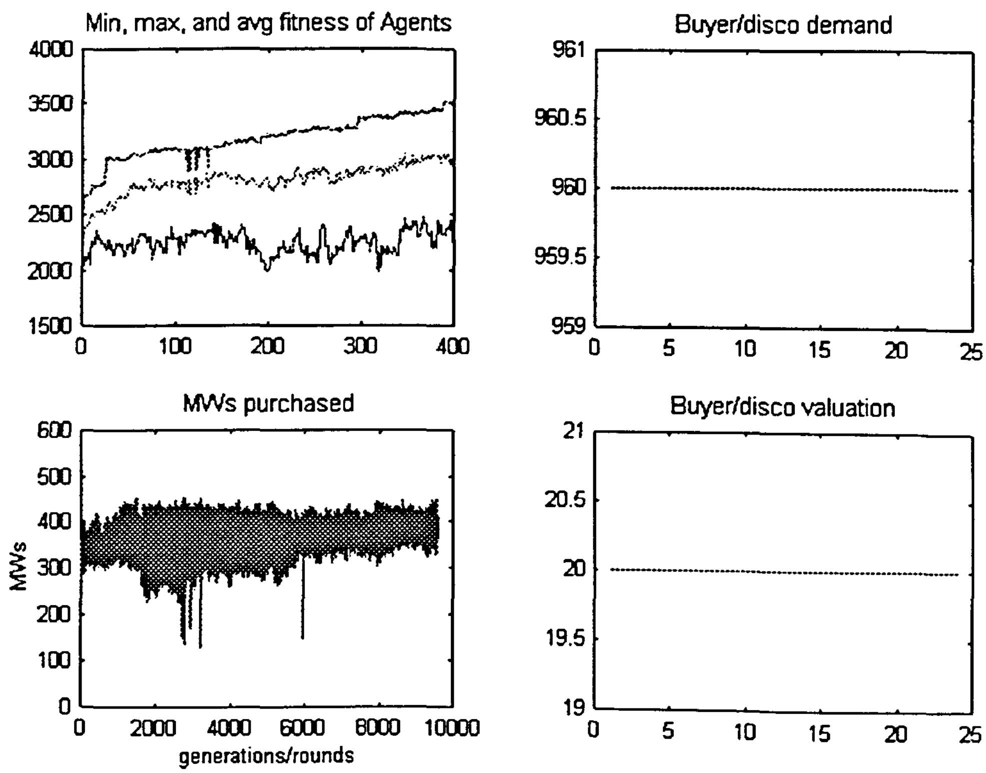

Figure 4.9. Case D plots

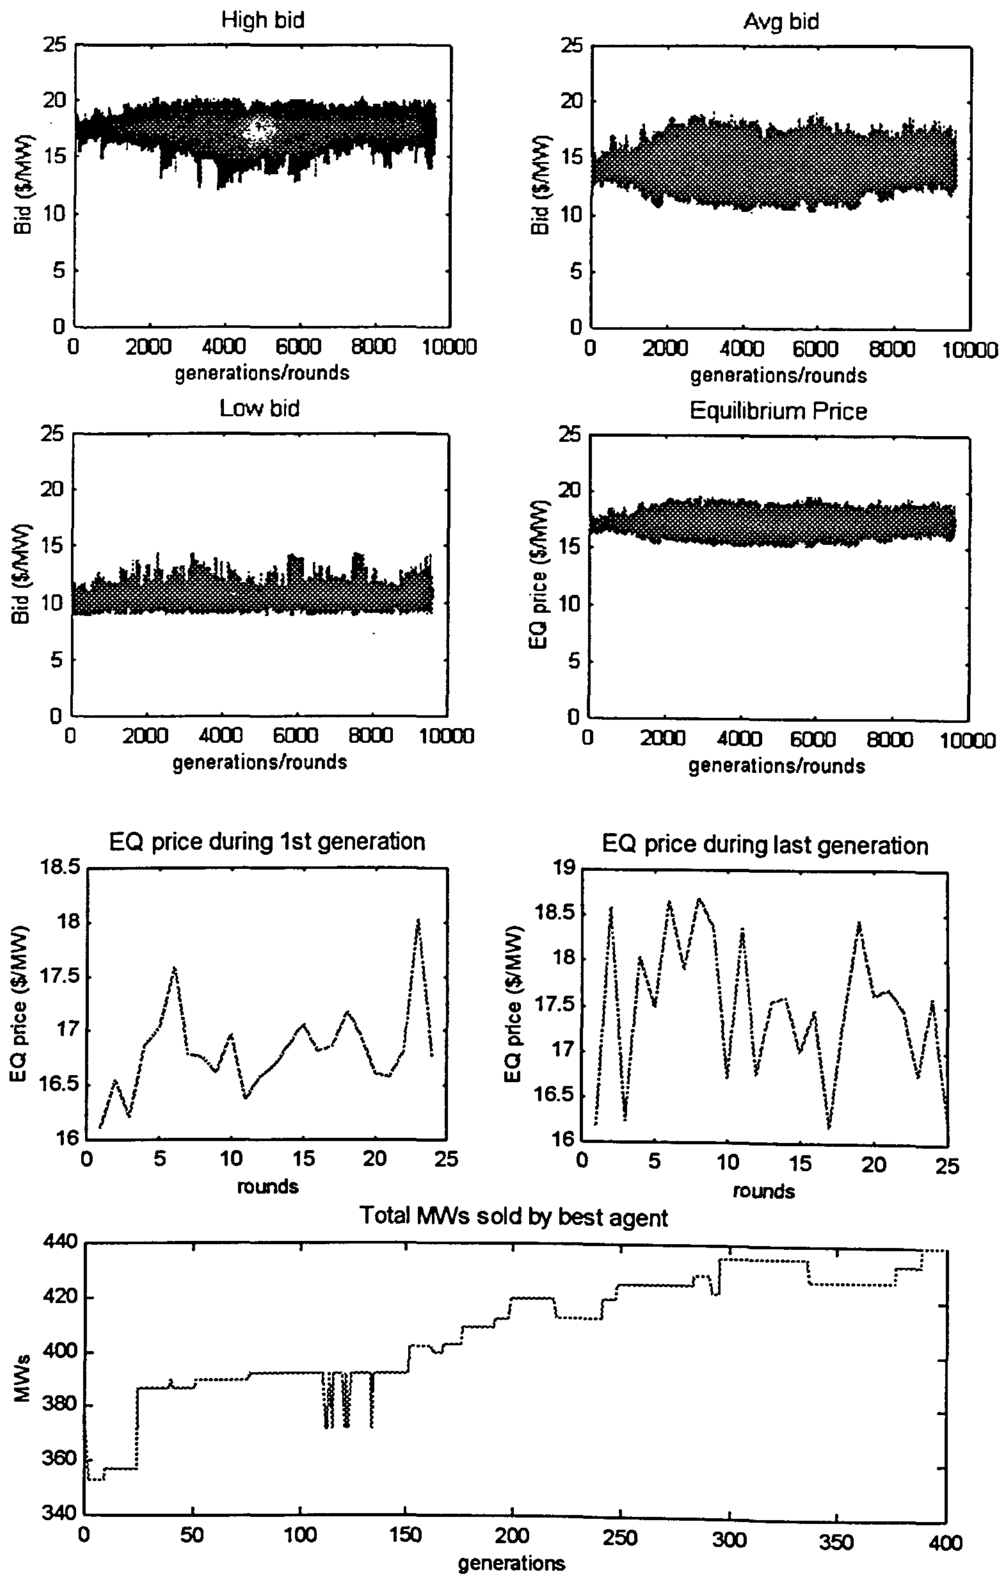

Figure 4.10. Case D plots (continued)

using this bid multiplier that is dependent on the equilibrium price. When they set their bid multipliers to bid higher than the equilibrium level, the equilibrium price tends to go beyond that which the buyer is willing to pay resulting in no contract. If all are doing this, the equilibrium price drops suddenly to zero. It is possible to see from the figures that the equilibrium price seems to go to about the $18 dollar range, where $20 is the maximum that the buyer is willing to pay.

Case E description

Since it is hard to tell where if the bid multiplier evolution is actually aiding the agents, for case E the number of contracts are fixed at the maximum for all rounds of bidding. This means that any improvement in profit will be due strictly to better bid multipliers. The parameters for this case are shown in Table 4.6. No prediction techniques are used to develop the expected equilibrium price. Separate bid multipliers are used for each round of bidding. See Figures 4.11 to 4.12 for graphical results.

Table 4.6. Parameter setpoints for case E

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	1:1*
popsize	24	Buyer bids ($)	20
number new	12	Case ID	E

* In this case, the number of contracts was fixed at maximum.

Case E analysis

From the graph of the profits or fitnesses of the agents versus the generation, it is observed that the bid multipliers have evolved and do have influence over how much profit the agent is making. From the cases that have been run, it can be noted that the evolution of

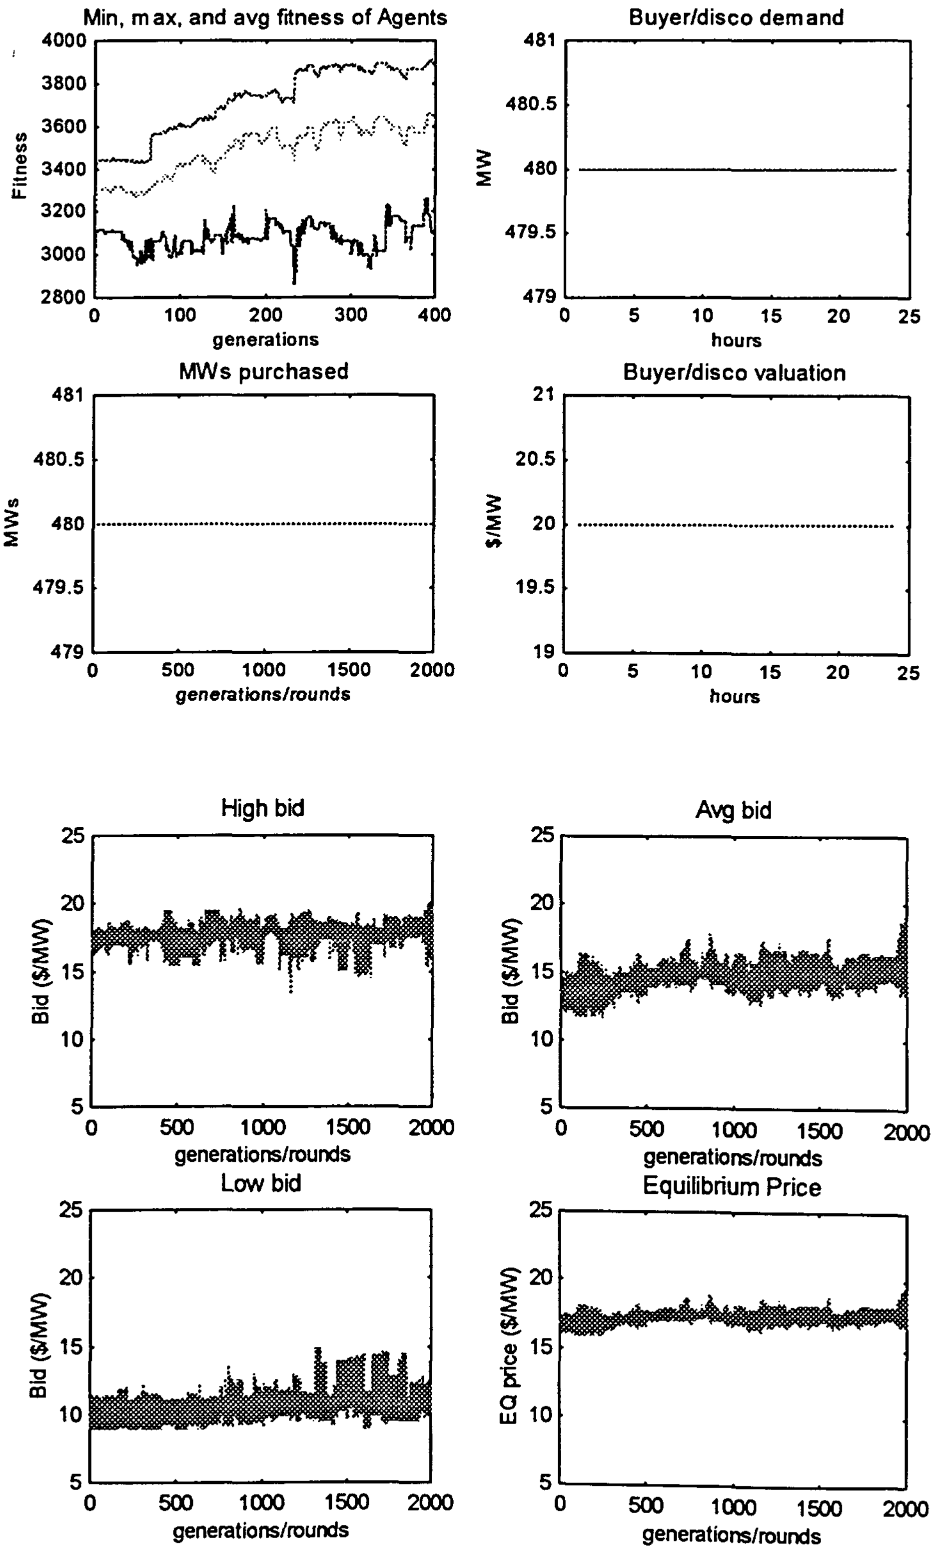

Figure 4.11. Case E plots

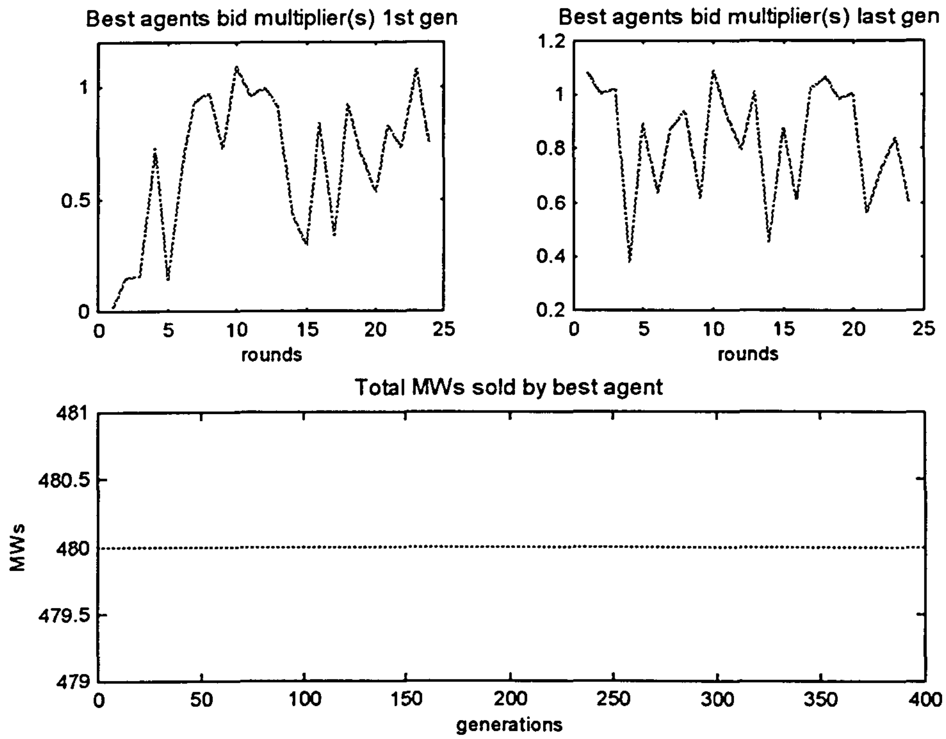

Figure 4.12. Case E plots (continued)

the bid multipliers is slower than the evolution of the contract amount gene when there is a large demand in relation to the supply of electricity.

Case F description

For this case, demand for electricity was set to two times the supply. The agents were allowed to choose a prediction technique to determine their expected equilibrium price. Multiple bid multipliers were used, i.e. a separate one for each round. Their bids were made by multiplying the bid multiplier by the expected equilibrium price. The bid multipliers ranged from 0.75 to 1.25, as opposed to the previous trials where they ranged from the cost to the expected equilibrium price. See Table 4.7 for the parameter setpoints and Figures 4.13 and 4.14 for the graphical results.

Case F Analysis

Although the fitness increases over the generations, the number of MWs actually sold to the buyer is well below the amount that the buyer would like to have. A number of agents are bidding well above what the buyer is willing to pay. The predicted equilibrium price is

Table 4.7. Parameter setpoints for case F

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	2:1
popsize	24	Buyer bids ($)	20
number new	12	Case ID	F

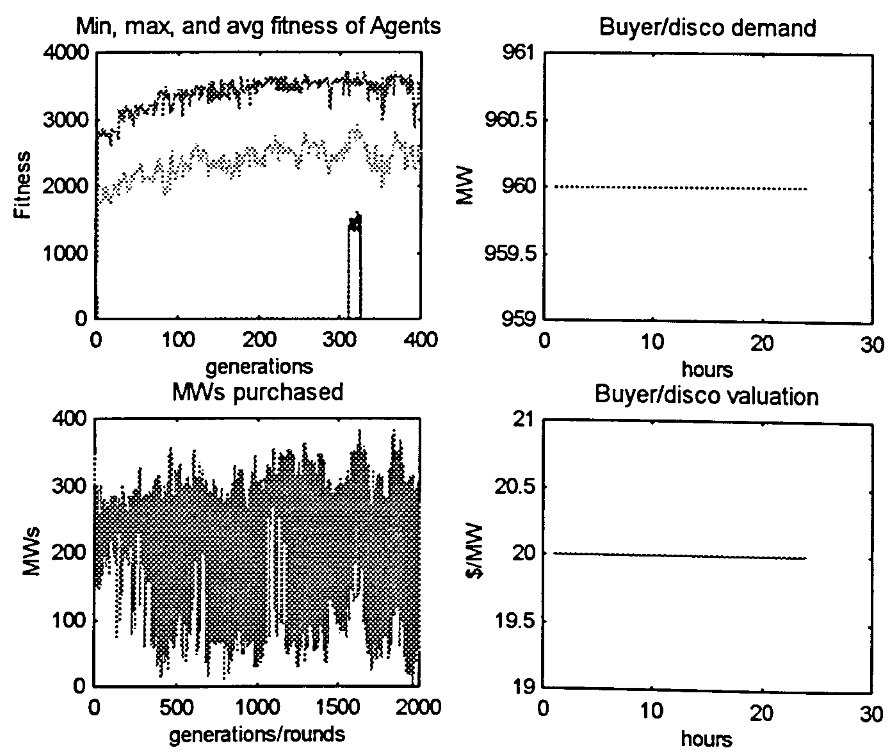

Figure 4.13. Case F plots

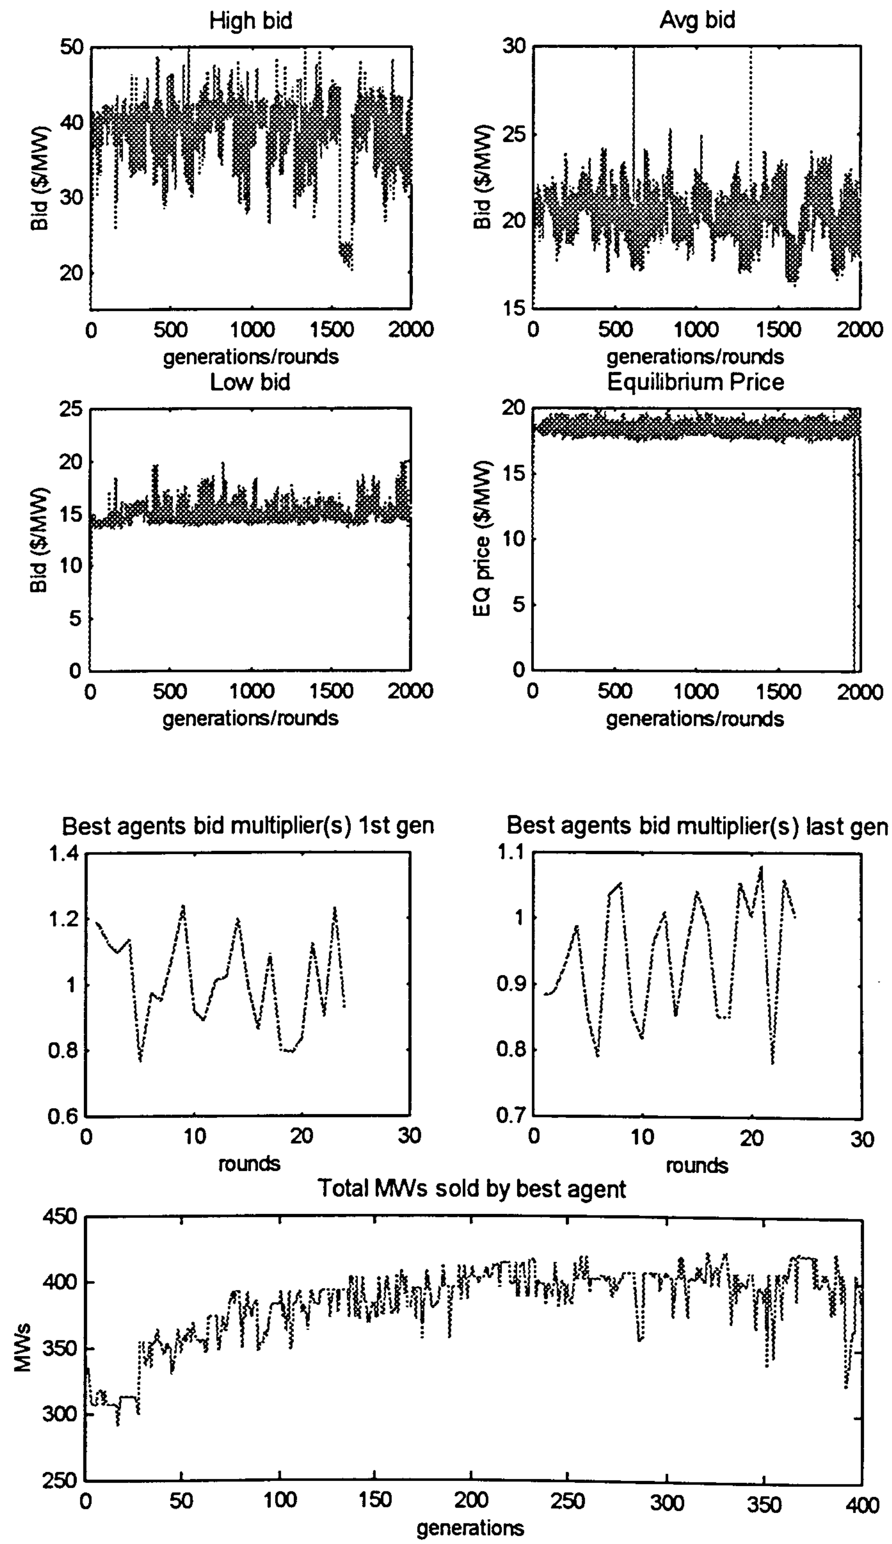

Figure 4.14. Case F plots (continued)

much to high to obtain a contract with the buyer. The best agent is selling less than half of the MWs that it could if the predicted prices were better.

Case G description

For case G, the scenario is the same as in case F, except that this time the demand is one half of the supply. See Table 4.8 for the parameter setpoints, and Figures 4.15 and 4.16

Table 4.8. Parameter setpoints for case G

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	1:2
popsize	24	Buyer bids ($)	20
number new	12	Case ID	G

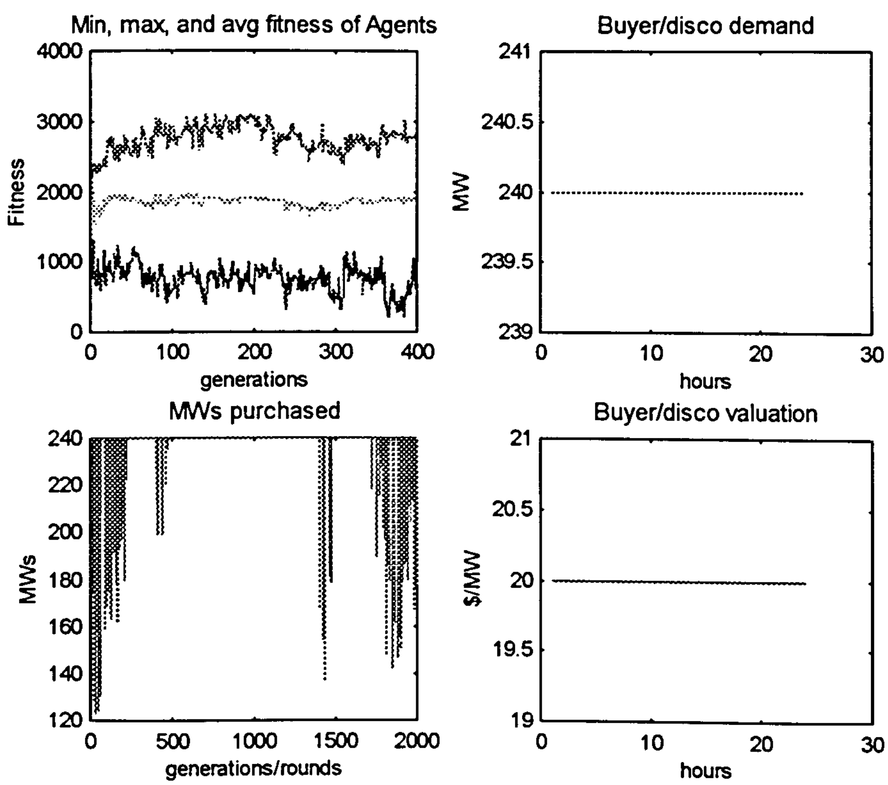

Figure 4.15. Case G plots

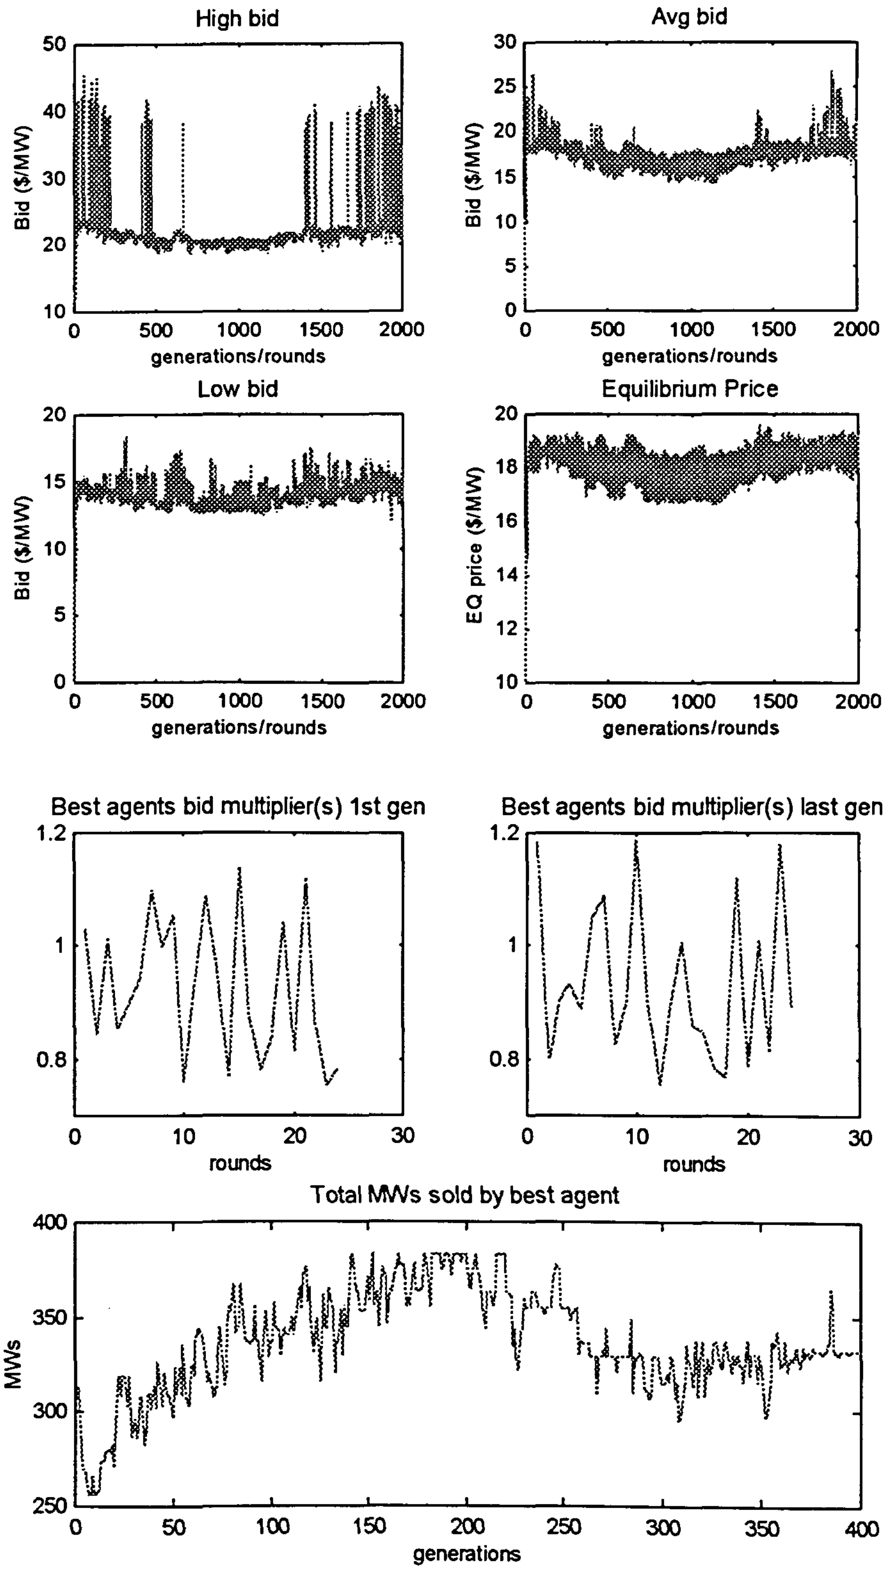

Figure 4.16. Case G plots (continued)

for the graphical results. Multiple bid multipliers are used again, and agents are given their choice of prediction parameters.

Case G analysis

From the graphs we can see that the fitness is a little lower this time as would be expected when the demand is less. The equilibrium price however is staying up higher than we might anticipate. This is due to the number of gencos bidding too high. The lowest bids are around $15/MW which is not even close to cost. The prediction techniques selected are producing optimistic price projections.

Case H description

For case H the demand was set to two times the supply of electricity. Agents were given their choice of prediction techniques to determine the expected equilibrium price. A single bid multiplier was used in combination with the predictions to determine the bids at each round of bidding. Parameters were set as shown in Table 4.9. See Figures 4.17 and 4.18 for the graphical results.

Case H analysis

The agents who actually got contracts in the auction for case H seemed to perform rather well. However, the amount of MWs actually purchased indicates that again the many

Table 4.9. Parameter setpoints for case H

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	2:1
popsize	24	Buyer bids ($)	20
number new	12	Case ID	H

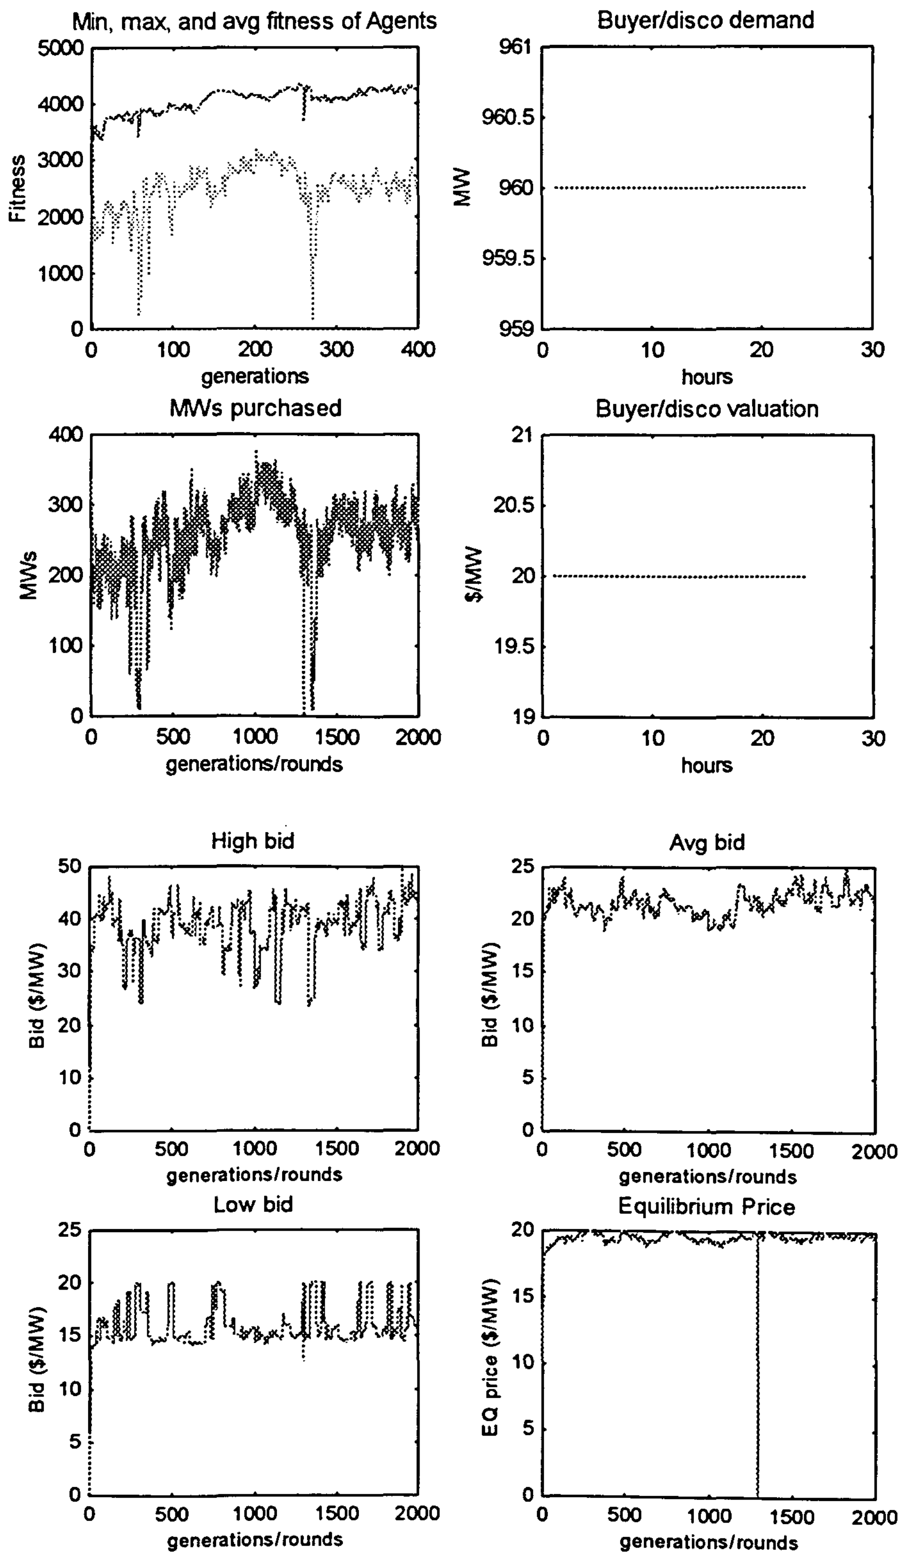

Figure 4.17. Case H plots

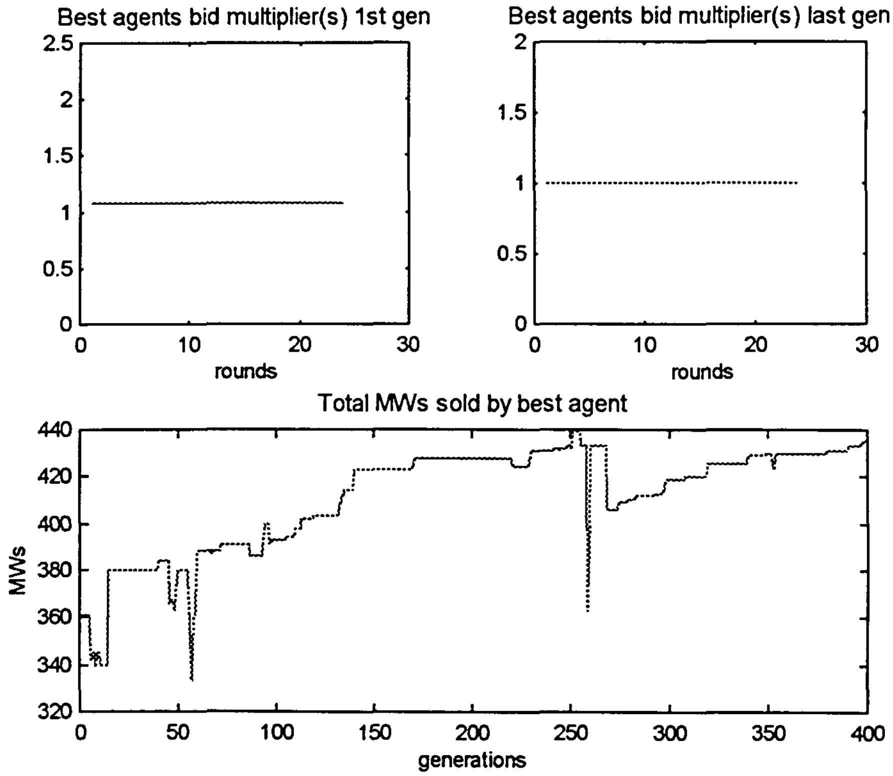

Figure 4.18. Case H plots (continued)

of the agents are bidding to high, and consequently there are several agents who end up without contracts for certain rounds of bidding. The equilibrium point is right at the buyer's evaluation. The best agent sells close to his full amount which means that the contract gene is evolving.

Case I description

For case I, the demand is decreased so that it is one half of the supply. A single bid multiplier is used along with the price prediction method selected by the agents to develop bids. The parameters are as shown in Table 4.10, and the results are as shown in Figures 4.19 and 4.20.

Case I analysis

As the graph shows, the fitness of the selling agents has a downward trend. Over the generations, agents are undercutting each other's prices in order to get the contracts with the buyer. The buyer is getting all the power that it desires. As the sellers learn that they must lower their prices to survive, it's interesting to note that the best agent is actually selling the majority of his generation. This goes back again to the fact that many of the agents are bidding way too high, and consequently those that are bidding under the buyer's bid, are able to sell power for a profit.

Table 4.10. Parameter setpoints for case I

Parameter	Value	Parameter	Value
generations	400	multiplier precision	10
rounds of bidding	24	mutation rate	20
maximum bid cycles	20	Demand:Supply ratio	1:2
popsize	24	Buyer bids ($)	20
number new	12	Case ID	I

General Observations

The equilibrium prices and the bids, due to their link with the prices, fluctuate over time. In fact we see more price movement within the rounds of bidding than we do throughout the generations. In fact, if only the equilibrium prices obtained at a specific round of bidding each generation are plotted, a much smoother curve will result.

Using this bid multiplier range between the cost and the equilibrium price for the gencos while the disco keeps its price fixed will result in a equilibrium price somewhere between the half way point and the fixed price that the disco is willing to pay. If both sides of the auction were using this same strategy the equilibrium price would tend to hang around the midpoint, or at least would be bounded by the buyer's valuation and the seller's cost. Although this shows that the auction scenario can arrive at an equilibrium price at the

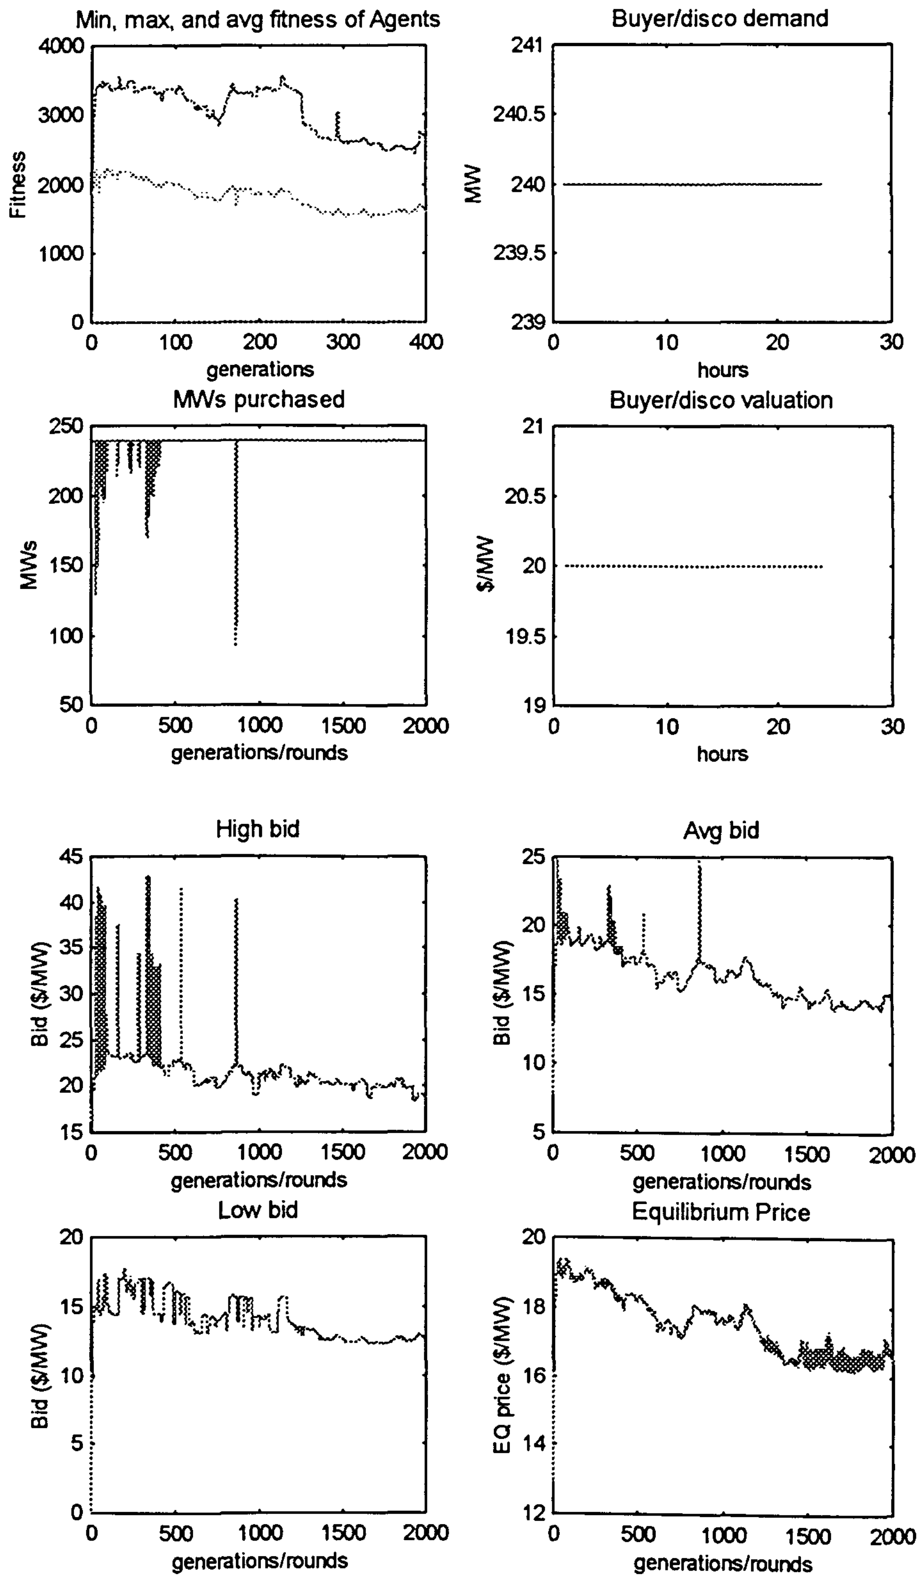

Figure 4.19. Case I plots

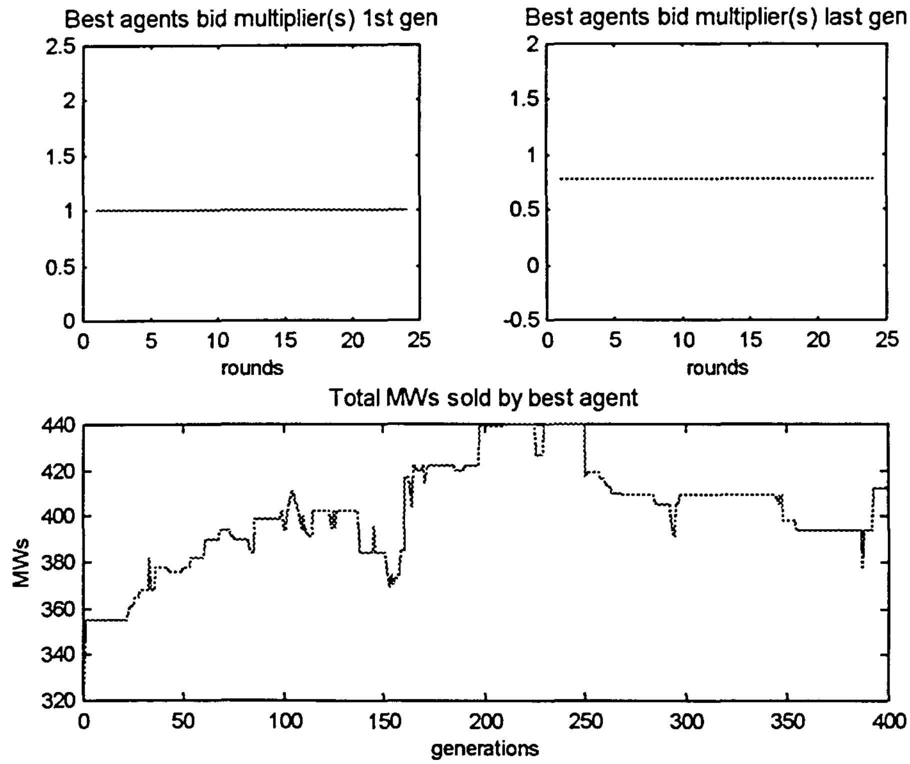

Figure 4.20. Case I plots (continued)

midpoint of the bids, this assumes that prices are not largely affected by the actions of a single market player. This may be a fairly good assumption.

The other method of making bids that multiplies the expected price times a bid multiplier requires that agents have access to a good forecast equilibrium price. The predicted prices obtained by the prediction techniques selected by many of the agents seemed to be much too high. Those agents that were selecting methods that developed more accurate predictions seemed to do great among all of the agents with poor predictors.

In order to determine whether it is beneficial to have a different multiplier for each of the rounds of bidding, cases were conducted using a single multiplier and the equilibrium curves looked fairly similar to those obtained when using the multiple bid multiplier gene. The fitnesses of the agents using the single multiplier seemed to be a little higher. Under the multiple bid multiplier scenario, when the number of contracts were held constant, the agents were able to adjust their bid multipliers to increase their fitness. It would be natural to assume that the separate multipliers give the agents more flexibility to decide when to take more profit, and when to hold off.

In the plots of this section, it is often difficult to tell exactly what is occurring as the graphs appear somewhat noisy. Much of the fluctuation in the equilibrium price comes from the nature of the bidding scenarios that were been used. Within the generations, there tends to be a 1-exp(-rounds) form to the equilibrium price, while from generation to generation, there is more gradual movement to what one might expect to be the equilibrium point. Each of the bid multipliers adds puts a wrinkle in the equilibrium price and consequently with the single bidding multiplier there are much smoother equilibrium prices curves.

CHAPTER V. CONCLUSIONS AND FUTURE RESEARCH

Conclusions

Based on the results presented in the previous chapter, the reader is able to draw some conclusions. First of all, the algorithms developed for this research are operating as expected and producing the expected outcomes. The supply and demand market forces arrive at prices that one might consider fair. The genco agents are modifying their behavior over time to increase their profit.

When the supply is much greater than the demand, the buyer is able to purchase power less expensively than when the demand is much greater than supply. The agents quickly learn to increase the number of contracts to offer during the bidding process in order to take advantage of the buyer in situations where supply is less than the amount of electricity demanded. However, when the supply is much greater than demand, they don't really modify the number of contracts to offer in any consistent manner.

Agents modify their bid multipliers and their choice of prediction techniques to increase their fitness if possible. Agents use the prediction techniques coupled with a bid multiplier to aid in making profitable bids. Each agent's choice of prediction technique and its own bid multiplier gives it a unique bidding behavior.

In situations when the supply is much greater than demand, it is difficult for an agent to get its bid accepted among all of the other sellers' bids. The learn that in order to get a contract, they must under bid the other gencos. They do this by decreasing their bid multipliers so that they are bidding slightly above their cost of producing power. This gives an obvious advantage to those producers who can generate electricity inexpensively. Another way to improve an agent's chances at getting a contract could be to bid high the first rounds of bidding, saving its inexpensive generation for later when others are offering electricity that has been generated at a higher cost.

The use of separate bid multipliers for each round of bidding allows the agents to change their bids over the course of the rounds of bidding. Using separate bid multipliers each round has the effect of allowing each round of bidding to be fairly different from the previous round (depending on the range of the bid multipliers). However, when the multiple bid multipliers were used at the same time as the agents were given access to prediction techniques, the agents did not seem to evolve both of these parameters well. Perhaps optimizing both the prediction parameters and the bid multipliers is too much for the agents. In fact, alone each of these evolving parameters of the agents maybe doing roughly the same thing, while together they are acting against each other. Therefore, one conclusion of this research is to keep the agent's genetic structure as simple as possible. Do not include the prediction parameter evolution if already including a separate bid multiplier for each round of bidding.

Two different methods of developing bids for the auction were investigated. The first method involved giving each agent access to a public prediction of the equilibrium price. The agent then used an individual bid multiplier or multipliers to make his bid based on that predicted price. This seemed to result in a fairly smooth equilibrium price over the rounds of bidding, which would be expected. It also seemed to react rather quickly to changes in supply and demand. The agents that were successful in getting a contract benefited by receiving more profit. However, from a social welfare standpoint, frequently many of the agents offered power at a price much greater than what the buyer was willing to pay. This resulted in those agents receiving no profit for those rounds, and the buyer ended with a shortage of power.

The second method of developing bids did not use a price forecaster explicitly. Rather it took the previous round's equilibrium price as its expected price. In a market with a large number of players, where each players actions do not greatly affect the equilibrium price, this is a fairly valid assumption. The bid multiplier was then used to determine where to place the bid in the range from the seller's cost to the expected price. The minimum bid that a seller can make is the cost of generation, and the maximum bid that the seller can make is the previous equilibrium price. In developing the actual equilibrium price, the auctioneer sets the price at the midpoint of the buyer's and seller's bids. If the buyer's bid remains constant as was assumed for the results presented in Chapter IV, this limits the movement of the equilibrium price to the midpoint of the cost and the buyer's bid for the lower range to an upper range approaching the buyer's bid. In cases where the demand is greater than supply, and all selling agents learn to use the expected price as their bid, the price approaches the buyer's bid. But as the agents evolve, new agents are created via crossover and mutation that bid lower than this. The lower bids prohibit the equilibrium price from ever actually reaching the buyer's bid.

Allowing the agents to bid above the expected price will not force the equilibrium price to stay at the buyer's bid in this scenario. The bids will actually exceed the buyer's bid and these agents will be denied a contract. They receive a low profit and are replaced. Choosing the expected price as the upper bound is an interesting selection. With the buyer's bid fixed, the resulting behavior is not as interesting as it would be if it was allowed to change. If multiple buyers were included and were allowed to bid in a similar fashion, the equilibrium price could move within the sellers cost and the buyer's upper bound.

This second scenario might lend itself well to bidding rules developed by Ashlock [1] for the Divide the Dollar game. Where instead of bidding between 0 and 100, here the limits are the seller's cost and the buyer's valuation of electricity. Another difference is that here we have multiple players where in the Divide the Dollar game there are only two bidders.

Future Research

The research and results that have been described in chapter IV and the conclusions that have been drawn based on those results, are based on some assumptions. One assumption is that the generation companies all desire to sell power from each generating unit at each hour. In order to more closely model the conditions observed in the real utility industry, the startup costs, ramp constraints, transition costs, and shutdown costs should be included in the determination of the genco's profit. Another assumption that needs adjusting is that one disco can accurately model the actions of many discos. With a little additional code, the buyers could be represented by multiple discos, allowing more intricate interaction between buyers and sellers.

The author intends to continue research in this area of investigation. In the future research there are several areas that would be natural extensions this thesis including the following areas:

- The bidding strategies of both buyers and sellers will be investigated.
- Consideration of a generic agent not labeled a genco or disco that would be allowed to buy and sell electricity in much the same way as a power marketer.
- In addition to the GA used here, the use of genetic programming could prove to be a useful tool in developing strategies for bidding.
- The neural network could be used to predict the equilibrium price.

A more realistic model

To accurately reflect the choices gencos have in developing their bids, more accurate models must be used. Including ramp rates, minimum up-times, minimum down-times, environmental constraints, and startup and shutdown constraints will influence the way bids are made and the profitability of the resulting contracts. For example, suppose there is a generating unit that needs to be taken off-line for routine maintenance. It may be more profitable to bring the unit down during a traditionally higher cost weekday than to bring it down on the weekend when costs are low, provided the owner is able to sell the power during the weekend for a great price. Minimum up and down times add a whole new dimension to the problem of developing profitable bids.

Along with the above constraints, the dependence of time and risk on a bid should be considered. An accurate model of a futures market should be included. Utilities should not have to risk their production based on fuel prices three months into the future. Gencos could purchase a futures contact guaranteeing a maximum cost of electricity for future months. This would affect the prices that they can ask, and the gencos profitability. This should not only be done by gencos, but discos should investigate this as well.

The research described in this thesis has assumed that transportation costs are negligible. Incorporating realistic transportation costs, or prices depending on which model is used, can drastically change the profitability of a contract. If transmission owners were allowed to participate in the auction, and make bids on the price of transmission services, the auction would become more complicated. This should be investigated in future research.

Consideration of discos

By including multiple agents to represent discos, the simulation might more closely represent what is observed in the real world. In areas where the volume of trade is low, and the number of gencos and discos participating in the auction or commodity exchange is small, there may exist opportunities to take advantage of mistakes made by rival bidders. Although it may be possible to prove that in a theoretical market, with ideal bidders, the result is an efficient market and consumers benefit, utilities should be interested in taking advantage of every profit opportunity. Developing successful bidding strategies requires using all information available and that includes the knowledge of what one's competition is doing.

Power marketers

A natural extension of this work is to observe what occurs when agents are allowed to be both buyer and seller. This is what power marketers are doing in the electric energy industry now, and what many of the vertically integrated utilities should be doing to increase their profitability. This has the effect of leveling the price of power across regions. If a genco can purchase power from a neighboring genco or power distributor less expensively than it can produce power with its own generator, then it should do so. In a totally competitive market, power producers may no longer cooperate as they do today, which means that pool-wide economic dispatch of generating units could be eliminated. If it is eliminated, power marketers help keep the incremental cost of generation fairly level across the power system. In the competitive market scenario, utilities should not aim to keep the system incremental cost equal across the country. They should be interested in profit for the company, and dividends for the shareholders.

Genetic programming

Developed after the emergence of genetic algorithms, genetic programming (GP) uses the same principles of natural selection that genetic algorithms use. Rather than evolving the contents of a fixed data structure like a genetic algorithm, GP can actually evolve small computer programs. The programs can be represented by parse trees. See Figure 5.1. Parse trees are generated randomly during the initialization part of the program. Only valid combinations are allowed. Only operator nodes may be connected to terminal nodes, and no terminal nodes shall have subtrees. Fitness is calculated based on some performance index. If we are trying to maximize profit, then it should be proportional to the resulting profit. Parents are selected in the same manner as used in genetic algorithms.

Sample Parse Tree

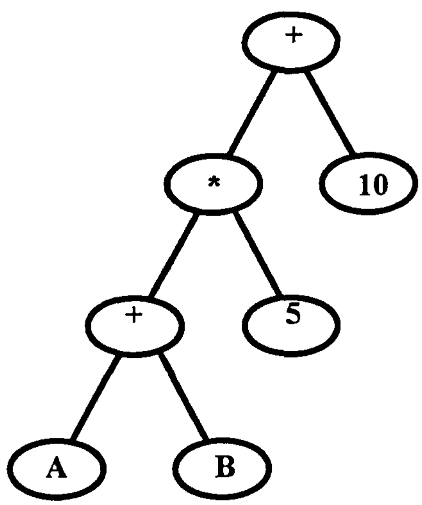

Nodes are either terminals or operators. Operators operate on the nodes linked to them from below. Operator nodes in this ex. are (+) addition and (*) multiplication. The remaining nodes are terminal nodes.

The parse tree can be decoded into the following equation:

$$(((A + B) * 5) + 10)$$

The tree is read starting at the left most operator node. At the lowest (+), we can read the terminals (A) and (B), and add them since the operator is addition.

Figure 5.1. Reading a GP parse tree

Crossover involves choosing a random node within each of the two parents. The subtrees defined by all nodes below this randomly chosen node are swapped. Mutation involves randomly choosing nodes and with a probability replacing that node with another valid operator or terminal. Another form of mutation eliminates the subtree below the chosen node, and randomly generates a new subtree.

One of the big differences between GPs and genetic algorithms is that GPs require the developer to develop a language of operators and terminals. In Figure 5.1, operators were chosen to be addition and multiplication, and these have the usual meaning. The example also allowed variables and integers. This gives way to a enormous combination of possible outcomes. Because of this, GPs can search the solution space, not just a predetermined portion of that solution space as is the case with GAs. Searching such a large solution space might never yield results. Therefore, careful selection of a valid and useful parse tree language is required for good results. Developing operators that work well together and that can actually get results can be rather difficult.

Recently researchers at Iowa State University have experimented with modifications to the normal GP. Ashlock [7] obtained interesting results using a hybrid GP-Finite State Automata for a fairly simple problem. Much has been published recently regarding GP and the author would like to use this new technique to help develop successful bidding strategies in competitive markets.

BIBLIOGRAPHY

1. Ashlock, Dan. (1995). "GP-automata for Dividing the Dollar." Mathematics Department, Iowa State University, Ames, IA
2. Barkovich, Barbara, and Hawk, Dianne. (1996). "Charting a new course in California" , IEEE Spectrum, July 1996, Vol. 33, p26. New York: IEEE
3. Fahd, George, Richards, Daniel, and Sheblé, Gerald. (1991). "The implementation of an energy brokerage system using linear programming." Presented at the 1991 IEEE/PES Winter Meeting, 91 WM 177-6 PWRS. New York: IEEE
4. Finlay, David J. (1995). Optimal Bidding Strategies in Competitive Electric Power Pools. Masters thesis, University of Illinois, Urbana-Champaign, IL.
5. Gilchrist, Warren. (1976). Statistical Forecasting. New York, NY: John Wiley & Sons.
6. Goldberg, David E. (1989). Genetic Algorithms in Search, Optimization & Machine Learning. Reading, MA: Addison-Wesley Publishing Company, Inc.
7. Hakin, Simon. (1994). Neural Networks: A Comprehensive Foundation. Englewood Cliffs, NJ: Macmillan Publishing Company.
8. Hunt, Sally and Shuttleworth, Graham. (1996) "Unlocking the Grid" , IEEE Spectrum, July 1996, Vol. 33, p20. New York: IEEE
9. Ioannides, Y.M. (1995) "Evolution of Trading Structures." Department of Economics, Tufts University, Medford, MA 02155.
10. Kazarlis, S.A., Bakirtzis, A.G., and Petridis, V. (1995) "A Genetic Algorithm Solution to the Unit Commitment Problem." 1995 IEEE/PES Winter Meeting 152-9 PWRS. New York: IEEE
11. Kumar, Jayant and Sheblé, Gerald. (1996). "Auction Game in Electric Power Market Place." Proceedings of the 58th American Power Conference, 1996. pp. 356-364.
12. Kumar, Jayant and Sheblé, Gerald. (1996). "Framework for Energy Brokerage System with Reserve Margin and Transmission Losses." 1996 IEEE/PES Winter Meeting, 96 WM 190-9 PWRS. New York: IEEE

13. LeBaron, Blake. (1995). "Experiments in Evolutionary Finance." University of Wisconsin - Madison, Madison, Wisconsin.

14. Makridakis, S., and Wheelwright, S.C. (1978) Forecasting Methods and Applications. New York, NY: John Wiley & Sons

15. Milgrom, Paul. (1989). Auctions and Bidding: A Primer. Journal of Economic Perspectives, Vol. 3, No. 3. Summer 1989, pp 3-22.

16. Mohammed, O., Park, D., Merchant, R., Dinh, T., Tong, C., and Azeem, A. (1994) "Practical Experiences with An Adaptive Neural Network Short-Term Load Forecasting System." Presented at the 1994 IEEE/PES Winter Meeting, 94 WM 210-5 PWRS. New York: IEEE

17. Morrison, Norman. (1969). Introduction to Sequential Smoothing and Prediction. New York, NY: McGraw-Hill Book Company.

18. Papalexopoulos, Alex D., Hao, S. and Peng, Tie-Mao (1994) "An implementation of a neural network based load forecasting model for the ems", Presented at the 1994 IEEE/PES Winter Meeting, 94 WM 209-7 PWRS. New York: IEEE

19. Post, Douglas L. (1994). Electric power interchange transaction analysis and selection. Master's thesis, Iowa State University, Ames, IA.

20. Post, D. L., Coppinger, S. S., and Sheblé, G. B. (1995). "Application of auctions as a pricing mechanism for the interchange of electric power." Presented at the 1995 IEEE/PES Winter Meeting, 95 WM 175-0 PWRS. New York: IEEE

21. Sheblé, Gerald B. (1996). "Priced Based Operation in an Auction Market Structure". Paper presented at the 1996 Power Engineering Conference. Baltimore, MD.

22. Sheblé, Gerald B. (1994). "Electric energy in a fully evolved marketplace." Presented at the 1994 North American Power Conference.

23. Sheblé G. B., and McCalley, J. D. (1994). "Discrete auction systems for power system management." Presented at the 1994 National Science Foundation Workshop, Pullman, WA.

24. Thompson, Gerald L., and Thore, Sten. (1992). Computational Economics: Economic Modeling with Optimization. San Francisco, CA: Scientific Press.

25. Tesfatsion, L. (1995) "A Trade Network Game with Endogenous Partner Selection." Economic Report Series, Department of Economics, Iowa State University, Ames, IA.

26. Vojdani, A.F., Imparto, C.F., Saini, N.K., Wollenberg, B.F., and Happ,H.H. (1994). "Transmission Access Issues." Presented at the 1995 IEEE/PES Winter Meeting, 95 WM 121-4 PWRS. New York: IEEE

27. Walters, D.C. and Sheblé, G.B. (1992). "Genetic Algorithm Solution of Economic Dispatch with Valve Point Loading." 1992 IEEE/PES Summer Meeting 414-3. New York: IEEE

28. Wood, Allen J. and Wollenberg, Bruce F. (1984). Power Generation, Operation, and Control. New York, NY: John Wiley & Sons.