1-1-2002

Exploring computational power markets with evolutionary algorithms

Valentin Tzankov Petrov
Iowa State University

Exploring computational power markets
with evolutionary algorithms

by

Valentin Tzankov Petrov

A thesis submitted to the graduate faculty
in partial fulfillment of the requirement for the degree of
MASTER OF SCIENCE

Major: Computer Engineering

Program of Study Committee:
Gerald Sheblé (Major Professor)
David Hennessy
Daniel Berleant

Iowa State University
Ames, Iowa
2002

Graduate College
Iowa State University

This is to certify that the Master's thesis of

Valentin Tzankov Petrov

has met the thesis requirements of Iowa State University

Signatures have been redacted for privacy

This work is dedicated to my mother.

TABLE OF CONTENTS

CHAPTER 1. INTRODUCTION

1.1 Overview

Throughout the world, the electric industry is in the midst of major changes designed to promote competition. No longer vertically integrated with guaranteed customers and suppliers, electric generators and distributors will have to compete to sell and buy electricity. The stable electric utilities of the past will find themselves in a highly competitive environment. Some countries and regions of the US (e.g., California, PJM, ERCOT) are already operating in a restructured environment. There does not yet appear to be a standardized final market structure that works for all areas, but each market that springs up adds to our experience and helps us make the next market implementation work a little better and more competitively. The author believes that, to some degree depending on the market implementation, regional commodity exchanges will play a key role in buying and selling electricity.

Figure 1.1 represents the new organizational market structure. The open market system consists of generation companies (GENCOs), distribution companies (DISTCOs), transmission companies (TRANSCOs) and brokers (BROCOs). The Independent System Operator (ISO) is in general ensuring the grid reliability, clearing the market and taking care of the market settlement.

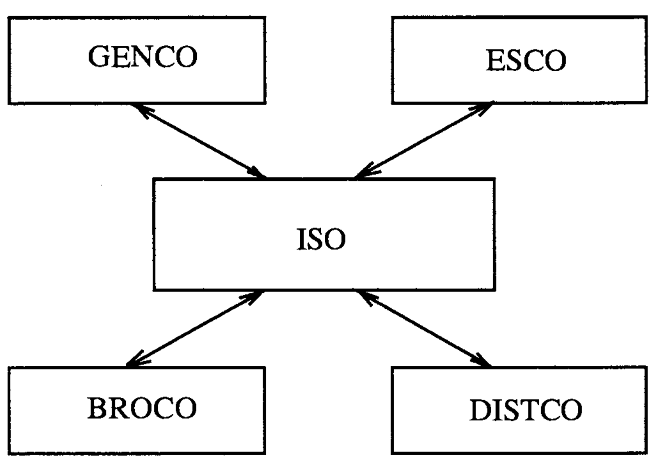

Fig. 1.1. Market structure

1.2 Main Market Role Players

In this work, a market with only GENCOs, ESCOs and an ISO is presented. A brief description of each of them is as follows [5].

Generation Company (GENCO)

The GENCO's primary role is to package standardized products at competitive prices. The GENCOs main concern is to dispatch its products (energy and ancillary services) against the market and maximize its gross margin. Trading mostly occurs at a wholesale level.

Energy Services Company (ESCO)

The primary goal of the ESCO is to purchase power at least cost when needed by consumers. In this respect, the ESCOs could be regarded as retail companies, purchasing at wholesale prices, prepackaging if necessary and reselling at retail prices to end consumers.

Independent System Operator (ISO)

ISOs primarily concern is ensuring the grid reliability. The means for doing that are mostly economical: clearing the market, and subsequently settling the market. The ISO also has the ultimate power to exercise full market control (forcibly dispatch units, shed load, etc.) as a final resort in order to keep the system stable. The ISOs exact role is generally described by the market protocols for each market.

1.3 Problem Motivation and Presentation

With an electricity market that promises to do nearly $100 billion of business each year, competition is expected to be stiff. The number of market players will be larger than ever before, and it would be impossible to guarantee that all participants will be trustworthy. It is likely that participants in this market will attempt to gain a competitive advantage if a means to do so arises. Corporate espionage that reveals the strategies of competitors could provide valuable information in developing one's own strategies. Disruptive practices need not be illegal. Loopholes in the regulation may allow a market participant to engage in activities of a questionable nature.

It is inevitable that some of the participants in the newly deregulated electric market will try to benefit by employing techniques dealing with the financial aspect of markets like predatory pricing to eliminate competition or purposely inflating the prices. Although the rules may be designed to encourage what is perceived to be "fair-play", the opportunity for profits will lead some to test the meaning of fairness. Other malicious actions may deal directly with the physical systems including causing congestion and blackouts/brownouts.

In this work attention is paid to the study of the effects of malicious gaming behavior during bidding in agent based power markets. The agents are modeled as autonomous learning entities utilizing different evolutionary programming techniques such as genetic algorithms, reinforced learning and genetic programming.

1.4 Main Contributions

The results of this work contributed in three ways to advancing our understanding of computational power markets. First, different evolutionary techniques were successfully implemented in the framework of a deregulated power market with gradually increasing degrees of complexity. Valuable insights were gained as to how the separate components would architecturally fit the computational market model. Second, important results about the market behavior were obtained by experimenting with various possible scenarios of market participants playing on the market. Third, two of the evolutionary techniques, namely the genetic algorithm (GA) and the Roth-Erev reinforced learning were improved. A better implementation of the GA was devised which allowed for more complex participant behavior. Also, the Roth-Erev reinforced learning algorithm was altered and considerably improved to allow for a very robust computational implementation.

CHAPTER 2. ECONOMICALLY DESTABILIZING ELECTRIC POWER MARKETS FOR PROFIT

2.1 Introduction

In this chapter a double auction market of electricity is modeled, and different scenarios for malicious activities from the sellers' side are considered and simulated. It is shown that under proper circumstances some of the players can significantly influence the market thus benefiting themselves.

Attention is paid to the study of the effects of malicious gaming behavior during bidding in agent based power markets. The chapter is organized as follows. Section 2.2 discusses the methods and techniques employed modeling the agents and markets for the simulations described here. Section 2.3 describes the design of the small experiment included in this work. Section 2.4 provides the results and analysis of the simulation. Section 2.5 talks about market power. Finally, Section 2.6 presents some ideas for future research.

2.2 Methods and Techniques

Different models have been built to simulate the electric marketplace. In this study the scheme outlined in Ref. [1] is used. A bilateral auction is constructed, with a fixed number of buyers and sellers of electricity. The bilateral contracts for fixed amount of electric energy are performed in fixed time intervals. The transmission system is lossless and has unlimited capacity. The trading agents use a fixed set of rules to change the bidding strategy [2,6]. The auction bid matching is performed by an independent third side. One round of bidding proceeds as follows. All buying/selling agents submit their bids/offers to a third party, an Independent Contract Administrator (ICA). ICA matches the bids using an approach similar to Wood and Wollenberg [4,5]. All bids and offers are sorted in descending and ascending order and juxtaposed. If a buy bid is higher than the corresponding sell offer, the two players are matched, a contract is approved and the players are notified. In the case where the buy bid is lower than

Table 2.1. An example of an auction bid matching of 5 sellers/buyers, and the resulting outcomes

Bid	Offer	Quantity	Match	Midpoint	Eq. Price
10.91	10.01	1 MW/h	Yes	10.46	10.48
10.76	10.07	1 MW/h	Yes	10.41	10.48
10.74	10.38	1 MW/h	Yes	10.56	10.48
10.26	10.53	1 MW/h	No	N/A	N/A
10.05	10.67	1 MW/h	No	N/A	N/A

the sell offer, the players are rejected and notified. There is no second call for bidding regardless of the number of contracts approved. A bid and offer midpoint, as well as an equilibrium price is calculated and the next round is called. Table 2.1 shows an example of 5 buyers and sellers after the bids and offers were matched by ICA.

Different pricing techniques can be studied using the model described here. Midpoint prices are calculated, which is used if discriminatory pricing is desired, and if non-discriminatory pricing is desired, the valid matches weighted by quantity are used to determine equilibrium price. Meanwhile the players update their strategy and the ICA calls for the next round. Each player updates its strategy according to the outcome of the previous round. The player knows only the outcome of its own last action (e.g. contract accepted/rejected) and remembers the bid/offer value. If the offer was rejected, next time the seller lowers the price, choosing the value randomly within an interval it calculates based on a rule including its own last offer. Accordingly, the buyer bids are higher than the preceding round. If the offer was accepted, the seller offers a higher price based on a rule including the last accepted offer, and the buyer bids are reduced. All players start trading around a fixed equilibrium price submitted by ICA. This is done for convenience. In a real world situation the players can start bidding around the closing price for the previous day, for example. After the first round the ICA takes care only of matching players.

Pseudocode is given on the next page, (Fig. 2.1) describing the actions of a seller and buyer depending on the outcome of the last round. In any event, we must set upBid < downBid, and upOffer < downOffer. If it were otherwise the price would not converge.

```
if (seller)
  if (oldOffer accepted)
    newOffer = oldOffer + upOffer;
  else
    newOffer = oldOffer - downOffer;
  end-if
end-if
if (buyer)
  if (oldBid accepted)
    newBid = oldBid - downBid;
  else
    newBid = oldBid + upBid;
  end-if
end-if
```

Fig. 2.1. Pseudocode describing the actions of a seller and buyer depending on the outcome of the last round.

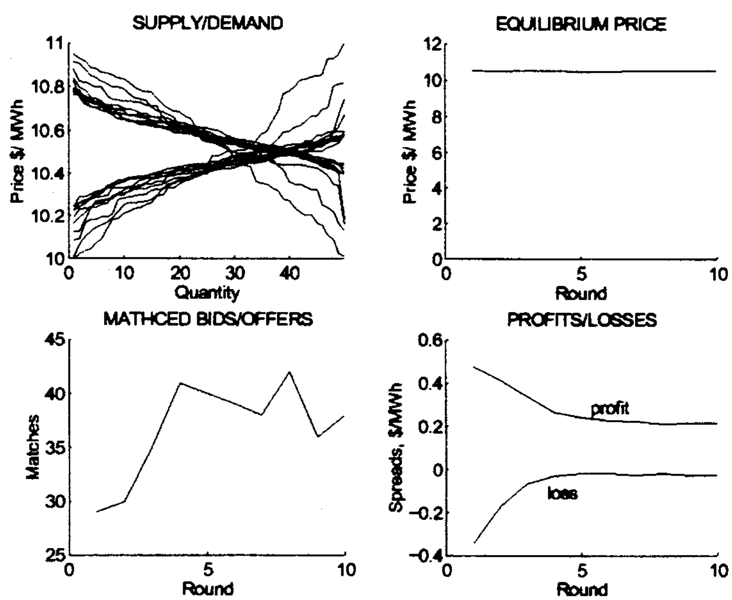

Fig. 2.2. Supply/Demand, equilibrium price, offers accepted and profit/loss spread for 50 buyers/sellers of electric power after 10 bidding rounds.

2.3 Experimental Design

The described system fairly quickly reaches equilibrium, and the buy/sell spread, the midpoint price and the equilibrium price converge to a limiting value. Fig. 2.2 on the previous page shows the pictorial representation of an auction consisting of 100 players, evenly divided into buyers and sellers of electricity. The auction went through 10 rounds. In this case downBid and upOffer were set at 5% of the last unsuccessful bid/offer and downOffer and upBid were set to 3% of the last successful bid/offer. All four price updates were tied up to a random generator, i.e. the players may choose price updates with equal probability within the prescribed price update interval.

Such seemingly simple market behavior is very similar to what was observed in experimental double auction markets played by people [4]. It can be seen that the initial price submitted by ICA remained stable, and the number of accepted offers quickly reached about 70% of the total number of offers submitted, up from about 50%. After only a few rounds the players minimized the losses to almost zero. However, the profits also dropped. The system did not diverge even after simulating a few years of bidding rounds, assuming that each round lasted an hour and players traded fixed amounts of electric power.

2.4 Results and Discussion

After building a stable model of the auction, simulations were run where some of the players misbehaved and attempted different strategies. A few cases were considered. Also, ICA was not allowed to intervene, except to match the submitted bids/offers.

First, one of the sellers attempted consistently to offer lower prices than the others (predatory pricing). The predatory pricing would drive the competition out, justified later by increasing the price by the remaining player, thus creating certain state of monopoly. After calculating the new offer, the resulting price was additionally decreased by 5%. For a 10 buyer/seller auction the effect of one single player occasionally offering lower prices may seem negligible. However, it turned out that the seller was able to drive the price down by

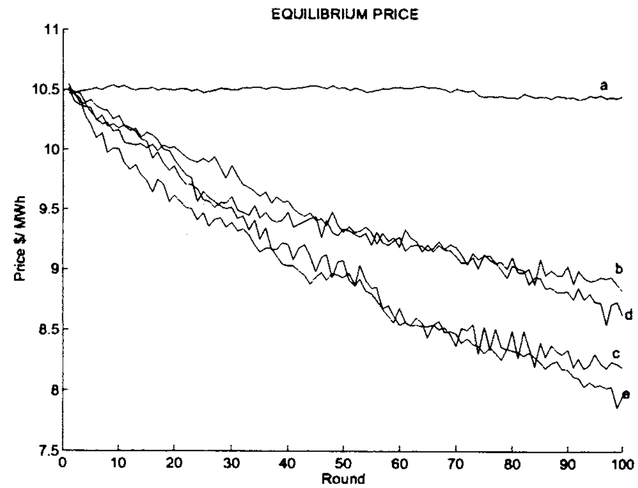

Fig. 2.3. The effect of predatory pricing on a 10 buyer/seller auction; a) no price decreasing, b), c), d and e) predatory pricing by one, two, three and four sellers.

more than 10% only after 100 rounds. The equilibrium price fell off exponentially, approaching zero. Fig. 2.3 shows the equilibrium price when the seller's downOffer is on average twice higher than the regular downOffer decrement no matter the out-come of the previous sell offer. The results were averaged over 10 auction runs.

The tactic was consequently employed by one, two, three and four sellers. It is interesting to see that after the initial sharp reduction of the equilibrium price by more that 10%, there was no significant difference when the second, third and fourth player engaged in predatory pricing. This suggests that if a company has a few generators, each of them participating in the market, the company may randomly designate each generator as a predator during every round, making it difficult to spot such malicious behavior, yet achieving the goal of lowering the price. This would allow the seller to efficiently eliminate other sellers within some period of time. The effect on the number of successful contracts is as follows: equilibrium market - 62%, one player - 64% and four players - 69% of the contracts were approved. The reason for the increase is that as the sellers decrease the price, the supply shifts down, allowing more buyers to make contracts.

Next, a scenario where five and nine of the sellers attempted predatory pricing, thus trying to eliminate the rest of the sellers, was tested. Each time an unsuccessful offer was made, the seller's price was additionally dropped by 5%. Fig. 2.4 shows the averaged

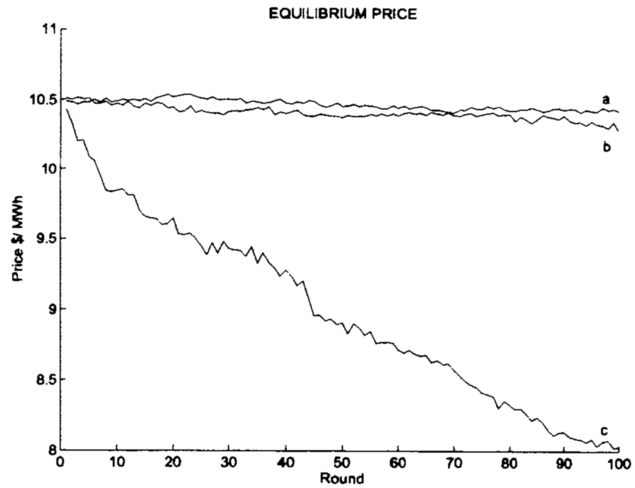

Fig. 2.4. The effect of predatory pricing: a) equilibrium market, b: five players against five, c) nine players against one.

equilibrium price during 100 bidding rounds. The strategy was even more elaborated by allowing each of the predators to offer a consistently lower price only when its last legal bid was rejected. The price is plotted against an equilibrium market price. It is seen that when 5 players attempt this, there is virtually no effect.

It was established that the strategy works only when more than 60% of the players were engaged. It should be noted that 60% is the lowest percentage of players successfully making deals on an equilibrium market without predatory pricing. The effect on the number of successful contracts was somewhat similar to that of the previous experiment: equilibrium market - 62%, five players - 62% and nine players - 68% of the contracts were approved.

Another case was considered, in which one or more of the sellers are trying to increase the price by consistently offering a 5% higher price than the rest. Fig. 2.5. shows the price plotted versus an equilibrium market and versus the case when the same player is trying to lower the price using the same value which was used for overbidding.

It is interesting to see that price increasing during the auction is much easier that price dumping. At the beginning of the auction the price is easy to lower, but as time advances, it is easier to increase and more difficult to decrease. The reason is that the auction becomes exponentially stable during price dumping (the price tends to 0) and exponentially unstable during price increasing (the price tends to infinity). This is so because the inherent structure

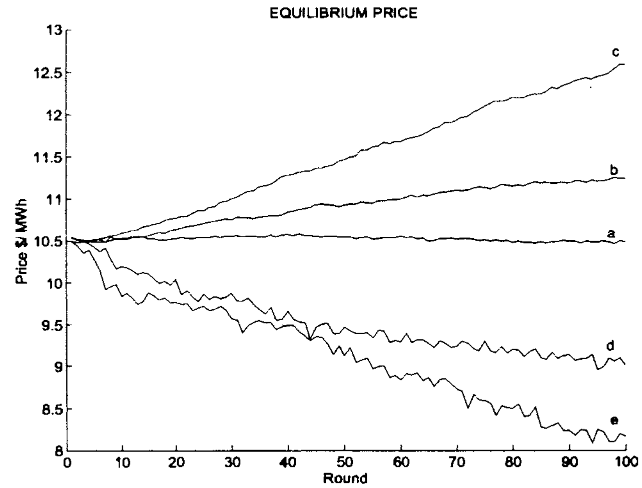

Fig. 2.5. The effect of price increase: a) equilibrium market; b) one seller increase; c) two sellers increase; d) one seller decrease; e) two sellers decrease

of the model does not allow for negative bidding. In that case the number of successful contracts is as follows: equilibrium market - 62%, one player - 60% and four players - 57% of the contracts were approved. The reason for the decrease is that as the sellers increase the price, the supply shifts up, allowing less buyers to make contracts.

2.5 Market Power

Consideration should be given to assessing the market power of each seller. A seller is said to have unilateral market power if a unilateral deviation from a competitive equilibrium is profitable for that seller, given that all other traders continue to use the strategies that generated the competitive equilibrium [7]. The market power in this case is directly proportional to the market share. In a 10 seller market each seller has 10% of the market share, hence 10% market power, given that all sellers trade fixed amounts of electricity and all buyers are willing to buy regardless of the price. The situation would get complicated if unequal quantities of electricity are traded, if the buyers also have market power, if the number of buyers does not equal the number of sellers and each traded unit of electricity has different characteristics (reliability for example). Market power study is treated in more details later.

2.6 Future Work

Further work has to be done to asses the effect of buyers on such an auction market. Another direction into which the research might go is estimating the market power of agents under different market conditions. Also, probabilistic analysis of the auction model similar to Ref. [4] should be carried out, which will allow for a priori estimation of the effect of different actions of buyers and sellers on the market. It is important to note that this is a work in progress and future improvements of the model, like implementing a power grid, modeling congestion and including smart agents, is considered.

CHAPTER 3. GAMING STRATEGIES ON ELECTRIC POWER MARKETS DRIVEN BY AGENTS

3.1 Introduction

This chapter reports on efforts that show how agents can use gaming strategies and predatory pricing to increase their profits at the expense of others. The work described here assumes a framework in which buyers and sellers submit bids to a central coordinating body. The buyers are termed energy service companies (ESCOs) and the sellers are termed generation companies (GENCOs). Prices and quantities are determined in an auction process after buyers and sellers have submitted bids to a power exchange. The power exchange works closely with the Independent System Operator (ISO) to ensure that matched contracts are feasible.

One of the major ISO functions is congestion management. The congestion management ensures that the power grid does not violate operating constraints like transmission capacity, generation limits on active and reactive power and such. Ensuring the integrity of the power network is vital for the operation of the electric power market.

3.2 Methods and Techniques

To keep the effect of different strategies obvious the modeled system is fairly small and easy to understand. The model consists of GENCOs and two ESCOs. Each GENCO owns three generators attached to a simple 3-bus network. Fig. 3.1 on the next page depicts the network. The power grid has two transmission lines with capacities T1 and T2, and two loads L1 and L2. The two GENCOs compete for the two customers by setting their own prices, and an independent entity, a Power Exchange, matches the sellers and buyers.

Different models have been built to simulate the electric marketplace. In this study the scheme outlined in Ref. [14] is used. A double-sided auction is constructed, with two buyers and two sellers of electricity. Computerized agents model the GENCOs and ESCOs as they participate in the simulated auction. The contracts for fixed amounts of electric energy are

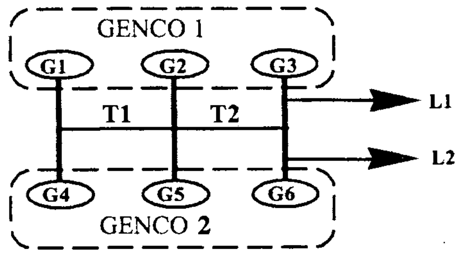

Fig. 3.1. A bilateral model with two generating companies and two customers.

created for fixed time intervals. One of the trading agents uses a fixed set of rules to change the bidding strategy [15, 11]. The other is equipped with an intelligent searching mechanism. The auction bid matching is performed by an independent entity called an Independent Contract Administrator (ICA) which handles the functions of the ISO.

Fig. 3.2 gives a flowchart describing the market simulation. One round of bidding proceeds as follows. The selling agents determine a bidding scheme based on certain criteria, such as

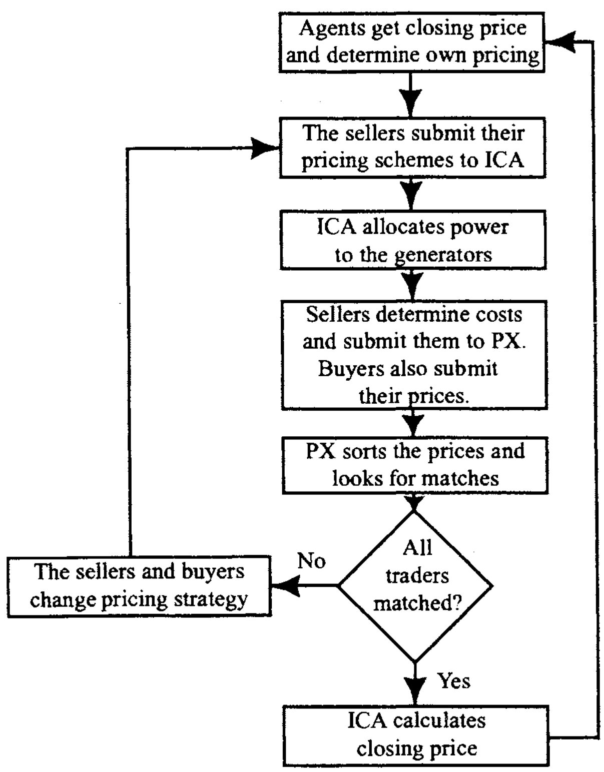

Fig. 3.2. Market flowchart

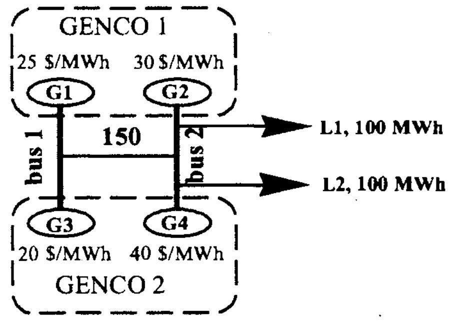

Fig. 3.3. A simple 3-bus example.

load forecasting, estimation of the available transmission, etc. Then they submit their bids. The ICA calculates the power each generator can deliver to the customer while attempting to minimize the total cost and enforce all network constraints. Thus the ICA allocates transmission rights to the cost-effective users. The ICA also determines the cost of transmission used to price the transmission network. Each user pays the cost for using the network.

A very simple example can illustrate the role ICA plays on the market. Let us consider the 2-bus network on Fig. 3.3 with 4 generators serving 2 loads [8]. Two GENCOs own two generators each with 10.0 MWh minimum and 100 MWh maximum load. The transmission limit of the line from bus 1 to bus 2 is set at 150 MWh. Each GENCO is serving a separate 100 MWh load. Both GENCOs produce bids for its generators and submit the bidding scheme to the ICA. If there were no limits imposed on the transmission line from bus 1 to bus 2, the ICA would allocate 100 MWh transmission to G1 and G3 each, because they are priced lower that G2 and G4 and this would minimize the total cost. Thus, the cost for L1 would be $2500/h and the cost for L2 would be $2000/h.

However, if there is 150 MWh transmission limit on the transmission line, ICA would allocate 50 MWh to G1, 50 MWh to G2, and 100 MWh to G3. Thus, the cost for L1 becomes $2750/h, and for L2 is $2000/h plus they have to pay for congestion transmission capacity. The congestion transmission costs are transferred to the customers and included in the final price. They are calculated as follows. If ICA were to substitute 50 MWh from G3 with the same amount from G4, this would cost $20/MWh to GENCO 2. Substituting 50 MWh from G1 with generation from G2 costs $5/MWh to GENCO 1, therefore the transmission from bus 1 to bus 2 is much more valuable to GENCO 2, and ICA minimized the cost by rescheduling the generators of GENCO 1 instead of GENCO 2. Thus, GENCO 1 appears as the marginal user of the congested transmission line and sets the marginal cost of transmission at $5/MWh. In that case GENCO 1 has to pay $5/MWh × 50MWh = $250.0 congested transmission costs for carrying 50 MWh over the congested line, and in the same fashion GENCO 2 has to pay $500.0 for carrying 100 MWh over the congested line. These amounts are calculated in the total price the customer pays, and transferred to the transmission owner by the ICA. This pricing scheme is very important because it unbundles energy and transmission costs, thus creating two separate markets. This allows the generation companies to properly arrange their pricing schedule in a way that minimizes costs and assures available transmission. The process of forecasting and managing the transmission by the ICA is called congestion management [8].

After the GENCOs submit their bids and are allocated transmission rights, they calculate the price and submit their offers. Respectively, the buyers submit their bids. An independent entity, a power exchange, matches the bids and offers using an approach which maximizes the profit, similar to Wood and Wollenberg [10, 15].

All bids and offers are sorted in descending and ascending order and juxtaposed. If a buy bid is higher than the corresponding sell offer, the two players are matched, a contract is approved and the players are notified. In the case where the buy bid is lower than the sell offer, the players are rejected and notified. The rest of the players are also notified that the market was not cleared and prompted to bid again. This is repeated until all players are matched and the market cleared, or it is deemed that enough bidding cycles have occurred and the matched players granted contracts. The rest of the players are rejected and notified.

Each cycle the sellers resubmit offers by providing the contract administrator with the new bidding scheme and it recalculates the congestion prices, reschedules generators if necessary and the players submit bids based on the transmission rights. A matched bid and offer average for each contract, as well as an equilibrium price is calculated and the next round is called.

Table 3.1. An example of an auction matching scheme of 2 sellers and buyers, and the resulting outcomes.

Bid	Offer	Quantity	Match	Midpoint	Eq. Price
10.91	10.01	1 MW/h	Yes	10.46	10.48
10.76	10.07	1 MW/h	Yes	10.41	10.48

Table 3.1 shows an example of 2 buyers and sellers after the bids and offers were matched by the power exchange. Different pricing techniques can be studied using the model described here. Midpoint prices are calculated, which is used if discriminatory pricing is desired, and if nondiscriminatory pricing is desired, the valid matches weighted by quantity are used to determine equilibrium price. Meanwhile the players update their strategy and the power exchange calls for the next round.

Each player updates its strategy according to its own forecasting techniques like load prediction, congested transmission forecast, etc. The players knows the outcome of its the past auctions (e.g. contract accepted/rejected), the bid/offer value, the congestion price and the closing price. If the offer was rejected, next time the seller lowers the price, choosing a pricing scheme it believes will bring maximum profit. Accordingly, the buyer bids are higher than the preceding round giving the agent a better chance to get a contract. If the offer was accepted, next time the seller offers a higher price, and the buyer bids are reduced. The players start trading around the previous round's closing price.

3.3 Experimental Design

The model was simulated using C++. Separate objects were created for the ICA, traders, and the power exchange. This allowed for quickly and conveniently changing the rules which the separate entities on the market have to follow.

At first both agents were equipped with similar bidding strategies. At the first step the agents get the closing price from the previous bidding round. Then the sellers set a bidding scheme that would result in a bid slightly below the closing price if all generators were allocated

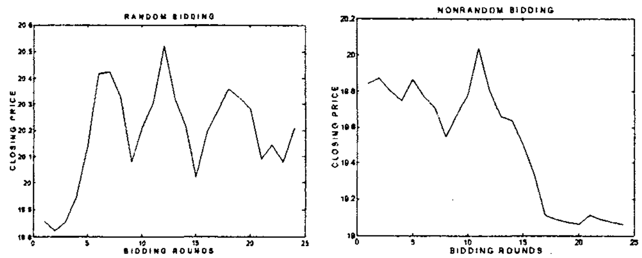

(a) (b)

Fig. 3.4. Closing price as a result of random (a) and nonrandom (b) bidding after 24 rounds

equal amount of power. In a real situation the bids for the separate generators of a particular GENCO would differ from each other due to numerous reasons like amount of rainfall, current weather conditions, etc. This was modeled by perturbing the bids around an average bid that would be slightly lower (3% in this case) than the closing price from the previous round. In the same fashion the buyers set their bids slightly higher (3% again) than the closing price. After the power was allocated by the ICA and bids/offers matched, if there were no price discovery, the sellers would lower the offers with 1% proportionally. The buyers in the same vein would raise the bids proportionally with 1%. This strategy usually led to price discovery on the 2-nd call at most.

After the round is completed, the ICA would announce the closing price calculated as the average of the two median prices of the contracts and call for the next round. Meanwhile the sellers would try to "forecast" the demand by repeating step one. The same was valid for the buyers. At first glance two randomly bidding GENCOs against each other do not provide much insight in the market behavior. Eventually, we could include results from previous runs, and employ elaborate forecast techniques, but for now we keep the model simple. Later it becomes clear that the random "forecast" and also randomizing the bidding scheme is very beneficial in sense that it allows for unbiased checking of different strategies when one of the GENCOs is trying to outsmart the other. Also the routines for determining the bidding strategy could very easily be fed with actual data.

The described system fairly quickly reached equilibrium (Fig. 3.4, a), and the equilibrium price converged to a limiting value. Such seemingly simple market behavior is very similar to what was observed in experimental double auction markets played by people [10, 11]. It can be seen that the initial price remained stable, and each round the market got cleared after the first or second call. The system did not diverge even after simulating a few years of bidding rounds, assuming that each round lasted an hour and players traded fixed amounts of electric power.

3.4 Results and Discussion

After building a stable model of the market, simulations were run where one of the players misbehaved and attempted different bidding strategies. One of the goals of that player was to set such a bidding scheme that ICA would always allocate any power or as much as possible to a particular generator of that seller. For that purpose the player was granted knowledge to the bidding scheme of the other GENCO. It was also equipped with full knowledge of the power grid and the same algorithm ICA uses to allocate available transmission resources. The problem was looked upon as an optimization process. The player would optimize its bidding scheme subject to maximizing the power allocated to that particular generator. Because of the discrete nature of the problem, the player was equipped with an algorithm which performed a random search of the solution space. It would randomly try different bidding schemes until one of them fulfills a condition. In most of the cases the algorithm found the best solution only in a few steps.

Table 3.2 represents a case when GENCO 1 obtained a solution where G1 was allocated power above the minimum limit of 10.0 MWh. However, some of the pricing schemes of the

Table 3.2. A sample solution, GENCO 1 / G1 gets power.

GENCO 1	G1	G2	G3
Bid	19.85	19.78	20.36
Power	79.88	100.0	20.11
GENCO 2	G1	G2	G3
Bid	19.29	19.69	19.75
Power	69.09	39.84	91.06

Table 3.3. A sample solution, GENCO 2 / G1 does not get power.

GENCO 1	G1	G2	G3
Bid	19.99	19.60	19.83
Power	10.00	100.00	90.00
GENCO 2	G1	G2	G3
Bid	19.85	19.61	19.48
Power	10.00	10.00	180.00

other company were intractable in sense that it would take unreasonable amount of time to reach a solution.

Other simulations were run where GENCO 1 tried to find such pricing schemes so that one of the generators of the other company would not get any power allocated. Table 3.3 represents such a solution. Another interesting result, which can be seen in that table is that it was possible for GENCO 1 to obtain a solution where although the transmission limit of T1 was 150 MWh, neither of G1 got any power allocated above the minimum. As before, in some cases, solutions were easily found, while in others the pricing scheme of the other company seemed practically intractable.

During such a bidding strategy however the equilibrium price was observed to drop and the equilibrium condition was violated (Fig. 3.4, b). for example, in the case on Table 2 the closing price dropped from \$19.92/MWh to \$17.54/MWh. In the case on Table 3 the closing price dropped from \$19.99/MWh to 18.20/MHw. This was so because in order to get power allocated one of the generators had to consistently bid under the equilibrium price. On the other hand, this led to the need of 2, 3, and sometimes up to 4 bidding cycles to price discovery, depending on how difficult the other bidding scheme was. A way to partially compensate this could be to raise the bids of the other two generators, which was not done in this simulation.

Such behavior is justified in cases where the malicious GENCO may first cut down one of the generators of the other company, then apply pricing schemes which will benefit its own generators, while keeping down the generator of the other GENCO. It would be unprofitable for this GENCO to start-up and shut-down these generators frequently even if it can price it in a way that will ensure power allocation from ICA for a short time.

During some of the runs it was observed that in some cases the malicious company obtained solutions which allowed the ICA to allocate power in a scheme that prevented the other company to fulfill the contract. Although this was not the primary intention of the simulation, such cases should also be considered. This was possible only under heavy loads, however. This situation arises when one of the generators is not allocated any power, and the load is so high that the rest of the generators could not make up for the total load contracted. This could be seen on Table 3.3. Because the first two generators did not get enough power allocated, the ICA transferred all the power to G3. Had G3 had a limit less that 180 MWh, this would be a problem. This was the reason why the maximum limit for G3 was set that high. It is possible, however, that an exclusion to be made in the algorithm and accommodate such situations.

Various market runs were performed with different loads, and transmission constraints. When the loads were high, it was very easy for the malicious company to reach a solution. However, as the loads were lowered below the level leading to congestion, solutions did not exist. In such cases the ICA simply did not reschedule any of the generators. This shows that the best opportunities for gaming are during peak demands for power. The results presented here were obtained with T1 = 150 MWh, T2 = 300 MWh, L1 = 200 MWh and L2 = 200 MWh. Also there were minimum loads of 10 MWh for all generators and maximum loads of 100 MWh for G1 and G2 and 200 MWh for G3 for each GENCO.

3.5 Market Power

Consideration should be given to assessing the market power of each seller. A seller is said to have unilateral market power if a unilateral deviation from a competitive equilibrium is profitable for that seller, given that all other traders continue to use the strategies that generated the competitive equilibrium [12]. The market power in this case is directly proportional to the market share. In a 10 seller market each seller has 10% of the market share, hence 10% market power, given that all sellers trade fixed amounts of electricity and all buyers are willing to buy regardless of the price. In this case both traders had equal market power (50% of the market share) which greatly simplified the modeling of the market and especially the congestion management of the power grid. However, in estimating the market power consideration here should be given to not only the market share of the company but also the information in possession of the company necessary to control a certain share of the market. If the malicious company did not possess information about the pricing schemes of the other company it would be impossible to come up with a solution leading to such devious behavior.

The situation would get complicated if unequal quantities of electricity are traded, if the buyers also have market power, if the number of buyers does not equal the number of sellers and each traded unit of electricity has different characteristics.

CHAPTER 4. PREDATORY GAMING STRATEGIES FOR ELECTRIC POWER MARKETS

4.1 Introduction

This chapter uses the same framework as in the previous one. Here a realistic power network is discussed.

4.2 Methods and Techniques

To keep the effect of different strategies obvious the modeled system is fairly small and easy to understand. The model consists of two GENCOs and two ESCOs. Each GENCO owns three generators attached to a simple 6-bus network. Fig. 4.1 depicts the network. G1, G3 and G5 belong to GENCO_1, and G2, G4 and G6 belong to GENCO_2. ESCOs are represented as two loads, L1 and L2, belonging to ESCO_1 and ESCO_2, respectively. The characteristics of the power network are given in Table 4.1 on the next page.

Different models have been built to simulate the electric marketplace. In this study the scheme outlined in Ref. [21] is used. A double-sided auction is constructed, with two buyers and two sellers of electricity. Computerized agents model the GENCOs and ESCOs as they participate in the simulated auction. The contracts for fixed amounts of electric energy are created for fixed time intervals. One of the trading agents uses a fixed set of rules to change the bidding strategy [23, 18]. The other is equipped with an intelligent searching mechanism.

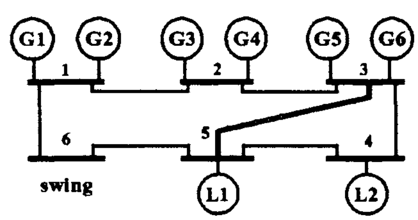

Fig. 4.1. A power network with two GENCOs' [(G1, G3, G5) and G2, G4, G6)] and two ESCO's [L1, L2].

Table 4.1. Power network data.

Bus Data

Bus Number	Mw Load	Mva Load	Voltage High Limit	Voltage Low Limit
1	0.0	0.00	1.05	0.95
2	0.0	0.00	1.05	0.95
3	0.0	0.00	1.05	0.95
4	100.0	0.00	1.05	0.95
5	100.0	0.00	1.05	0.95
6	0.0	0.00	1.05	0.95

Branch Data

From Bus	To Bus	R	X	Line Charg.	Mva Limit
1	2	0.01	0.01	0.01	1.0
2	3	0.01	0.01	0.01	1.0
3	4	0.01	0.01	0.01	1.0
3	5	0.01	0.01	0.01	1.0
4	5	0.01	0.01	0.01	1.0
5	6	0.01	0.01	0.01	1.0
6	1	0.01	0.01	0.01	1.0

The auction bid matching is performed by an independent entity, ICA, which handles the functions of the ISO.

Fig. 4.2 on the next page gives a flowchart describing the market simulation. One round of bidding proceeds as follows. The selling agents determine a bidding scheme based on certain criteria, such as load forecasting, estimation of the available transmission, etc. Then they submit their bids. The ICA then calculates the power each generator can deliver to the customer while attempting to minimize the total cost and enforce all network constraints. Thus the ICA allocates transmission rights to the cost-effective users. The ICA also determines the cost of transmission used to price the transmission network. Each user pays the cost for using the network.

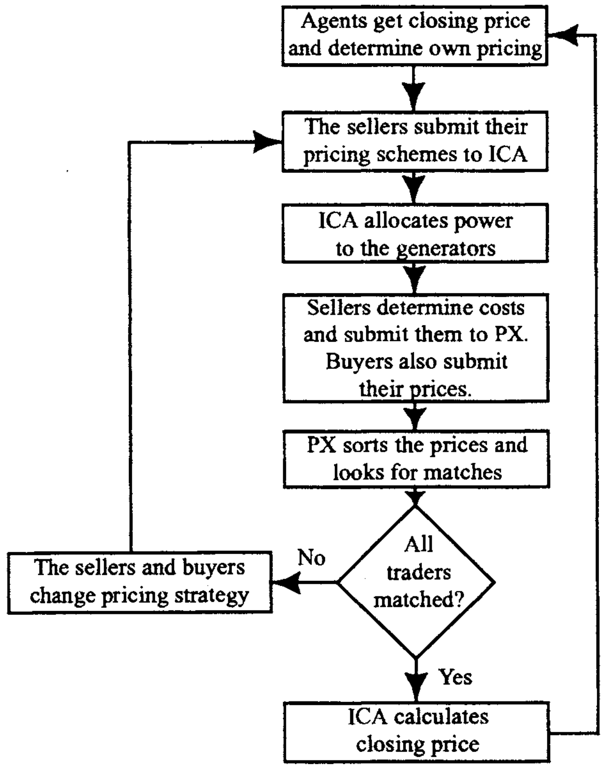

Fig. 4.2. Market flowchart.

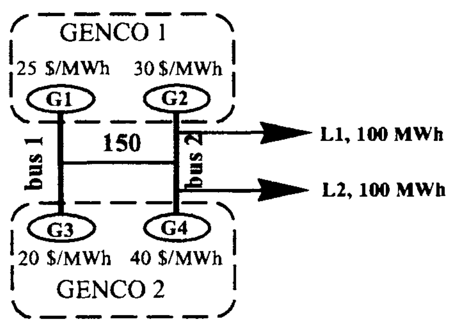

Fig. 4.3. A simple 3-bus example.

A very simple example can illustrate the role ICA plays on the market. Let us consider the 2-bus network on Fig. 4.3 on the next page with 4 generators serving 2 loads [16]. Two GENCOs own two generators each with 10.0 MWh minimum and 100 MWh maximum load. The transmission limit of the line from bus 1 to bus 2 is set at 150 MWh. Each GENCO is serving a separate 100 MWh load. Both GENCOs produce bids for its generators and submit the bidding scheme to the ICA. If there were no limits imposed on the transmission line from bus 1 to bus 2, the ICA would allocate 100 MWh transmission to G1 and G3 each, because they are priced lower that G2 and G4 and this would minimize the total cost. Thus, the cost for L1 would be \$2500/h and the cost for L2 would be \$2000/h.

However, if there is 150 MWh transmission limit on the transmission line, ICA would allocate 50 MWh to G1, 50 MWh to G2, and 100 MWh to G3. Thus, the cost for L1 becomes \$2750/h, and for L2 is \$2000/h plus they have to pay for congestion transmission capacity. The congestion transmission costs are transferred to the customers and included in the final price. They are calculated as follows. If ICA were to substitute 50 MWh from G3 with the same amount from G4, this would cost \$20/MWh to GENCO 2. Substituting 50 MWh from G1 with generation from G2 costs \$5/MWh to GENCO 1, therefore the transmission from bus 1 to bus 2 is much more valuable to GENCO 2, and ICA minimized the cost by rescheduling the generators of GENCO 1 instead of GENCO 2. Thus, GENCO 1 appears as the marginal user of the congested transmission line and sets the marginal cost of transmission at \$5/MWh. In that case GENCO 1 has to pay \$5/MWh × 50MWh = \$250.0 congested transmission costs for carrying 50 MWh over the congested line, and in the same fashion GENCO 2 has to pay \$500.0 for carrying 100 MWh over the congested line. These amounts are calculated in the total price the customer pays, and transferred to the transmission owner by the ICA. This pricing scheme is very important because it unbundles energy and transmission costs, thus creating two separate markets. This allows the generation companies to properly arrange their pricing schedule in a way that minimizes costs and assures available transmission. The process of forecasting and managing the transmission by the ICA is called congestion management [16].

After the GENCOs submit their bids and are allocated transmission rights, they calculate the price and submit their offers. Respectively, the buyers submit their bids. An independent entity, a power exchange, matches the bids and offers using an approach which maximizes the profit, similar to Wood and Wollenberg [17, 22] (Table 4.2). In that case the ICA does not analyze the market power each participant could have and it does not account for that.

All bids and offers are sorted in descending and ascending order and juxtaposed. If a buy bid is higher than the corresponding sell offer, the two players are matched, a contract is

Table 4.2. An example of an auction bid matching of 5 sellers/buyers, and the resulting outcomes

Bid	Offer	Quantity	Match	Midpoint	Eq. Price
10.91	10.01	1 MW/h	Yes	10.46	10.48
10.76	10.07	1 MW/h	Yes	10.41	10.48
10.74	10.38	1 MW/h	Yes	10.56	10.48
10.26	10.53	1 MW/h	No	N/A	N/A
10.05	10.67	1 MW/h	No	N/A	N/A

approved and the players are notified. In the case where the buy bid is lower than the sell offer, the players are rejected and notified. The rest of the players are also notified that the market was not cleared and prompted to bid again. This is repeated until all players are matched and the market cleared, or it is deemed that enough bidding cycles have occurred and the matched players granted contracts. The rest of the players are rejected and notified.

Each cycle the sellers resubmit offers by providing the contract administrator with the new bidding scheme and it recalculates the congestion prices, reschedules generators if necessary and the players submit bids based on the transmission rights. A matched bid and offer average for each contract, as well as an equilibrium price is calculated and the next round is called. Table 2 shows an example of 2 buyers and sellers after the bids and offers were matched by the power exchange. Different pricing techniques can be studied using the model described here. Midpoint prices are calculated, which is used if discriminatory pricing is desired, and if nondiscriminatory pricing is desired, the valid matches weighted by quantity are used to determine equilibrium price. Meanwhile the players update their strategy and the power exchange calls for the next round.

Each player updates its strategy according to its own forecasting techniques like load prediction, congested transmission forecast, etc. The players knows the outcome of its the past auctions (e.g. contract accepted/rejected), the bid/offer value, the congestion price and the closing price. If the offer was rejected, next time the seller lowers the price, choosing a pricing scheme it believes will bring maximum profit. Accordingly, the buyer bids are higher than the preceding round giving the agent a better chance to get a contract. If the offer was accepted, next time the seller offers a higher price, and the buyer bids are reduced. The players start trading around the previous round's closing price.

4.3 Experimental Design

The model was simulated using C and C++ modules. Separate objects were created for the ICA, traders, and the power exchange. This allowed for quickly and conveniently changing the rules which the separate entities on the market have to follow. Real / reactive power flow was run to calculate the line transmission and the limit on the line from bus 3 to bus 5 was set on 30.0 MW. The ICA minimized the total cost of the amount of electricity sold subject to that line constraint. There was not limit set on the rest of the power lines. The algorithm for calculating the power flow followed Refs. [17, 22].

At first both agents were equipped with similar bidding strategies. At the first step the agents get the closing price from the previous bidding round. Then the sellers set a bidding scheme that would result in a bid slightly below the closing price if all generators were allocated equal amount of power. In a real situation the bids for the separate generators of a particular GENCO would differ from each other due to numerous reasons like amount of rainfall, current weather conditions, etc. This was modeled by perturbing the bids around an average bid that would be slightly lower (3% in this case) than the closing price from the previous round. In the same fashion the buyers set their bids slightly higher (3% again) than the closing price. After the power was allocated by the ICA and bids/offers matched, if there were no price discovery, the sellers would lower the offers with 1% proportionally. The buyers in the same vein would raise the bids proportionally with 1%. This strategy usually led to price discovery on the 2-nd call at most.

After the round was completed, the ICA would announce the closing price calculated as the average of the two median prices of the contracts and call for the next round. Meanwhile the sellers would try to "forecast" the demand by repeating step one. The same was valid for the buyers. At first glance two randomly bidding GENCOs against each other do not provide much insight in the market behavior. Eventually, we could include results from previous

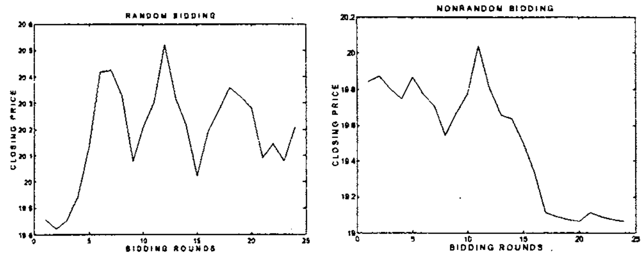

(a) (b)

Fig. 4.4. Closing price as a result of random (a) and nonrandom (b) bidding after 24 rounds

runs, and employ elaborate forecast techniques, but in that case the model was kept simple. Later in this work it becomes clear that the random "forecast" and also randomizing the bidding scheme is very beneficial in sense that it allows for unbiased checking of different strategies when one of the GENCOs is trying to outsmart the other. Also the routines for determining the bidding strategy could easily be fed with actual data.

The described system fairly quickly reached equilibrium (Fig. 4.4, a), and the equilibrium price converged to a limiting value. Such seemingly simple market behavior is very similar to what was observed in experimental double auction markets played by people [17, 18]. It can be seen that the initial price remained stable, and each round the market got cleared after the first or second call. The system did not diverge even after simulating a few years of bidding rounds, assuming that each round lasted an hour and players traded fixed amounts of electric power.

4.4 Results and Discussion

After building a stable model of the market, simulations were run where one of the GENCOs misbehaved. A simulation was run where GENCO 1 tried to find such pricing schemes so that one of the generators of the other company would get less that 10.0 MW power allocated. Table 3 represents two such solutions. Solutions in most cases were easily found, while in others the pricing scheme of the other company seemed practically intractable. In the case on Table 4.2, the targeted generator was G3 of GENCO 2. As we see, although it

Table 4.3. Two solutions, GENCO 2 / G3 gets no power.

GENCO 1	G1	G2	G3
Bid	19.77	19.54	19.65
Power	58.48	17.00	24.51
GENCO 2	G1	G2	G3
Bid	19.74	19.82	19.55
Power	67.70	31.22	1.068
GENCO 1	G1	G2	G3
Bid	19.97	19.12	19.55
Power	61.94	33.31	4.743
GENCO 2	G1	G2	G3
Bid	19.77	19.64	20.07
Power	85.94	11.35	2.701

submitted the lowest bid from all 3 generators, it got almost no power allocated, due to the pricing scheme submitted by the other company.

During such a bidding strategy however the equilibrium price was observed to drop and the equilibrium condition was violated (Fig. 4.4, b). For example, in the case on Table 2 the closing price dropped from $20.03/MWh to $19.41 /MWh in just 5 bidding periods. This was so because in order to get power allocated one of the generators had to consistently bid under the equilibrium price. On the other hand, this led to the need of 2, 3, and sometimes up to 4 bidding cycles to price discovery, depending on how difficult the other bidding scheme was. A way to partially compensate this could be to raise the bids of the other two generators, which was not done in this simulation.

Such behavior is justified in cases where the malicious GENCO may first cut down one of the generators of the other company, then apply pricing schemes which will benefit its own generators, while keeping down the generator of the other GENCO. It would be unprofitable for this GENCO to start-up and shut-down these generators frequently even if it can price it in a way that will ensure power allocation from ICA for a short time.

Another problem that arises is that the solutions are far from optimal. As we can see in Table 3, for both GENCOs the generators submitting the lowest offers get less power allocated that the ones submitting more expensive offers. Which means that such gaming behavior not only would hurt the CENGO targeted, but also would rise the cost of electricity produced in very short run.

4.5 Market Power

Consideration should be given to assessing the market power of each seller. A seller is said to have unilateral market power if a unilateral deviation from a competitive equilibrium is profitable for that seller, given that all other traders continue to use the strategies that generated the competitive equilibrium [19]. The market power in this case is directly proportional to the market share. In a 10 seller market each seller has 10% of the market share, hence 10% market power, given that all sellers trade fixed amounts of electricity and all buyers are willing to buy regardless of the price. In this case both traders had equal market power (50% of the market share) which greatly simplified the modeling of the market and especially the congestion management of the power grid. However, in estimating the market power consideration here should be given to not only the market share of the company but also the information in possession of the company necessary to control a certain share of the market. If the malicious company did not possess information about the pricing schemes of the other company it would be impossible to come up with a solution leading to such devious behavior.

The situation would get complicated if unequal quantities of electricity are traded, if the buyers also have market power, if the number of buyers does not equal the number of sellers and each traded unit of electricity has different characteristics.

CHAPTER 5. POWER AUCTIONS BID GENERATION WITH ADAPTIVE AGENTS USING GENETIC PROGRAMMING

5.1 Introduction

In this chapter attention is paid to the study of the possibility of developing an agent capable of discovering a set of rules that allow for profitable participation on the market. The agent is modeled using genetic programming techniques.

The chapter is organized as follows. Section 5.2 presents the techniques and methods used to model the auction market and the participants. Section 5.3 describes a simple experiment that illustrates how a rule-based decision tree could be developed. Section 5.4 presents an implementation of a full-blown agent playing on the market. Section 5.5 discusses the results achieved. Finally, Section 5.6 presents a conclusion and possible extensions to this topic.

5.2 Methods and Techniques

To keep the effect of different strategies obvious the modeled system is fairly small and easy to understand. The base model consists of ten GENCOs and ten ESCOs. Each GENCO has a fixed amount of electricity to sell and each ESCO has a fixed amount of electricity to buy. Different models have been built to simulate the electric marketplace. In this study the scheme outlined in Richter, Sheblé and Ashlock [6] is used. A bilateral auction is constructed, with a fixed number of buyers and sellers of electricity. The bilateral contracts for fixed amount of electric energy are performed in fixed time intervals. The transmission system is lossless and has unlimited capacity. The trading agents use a fixed set of rules to change the bidding strategy [25,29]. The auction bid matching is performed by an independent third side.

One round of bidding proceeds as follows. All buying/selling agents submit their bids/ offers to a third party, an Independent Contract Administrator (ICA). ICA matches the bids

Table 5.1. An example of an auction bid matching of 5 sellers/buyers, and the resulting outcomes

Bid	Offer	Quantity	Match	Midpoint	Eq. Price
10.91	10.01	1 MW/h	Yes	10.46	10.48
10.76	10.07	1 MW/h	Yes	10.41	10.48
10.74	10.38	1 MW/h	Yes	10.56	10.48
10.26	10.53	1 MW/h	No	N/A	N/A
10.05	10.67	1 MW/h	No	N/A	N/A

using an approach similar to Wood and Wollenberg [27,28]. All bids and offers are sorted in descending and ascending order and juxtaposed. If a buy bid is higher than the corresponding sell offer, the two players are matched, a contract is approved and the players are notified. In the case where the buy bid is lower than the sell offer, the players are rejected and notified. There is no second call for bidding regardless of the number of contracts approved. A bid and offer midpoint, as well as an equilibrium price is calculated and the next round is called.

Table 5.1 shows an example of 5 buyers and sellers after the bids and offers were matched by ICA. Different pricing techniques can be studied using the model described here. Midpoint prices are calculated, which is used if discriminatory pricing is desired, and if nondiscriminatory pricing is desired, the valid matches weighted by quantity are used to determine equilibrium price. Meanwhile the players update their strategy and the ICA calls for the next round.

Each player updates its strategy according to the outcome of the previous round. The player knows only the outcome of its own last action (e.g. contract accepted/rejected) and remembers the bid/ offer value. If the offer was rejected, next time the seller lowers the price, choosing the value randomly within an interval it calculates based on a rule including its own last offer. Accordingly, the buyer bids are higher than the preceding round. If the offer was accepted, the seller offers a higher price based on a rule including the last accepted offer, and the buyer bids are reduced. All players start trading around a fixed equilibrium price submitted by ICA. This is done for convenience. In a real world situation the players can start bidding around the closing price for the previous day, for example. After the first round the ICA takes care only of matching players. The actions of a seller and buyer depending on the outcome of the last round are described in [24]. In any event, we must set upBid < downBid, and upOffer < downOffer in order to reach convergence within reasonable time. If it were otherwise the price would not converge.

5.3 Experimental Design

The described system fairly quickly reaches equilibrium, and the buy/sell spread, the midpoint price and the equilibrium price converge to a limiting value. Fig. 5.1 shows the pictorial representation of an auction consisting of 100 players, evenly divided into buyers and sellers of electricity.

The auction went through 10 rounds. In this case downBid and upOffer were set at 5% of the last unsuccessful bid/offer and downOffer and upBid were set to 3% of the last successful bid/offer. All four price updates were tied up to a random generator, i.e. the players may choose price updates with equal probability within the prescribed price update interval.

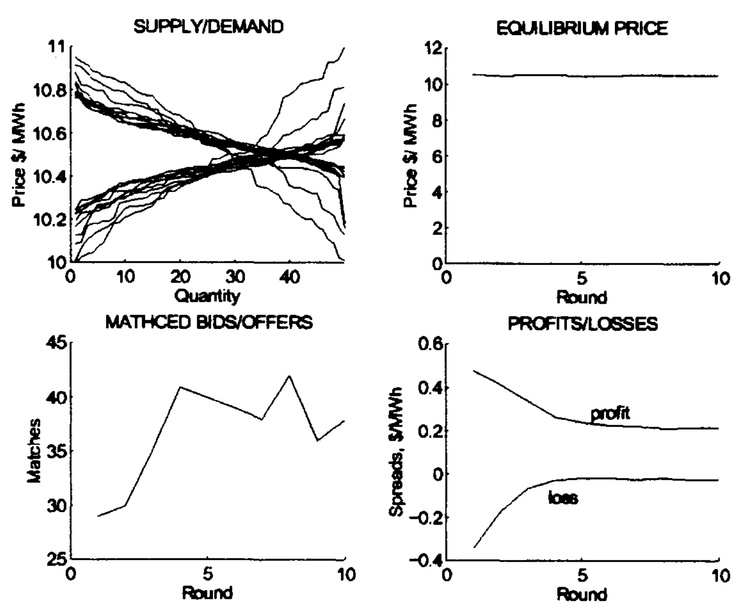

Fig. 5.1. Supply/Demand, equilibrium price, offers accepted and profit/loss spread for 50 buyers/sellers of electric power after 10 bidding rounds.

Such seemingly simple market behavior is very similar to what was observed in experimental double auction markets played by people [27]. It can be seen that the initial price submitted by ICA remained stable, and the number of accepted offers quickly reached about 70% of the total number of offers submitted, up from about 50%. After only a few rounds the players minimized the losses to almost zero. However, the profits also dropped. The system did not diverge even after simulating a few years of bidding rounds, assuming that each round lasted an hour and players traded fixed amounts of electric power. Fig. 5.1 depicts the auction behavior after 10 bidding rounds.

After building a stable fixed-rule auction, one of the players was replaced with an agent which evolves a set of rules. The agent evolved a set of rules using a genetic programming algorithm. A genetic program (GP) is an adaptive learning system based on many of the principles of the genetic algorithm. The rules are represented by a decision tree which output is the bid submitted to the auction. The genetic programming algorithm evolving the tree is based on the LISP work of John Koza [32]. The actual implementation was done using the lil-gp library from the Michigan State University [33]. The library is an implementation of the work of Koza in C and can handle very large problems.

Table 5.2. Functions used in building an agent.

Function	Symbol	Nodes	Return Value
Summation	+	2	Real
Subtration	–	2	Real
Multiplication	*	2	Real
Division	/	2	Real
Greater	>	2	Boolean
Less	<	2	Boolean
Equal	=	2	Boolean
If-Then-Else	if	3	Boolean
And	and	2	Boolean
Or	or	2	Boolean

The agent was built using functions and terminal nodes. For functions the following algebraic and logical operators were used: summation, subtraction, multiplication, division, "if-then-else", "and", "or". Table 5.2 below gives the functions used. Each function has terminal nodes and returns a value.

Tree terminal nodes were specified. The first was EqPrice, taking real values, containing the equilibrium price from the last run. The second one was LastContract, taking boolean values, reflecting the success or failure to make a contract during the last run. The third terminal node was a random ephemeral constant, taking any real value between 0 and 20.

The agent (seller) was attached to the auction and the auction was evaluated once for every individual of the GP population. In our case we had a population of 10000 individuals, and the following GP operators were performed on the population: reproduction, crossover and mutation. The probability of reproduction was 0.2, the probability of crossover was 0.8 and the probability of mutation was 0.1. For each individual the auction was run for 100 rounds and the fitness was directly proportional to the number of contracts the GP agent was able to make during the auction. The parents were selected based on the fitness of the individuals in the population and the tournament method of selection was used with a parent chosen as the best among 10 randomly picked individuals.

5.4 Results and Discussion

The agent performed very well and consistently surpassed the fixed rule agents during all separate runs of the simulation. The profits realized were the same as the fixed rule agents. On average, however, the GP agent scored 91 out of 100 possible contracts compared with 80 of the fixed-rule agents. Below is given a typical decision tree that scored 92 out of 100. The tree had the following parameters: generation: 38, nodes: 98, depth: 6, hits: 92 out of 100. The total calculation time was 0.044 MIPS years (2 hours on Pentium/233MHz).

GP tree:

(< (if EqPrice EqPrice LastContract)

```
(+ (+ (< (and (> LastContract EqPrice)
(* EqPrice EqPrice))
(* (if 2.598 LastContract 7.187)
(and 16.128 EqPrice)))
(or (< (> EqPrice LastContract)
(or LastContract 9.518))
(/ (or 16.687 LastContract) 7.246)))
(if (and (+ (if LastContract LastContract
(and (if 15.976 11.389 EqPrice)
(- LastContract 8.803)))
(and (+ EqPrice EqPrice)
(+ LastContract 14.034)))
(and (> EqPrice LastContract)
(if 11.176 7.207 5.685)))
(+ (- (< LastContract 0.068)
(and LastContract LastContract))
(and 19.414 LastContract))
(or (and (- EqPrice 3.832)
(or 16.847 EqPrice))
(or (> EqPrice EqPrice)
(- LastContract 3.629))))))
```

The only way a seller agent could perform better that the hardcoded-rule seller is if it:

1) "discovers" the hardcoded rule that the rest of the agents are using, and

2) Modifies the upBid and downBid using the results from the last big and the equilibrium price to his advantage

Although it would be extremely cumbersome to debug the above generated rule, a few things about it could be observed. First, it takes into account the fact that the last contract was accepted or rejected. Second, it takes into account the equilibrium price, and third, it involves a large number of generated constants which correspond to the upBid and downBid used by the hardcoded rule agents. Thus, there is a great similarity between the way the hardcoded rule is used and the generated rule is used.

A future direction of that work could be an attempt to introduce a "complexity punishment" in the fitness function. That is, a simple rule would be preferred before a complex rule and there would be a tendency for an agent to develop not only a good bidding rule, but also a simpler representation of that rule. The outcome would be a rule set that is very easy to debug, explain and programmatically implement. The other benefit of that would be using the experiment as a "rule discovery" engine. In other words, the goal is to easier discover the rules the rest of the agent are using which in the above work is a sub-product of the experimental results.

CHAPTER 6. ELECTRIC POWER MARKET AUCTION WITH A GENETIC ALGORITHM AS A LEARNING MECHANISM FOR AUCTION PARTICIPATION

6.1 Introduction

Genetic algorithms have proven to be useful for solving optimization problems in which the solution space has discontinuities rather than the monotonous and continuous functions preferred by most hill-climbing algorithms. The problem of bidding/asking for contracts in a double-sided auction without going bankrupt is an example of such a discontinuous search space. The fact that there are a finite number of players against whom an agent competes is partly responsible for the discretizing of the solution space. Double-sided auctions are one method proposed for determining the clearing prices for service and supply contracts in the deregulated electricity markets [34]. At present, the exact rules and mechanism of this auction are being tested in various markets in real time. It is of great interest to regulators and the companies who buy and sell electricity whether this auction mechanism can result in a fair market, whether the prices discovered are near to the theoretical equilibrium achieved through perfect competition, and if it is possible for a firm to exercise market power either by being able to supply a majority of the product or by having production costs much cheaper than its rivals [37].

The market simulator should provide the researchers with a means to study market power on electricity markets. Measuring market power of participants on markets with few players (three or less buyers and three or less sellers) is crucial to understanding the emerging market structure in a restructured environment.

One very important issue that can be explored with the new simulator is studying the difference between structural and strategic market power. The structural market power is due to market rules and the strategic power is due to the ability of individual agents to learn and adapt [36].

6.2 The Market Simulator

An electric market simulator was created similar to the EPRI simulator described in [36, 37]. The market simulator simulates a simple electricity auction market by modeling learning agents with a genetic algorithm (GA). It allows for simulating a double-sided auction of electricity in a day-ahead or hour-ahead market who operate on a fixed power grid. Sellers and buyers represent generating companies (GENCOs) and energy services companies (ESCOs). The power grid is a fully interconnected graph with buyers and sellers at the nodes and transmission lines as edges. Each node is assigned power to buy/sell and each edge is assigned certain available transmission capability (ATC).

The traders are divided into buyers and sellers and each trader is assigned certain ATC with respect to all other traders. The following parameters can be specified for each buyer: maximum amount of power to buy in MWh, revenue per MWh, and fixed cost. For each seller the following can be specified: maximum amount of power to sell in MWh, cost per MWh, and fixed cost.

All buyers and sellers trade electric power in a double auction run by an independent entity. The auction is performed in rounds. In each round all traders submit their bids/asks simultaneously and get matched. The criterion used for matching consists of maximizing the total profit of each round. For that purpose all buyers and sellers are sorted in descending/ ascending fashion by their bids/asks. Then they are matched and contracts assigned depending on the ATC and amount of available power they can trade.

The genetic algorithm (GA) evolves traders' actions (bids or asks.) Each action is represented by a 10 bit string mapping to a floating point number in the interval [0, 1.0) with step 1/2^10 . Thus, we can have at most 2^10 possible actions. Each resulting number is then multiplied by a constant to obtain the monetary markup amount traders will use to play the auction. For example, let the constant be $40.00. Seller asking prices are restricted to the range [marginal cost, marginal cost + $40.00). Buyers offer prices are restricted to the range (marginal revenue - $40.00, marginal revenue]. By preventing sellers from selling below their marginal cost and buyers from buying above their marginal revenue, the simulated agents bid in a the same range as a rational bidder would bid in the real world (i.e. so as not to lose money).

The fitness of each trader is proportional to the profit made in the auction round and is recalculated every round. It is used to select and sort traders. Each auction round represents a generation. The traders participating on the market are divided into buyers and sellers. Each trader evolves its own actions as a population. Three genetic operators are performed on the populations: reproduction, crossover and mutation. The probability and type of operators will be defined later. Information is exchanged solely within each trader. There is no explicit information exchange among separate traders.

A single auction round proceeds as follows. First, the population's bids/asks are initialized with random values. Then all ATCs of the power grid are specified at fixed values as well as the capacity levels that determine the maximum quantities that sellers and buyers can sell and purchase, respectively. All buyers and sellers are sorted by bids/asks in descending/ascending order. Before each sorting the players' order is randomized. This is necessary in case some buyers or sellers bid or ask the same amount. This is highly probable given the fact that the bid/ask space is limited and the population number is often low (6 to 12 players).

The buyer with the highest bid is matched with the sellers with the lowest ask. If there is nonzero ATC, the seller is matched with the buyer for the amount of electricity calculated as the minimum of three amounts: ATC, the amount the buyer can buy, or the amount the seller can sell. Thus, if the transmission capability is 5 MW, the seller can sell 20 MW and the buyer can buy 10 MW, the contract is made for 5 MW only because this is the maximum the grid can support. Then the carry-over transmission capability and carry-over amount to buy and sell for each case are calculated. The unit price for the contract is set at the midpoint of the bid and ask price. Then the next lower pair is matched.

Figure 6.1 shows an example of three sellers and two buyers. In this case, Seller 1 sells 8 MW to Buyer 1 at a price of $7.5/MW, Seller 1 sells his last 2 MW to Buyer 2 at a price of $6.5/MW, and Seller 2 sells 5 MW to Buyer 2 at $7.5/MW. Buyer 2 still lacks 3 MW but does not buy from Seller 3 because Seller 3 is too expensive for him.

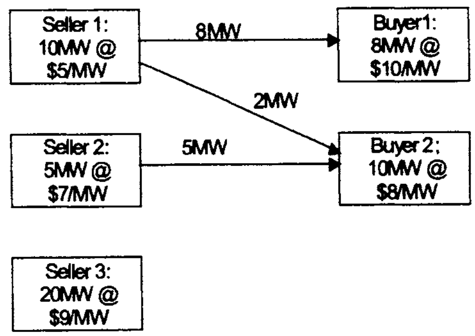

Fig. 6.1. Example of a simple auction.

After each paired matching of buyers and sellers, a few data are calculated: average price of the contract, total contract value and corresponding profits. The unit price of the contract is calculated as the average of the buy bid (bi) and sell ask (sj) as shown in Figure 6.2. The total contract value is obtained by multiplying the amount of electricity contracted by the unit price. Buyers' profit is calculated as the difference of total revenue and total contract amount. Sellers' profit is calculated as the difference of the total contract value and total cost of the contract.

After all buyers and sellers have been through the matching process, the average profit for all buyers and sellers, as well as the average bid and ask are calculated for statistical

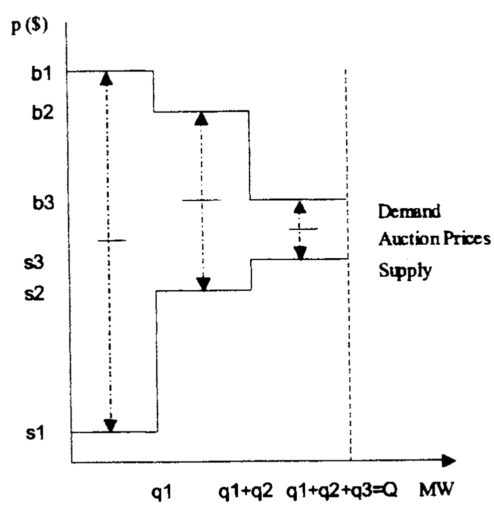

Fig. 6.2. Example of bids/asks in an auction for three sellers and three buyers.

purposes. Each buyer and seller is assigned a fitness score proportional to the profit earned by this trader in the preceding round.

6.3 Implementation

There are NS number of sellers and NB number of Buyers. Each one of sellers and buyers has certain amount of electricity to sell MWS and buy MWB respectively. Each seller has associated marginal cost MC and each buyer has associated marginal revenue MR. Also, each seller has associated fixed cost FCS and each buyer has also associated fixed cost FCB. The amount of electricity they can trade is constrained by the available transmission capability matrix ATC with size NS x NB.

Overall, a single run of the auction proceeds as follows. The sellers submit their ask prices for up to the amounts of electricity they can sell. The buyers submit their bid prices for up to the amounts of electricity they can buy. The prices submitted are fixed and do not depend on the final amount of electricity being traded. sellers' prices are sorted in ascending order and buyers' prices are sorted in descending order. Then the prices are matched. For each matched seller-byuer pair a contract is granted or refused and if a contract is granted, a final price for each contract is calculated. Then for all matched seller-buyer pairs the amount of electricity traded is calculated based on the ATC, MWB and MWS. After that the profits and revenues are calculated and applied to each trader's account.

Once the profits and revenues are applied, fitness measure for the particular bid or ask price is calculated based on the amount of profit or revenue generated during the run. Then the bid and ask prices are modified by a genetic algorithm in order to generate better ones which would eventually produce higher profits or revenues on the next run.

Each seller and buyer is implemented by a data structure consisting of six elements:

- 10 bit string of 0-s and 1-s (boolean)
- Index (integer)
- Real price (floating point)
- Fitness (floating point)

- Profits (floating point)

- Amount of electricity to trade (floating point)

The following data are supplied by the user.

(1) An ATC matrix containing the available transmission capacities that can be traded.

(2) Marginal cost vector MC containing the marginal cost for each seller.

(3) Marginal revenue vector MR containing the marginal revenue for each buyer.

(4) Fixed cost vector FCS containing the fixed cost for each seller.

(5) Fixed cost vector FCB containing the fixed cost for each buyer.

(6) Amount of electricity to sell vector MWS containing the available amount of electricity each seller can sell.

(7) Amount of electricity to buy vector MWB containing the available amount of electricity each buyer can buy.

(8) NS and NB integers specifying the number of sellers and buyers on the market.

(9) PricingType boolean specifying the pricing scheme used (Uniform or Discriminatory).

(10) A floating point number SM (soft multiplier) for scaling the price markup of each trader.

(11) An integer AR specifying the number of auction runs.

(12) An integer S specifying the number of prices (population number) a trader can pick from.

At the beginning all variables are initialized from a GUI. Then a data structure containing NS number of sellers and a data structure containing NB number of buyers are instantiated. Each seller and buyer has S number of prices he can pick from.

The data structures are initialized as follows. First all 10-bit strings are randomly initialized with 1-s and 0-s. Second, all Index fields are given unique consecutive numbers. After that the real ask and bid prices corresponding to each bitstring are calculated as follows. Each 10-bit string (which from now on we will call bitstring) is converted to a floating point number FPN in the interval [0, 1.0). Then the resulting number is multiplied by SM. The final ask and bid price is calculated in the following fashion. For the sellers, for each sellers' bitstring we have an ask price AP calculated as

AP = MC + SM * FPN + FCS

For the buyers, for each buyer's bitstring we have a bid price BP calculated as

BP = MR - SM * FPN - FCB

After that the fitness fields are initialized with zeroes, the profit fields are initialized with zeroes and the quantity fields for electricity to trade are initialized with the MWS and MWB values.

Once the traders are created and initialized, the auction may begin. Consecutively, NS sellers and NB buyers are chosen from within the populations. They are submitted to a function which matches them, grants or refuses contracts, computes trading prices, and calculates profits and revenues. The function proceeds as follows. First, the sellers are sorted in ascending order according to their ask prices. The buyers are sorted in descending order according to their bid prices. Then the sellers and buyers are matched.

(1) The highest ask seller is matched with the lowest bid buyer. If the bid price is higher than the ask, the amounts of electricity to sell and to buy and the ATC between these two traders is checked. If all of them are greater than zero, a contract is granted. The final amount of electricity traded is the minimum of all three amounts (MWS, MWB and ATC for these particular traders). The contracts are written to a boolean matrix containing all contracts for that run, and the quantities traded are written to another matrix containing all quantities traded for that run. Then for the seller the amount sold is subtracted from the amount available to sell and for the buyer the amount bought is subtracted from the amount available to buy.

(2) After that if the seller still has electricity available for sale, he is matched with the next buyer with lower bid price.

The above procedures 1 and 2 are repeated until the seller runs out of all electricity available for sale or runs out of buyers available to buy his electricity.

(3) The next seller with higher ask price is picked and the above procedures 1 and 2 are repeated.

Procedures 1, 2 and 3 are repeated until all sellers and buyers are matched. Once this is completed, the profits and revenues are calculated as follows. For every contract made a price is computed as the average of the bid and ask price. For each matched seller a profit P is calculated:

P = (Price - MC) * Quantity - FCS

For each matched buyer a revenue R is calculated:

R = (MR - Price) * Quantity - FCB

After calculating the profits and revenues, a new auction round is called. After certain number of auction rounds (auction batch) which equals the number of individuals in the population, each individual in the populations has accumulated certain profits or revenues. At the end of each batch of rounds, crossover and mutation are applied to each trader individually. For every trader's population of markups, two parents' bitstrings B1 and B2 are selected using tournament selection. The tournament is being conducted by picking two bitstrings at random and comparing their respective profits/revenues generated during the auction round.. The individual with the higher profit is the winner. After selecting the parents, a crossover is performed. There are 3 types of crossover available: single point, two point and uniform.

The single point crossover is performed as follows. An integer CP within the interval [1,10] is generated. Then a bitstring is created by taking the bits of B1 from position 1 to CP and by taking the bits of B2 from position (CP + 1) to 10. Then the bitstring is stored and this is repeated until as many children are generated as necessary to fill the new population of markups. After generating all new bitstrings, the old bitstrings are replaced with the new bitstrings.

The two-point crossover is performed as follows. Two integers CP1 and CP2 within the interval [1,10] are generated. Then a bitstring is created by taking the bits of B1 from position

1 to CP1, by taking the bits of B2 from position (CP1 + 1) to CP2 and by taking the bits of B1from position (CP2 + 1) to 10. Then the bitstring is stored and this is repeated until as many children are generated as necessary to fill the new population of markups. After generating all new bitstrings, the old bitstrings are replaced with the new bitstrings.

The uniform crossover is performed as follows. A bitstring is created by alternatively taking every odd bit of B1 and by taking every even bit of B2. Then the bitstring is stored and this is repeated until as many children are generated as necessary to fill the new population of markups. After generating all new bitstrings, the old bitstrings are replaced with the new bitstrings.

Then, a bitwise mutation is performed. The mutation is defined as flipping a bit from 0 to 1 or from 1 to 0. Each bit is mutated with certain probability, which is uniform for all bits.

Then all profits and fitnesses are reset to zero, the bid or ask prices are recalculated according to the new bitstrings and all amounts of electricity to buy or sell are reset to the initial values.

An elitism scheme is also implemented. The elitism scheme retains the top performing individuals from each population, and copies them to the new population. The rest of the population is filled with individuals generated by crossover and mutation as described above. The percentage of top performing individuals to be retained is set at the beginning of the auction run. The elite is not mutated.

6.4 Design Discussion

The EPRI market simulator discussed in [3] and [4] showed certain shortcomings. Information was exchanged among traders as a way of learning. However, each trader was solving his own optimization problem and information was particular to his parameters like marginal cost, amount of power to buy/sell and ATC. When exchanged, the information was assimilated as is, which impeded the learning process. Also, when there were very few players (three or less buyers and three or less sellers) on the market (monopoly, oligopoly), the GA in the way it was implemented was inherently unsuitable to solve the problem. In the design described in this work, these problems were overcame by assigning separate gene pools to each trader. The traders learn to participate on the market by evolving their own action pools. This separation of evolving action sets is much more suitable to solving real world trading of electricity that the previous EPRI market simulator.

6.5 Experiments

A set of experiments is presented here, which illustrates the advantages of the new design. It consists of a market with a single seller and many buyers. This experiment shows solving the problem with low number of traders. The seller's capacity to sell was gradually changed at 10 steps from less that the combined buying capacity of all buyers to more than the combined buyers' capacity. A graph (Fig. 3.) was constructed with the results from all separate runs, that shows the dependence of the average price on the ratio of sell/buy capacity. This experiment cannot be run by the EPRI market simulator.

The seller's capacity was changed from 10 MW/h to 100 MW/h. All buyers' capacities were fixed. Seller's marginal cost was set at 10 $/MW and all buyers' marginal revenues were set at 20 $/MW. Fig. 6.3 represents the result of the experiments. Ten separate runs were

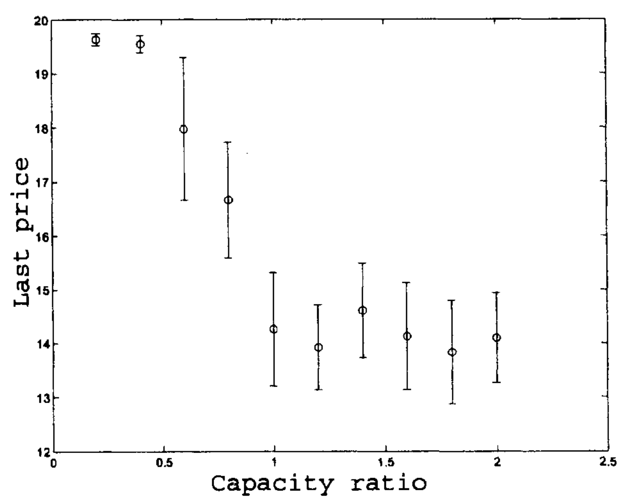

Fig. 6.3. A set of runs with different sellers/buyers capacity ratio.

performed, and each of them was repeated 100 times to gather data for statistics. The X axis represents the ratio of seller/buyer capacity. The Y axis is the price of the last contract matched at the end of each auction run. Each auction run consisted of 200 auction rounds. The population number for each trader was 20, the mutation rate was 3 %, the batch size was 20 rounds, the crossover was two-point, with no elitism. Discriminatory pricing scheme was used.

As we can see, the price at which the last contract was executed drops from 20 $/MW to about 13 $MW as the ratio of the seller/buyer capacity increases. This illustrates the supply/demand law. When the amount of electricity offered on the market is low, the price is as high as the buyers are able to pay. Here the price can be as high as the buyers' marginal revenue. As the amount offered increases, the price drops. However, in our case the price does not drop down to the marginal cost of the seller because there is only one seller which has a monopoly. There is no competition, so the price drops to a bit below the marginal cost / marginal revenue midpoint and stays there no matter how much is the excess supply of electricity on the market.

CHAPTER 7. BUILDING ELECTRIC POWER AUCTIONS WITH IMPROVED ROTH-EREV REINFORCED LEARNING

7.1 Introduction

The goal of this chapter is to model an agent driven bilateral power market auction where the players are represented by agents using a much-improved modification of the Roth-Erev reinforcement-learning algorithm. An Available Transmission Capacities (ATC) matrix represents the underlying physical power network. The agents have to learn how to bid on the auction and respectively develop a winning strategy. The improvements made in the reinforcement algorithm make it suitable for robust industrial applications. The main purpose of this work is to build a consistent power auction that can be experimented with. The power auction consists of buyers (i.e., energy service companies or ESCOs) and sellers (i.e., generation companies or GENCOs) that can intelligently participate on the power market. To this end, a computational electricity market is constructed that can be used as a laboratory for systematic experimentation.

7.2 Computational Design Methods

The computational electricity market models the trading of electricity by traders attached to an electric power grid. The power grid is a fully connected graph with traders as the nodes and transmission lines as the edges. Each trader is assigned electricity to buy or sell as well as a certain Available Transfer Capacity (ATC) with respect to each other trader. Traders with electricity to buy are referred to as buyers, and traders with electricity to sell are referred to as sellers. The following parameter values are specified for each buyer: capacity in MWh (maximum amount of electricity that can be resold in a secondary retail electricity market); marginal revenue per MWh purchased and resold; and fixed cost (set to zero). Also, the following parameter values are specified for each seller: capacity in MWh (maximum amount of electricity that can be produced); marginal cost per MWh produced; and fixed cost (set to zero). These parameter values are private to each trader. The buyers and sellers trade electricity in a discriminatory-price double auction run by an independent clearinghouse, henceforth simply referred to as a discriminatory auction.

The goal of each buyer and seller is to maximize its own profits. The discriminatory auction is performed in rounds. In each auction round, the buyers and sellers simultaneously submit their bids (offers to buy) and asks (offers to sell) to the clearinghouse. Each bid and ask consists of a single price-quantity pair. The clearinghouse then matches these bids and asks, using as its criterion the maximization of total profit. To accomplish this, all buyers and sellers are sorted in descending and ascending order by their bid and ask prices, respectively. As detailed more fully below, bids and asks are then matched in order, and contracts are assigned depending on the amount of available electricity that the buyers and sellers are willing and able to trade. The price of each contract is set at the midpoint of the buyer's bid price and the seller's ask price, i.e., at the midpoint of the bid-ask spread.

A single auction round proceeds as follows. First, the population's bids/asks are initialized with random values. Then all ATCs of the power grid are specified at fixed values as well as the capacity levels that determine the maximum quantities that sellers and buyers can sell and purchase, respectively. All buyers and sellers are sorted by bids/asks in descending/ascending order. Before each sorting the players' order is randomized. This is necessary in case some buyers or sellers bid or ask the same amount. This is highly probable given the fact that the bid/ask space is limited and the number of market participants is often low (6 to 12 players).

The buyer with the highest bid is matched with the sellers with the lowest ask. If there is nonzero ATC, the seller is matched with the buyer for the amount of electricity calculated as the minimum of three amounts: ATC, the amount the buyer can buy, or the amount the seller can sell. Thus, if the transmission capability is 5 MW, the seller can sell 20 MW and the buyer can buy 10 MW, the contract is made for 5 MW only because this is the maximum the grid can support. Then the carry-over transmission capability and carry-over amount to buy and sell for each case are calculated. The unit price for the contract is set at the midpoint of the bid and ask price. Then the next lower pair is matched.

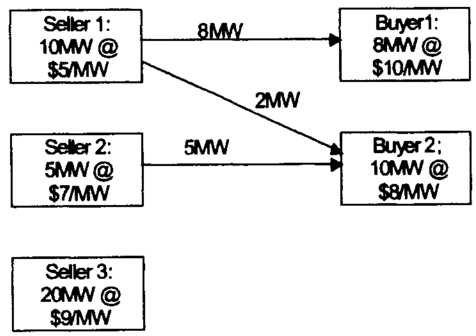

Fig. 7.1. Example of a simple auction.

Figure 7.1 shows an example of three sellers and two buyers. In this case, Seller 1 sells 8 MW to Buyer 1 at a price of $7.5/MW, Seller 1 sells his last 2 MW to Buyer 2 at a price of $6.5/MW, and Seller 2 sells 5 MW to Buyer 2 at $7.5/MW. Buyer 2 still lacks 3 MW but does not buy from Seller 3 because Seller 3 is too expensive for him.

At the end of each auction round, the performance of each trader is calculated. The performance of each trader is directly proportional to the profit earned by this trader in the preceding round; profit performance in prior rounds is not taken into account. This reflects the assumption that each trader is myopic; at any given time, a trader only knows his last bid or ask and the amount of profit earned on this bid or ask, so he can- not take advantage of information in any prior rounds. Also, in all experiments reported below, the linearity assumed for the revenue and cost functions ensures that an optimizing trader would always desire to bid or ask for its capacity quantity. Consequently, only the bid and ask prices of the traders are assumed to evolve. After the completion of this step, another auction round commences. A detailed description of the market mechanism is given in Chapter 6.

7.3 Improved reinforced learning algorithm

In the economics community, relatively limited amount of work is being done in developing learning algorithms for game strategies. The work done is largely restricted to single-player games. One exception is the work of Roth and Erev [38, 39] in what they call "cognitive game theory." They distinguish between "low" and "high" game theory. The low game theory assumes that the players have limited rationality but adapt, learning from their own experience by playing the game, thus deriving rational strategies. Also, the players may not consider all strategies available to them, and they may not be "subjective expected utility maximizers" [38]. The high game theory assumes that the players already have humanlike rationality and can deduce the rational strategy from the rules of the game. Roth and Erev focus their research on low rationality game theory, aiming to model the natural learning process usually observed in humans. Their goal is to better understand the nature of different economic games and the limitations of learning rational (i.e., optimal) strategies to play these games. According to Roth and Erev, if player P_i plays strategy S_{ij} and receives payoff r, then the fitness of that strategy S_{ik} is updated as

$$q_{ik} = (1-\phi)q_{ik} + E(k, R(\rho))$$

where ϕ is a recency, $R(\rho)$ is a mapping of payoff r into a nonnegative reinforcement signal (in the simplest case, $R(\rho) = \rho$), and E is an update value reflecting the experience gained from play. From here on we will refer to the fitness as propensity. Roth and Erev define E to be

$$E_j(k, R(\rho)) = \begin{cases} R(\rho)(1-\varepsilon) & \mathit{if}\ k = j \\ R(\rho)\dfrac{\varepsilon}{2} & \mathit{if}\ k = j \pm 1 \\ 0 & \mathit{otherwise} \end{cases}$$

where ε is a small, positive, user-defined parameter. Note that this model assumes neighboring strategies are "close strategies." By "close" it is meant that the strategies have something in common, and closeness could be quantized as a measure in the solution space. When no such closeness exists, E is modified as follows:

$$E_j(k, R(\rho)) = \begin{cases} R(\rho)(1-\varepsilon) & \mathit{if}\ k = j \\ R(\rho)\dfrac{\varepsilon}{m-1} & \mathit{otherwise} \end{cases}$$

where m is the number of available strategies. Given the propensities q_{ij}, the probability of selecting strategy j is:

$$p_{ij}(s) = \frac{q_{ij}}{\sum_{k=1}^{m} q_{ik}}$$

which defines a mixed strategy, that is, a finite set of strategies, each one of them with corresponding non-zero probability to be chosen. Roth and Erev studied the performance of this approach on three simple economic games with pure strategy equilibria [38] and eleven simple economic games with mixed-strategy equilibria [39]. They found that their simulated results modeled well several experiments involving human subjects.

The intent of their experiments was to model a psychologically plausible learning approach that could determine the ability of low rationality agents to learn game equilibria, given a variety of playing conditions. They concluded that their model provided an extremely good approximation of the biological learning process in a wide variety of contexts and believe such a simple model motivates additional development of low rationality, cognitive game theory.

In our experimental market, after implementing the original Roth and Erev idea, we observed very low efficiency. The efficiency is specified as the ratio of the profits achievable under competitive equilibrium conditions vs. the actual profits obtained. After observing the time series of the probabilities for each strategy to be chosen, it was clear that when the profits were very low or zero, the algorithm in its original form prevented from learning. Thus, each time a strategy was chosen, the corresponding propensity was updated, and the rest of the propensities were lowered, even if the strategy was not particularly successful one. That effect was especially strong at the very beginning of the auction, when all strategies had equal initial propensities.

In order to compensate for that, the following changes were made in the algorithm. The initial propensities were raised significantly over the ones suggested by the original algorithm. Thus, when a strategy was chosen, and a propensity was assigned, that propensity was lower than the rest. This assures that all strategies would be chosen initially at least once. Next, care had to be taken to fix the situation when profits became low or zero well within the auction run. We assume that in our case there are no neighboring strategies. Initially, according to the original algorithm, he propensity updating equations were:

$$q_{ik} = \begin{cases} (1-\phi)q_{ik} + R(\rho)(1-\varepsilon) & \textit{if } k = j \\ (1-\phi)q_{ik} + R(\rho)\dfrac{\varepsilon}{m-1} & \textit{otherwise} \end{cases}$$

This was changed to

$$q_{ik} = \begin{cases} (1-\phi)q_{ik} + R(\rho)(1-\varepsilon) & \textit{if } k = j \\ (1-\phi + \dfrac{\varepsilon}{m-1})q_{ik} & \textit{otherwise} \end{cases}$$

In this case essentially we have replaced R(r) with q_{ik}, whenever the strategy was not chosen. Thus, if we get zero profits, all propensities shrink, but the updated algorithm assures that the propensity for the strategy delivering zero profits (R(r) = 0) shrinks more than the rest, thus assuring that the other strategies will have higher chance of getting chosen during the next round.

7.4 Results and discussion

Two separate sets of experiments were ran. One of them was with the original Roth-Erev algorithm, and the other one was with the improved one. All parameters were the same, except with the change in the algorithm. There were three sellers and three buyers, each one of them with 30 strategies to choose from. The learning parameters were: e = 0.97, f = 0.50. Discriminatory pricing was done in both cases. For S_1 in the case of the original Roth-Erev S_1 = 9.00, and the improved one S_1 = 10000.00. The marginal price and power to trade are represented in Table 7.1. All ATCs were set to 100 MW. The total profit that could be contracted by each trader under competitive equilibrium was \$400 and the total profit for all traders was \$2400. There were 100 auction runs, each of which was repeated 3 times for collecting statistics. Table 7.2 shows the results of the experiment. All results in the table were rounded off to the decimal point.

As we can see in Table 7.2, the total efficiency of the original Roth-Erev algorithm is very low. Also, the standard deviation of the original algorithm shows how inconsistent the

Table 7.1. Properties of Sellers and Buyers.

	Seller 1	Seller 2	Seller 3
Power to Sell [MWh]	40	40	40
Marginal Cost [\$]	20	20	20
	Buyer 1	Buyer 2	Buyer 3
Power to Buy [MWh]	40	40	40
Marginal Revenue [\$]	40	40	40

Table 7.2. Results of 3 averaged rounds. The profits are in dollars [\$], and the efficiencies are in percentages [%].

	Original		Improved	
	Profit	Efficiency	Profit	Efficiency
Buyer 1	189	47	400	100
Buyer 2	171	43	401	100
Buyer 3	213	53	403	101
Seller 1	162	40	397	99
Seller 2	193	48	398	99
Seller 3	216	54	400	100
Total	1114	215	2400	100

profits are. In the runs with the original algorithm, about half of the outcomes for each trader were zero profits, that is, no match took place, and the trader was not granted even a single contract. On the other hand, the improved Roth-Erev algorithm shows nearly total efficiency for each individual trader, outstanding profit consistency plus excellent total profit efficiency of 100.00% across all traders. There was not even a single case of zero profits, and all available power that could be traded was traded, with no exception for each one of the 3 auction repeats!

From these results we can conclude that the original Roth-Erev reinforced learning algorithm, although being good for mimicking human behavior, is unfit when it has to be implemented to solve actual problems. In our case this was trading of electric power, but that problem could be easily translated to periodic auction trading of any commodity. The improvements made the algorithm robust enough to be computationally implemented to solving industrial problems.

CHAPTER 8. CONCLUSION

Although there does not yet appear to be a standardized final market structure that works for all areas, each market adds up to our experience and helps us make the next market implementation work a little better and more competitively. The goal of the presented computational techniques is to allow for better understanding, to some degree depending on the market implementation, of the present electricity markets and make it easier to model future market implementations.

In this work attention was paid to the study of the effects of different gaming behavior during bidding in agent based power markets. The agents were modeled as autonomous learning entities utilizing different evolutionary programming techniques such as genetic algorithms, reinforced learning and genetic programming. In the course of work two of the models were significantly improved, in particular, the market utilizing genetic algorithm as a learning mechanism, and the market utilizing the reinforced Roth-Erev learning algorithm.

Also it was proven that under certain assumptions the market participants could take advantage of the existing market rules and both gain an unfair advantage over the rest of the participants and at the same time be able to adversely affect the bidding behavior of the other market participants by leveraging certain aspects of the market rules.

Perhaps the major conclusion that could be drawn from the work is that the best way to prevent such predatory behavior is to limit the amount of information available to the market participants. Almost all experiments were conducted under the assumptions that the market participants have almost full information about each other, which is a very close approximation of the real-life scenario, where market participants can obtain almost complete information about their competitors' capabilities from public sources. Which brings up the question to what extend the ICA's such as CAISO, ERCOT and PJM should make such information publicly available and to what extend the market participants should protect their own information.

The question of market information transparency should certainly be investigated and such a direction is worthwhile pursuing.

APPENDIX A. LIST OF ABBREVIATIONS

AI	Artificial Intelligence (widely abused term)
ATC	Available Transmission Capacity
CAISO	California Independent System Operator (Control Area)
EPRI	Electric Power Research Institute
ERCOT	Electric Reliability Council Of Texas (Control Area)
ESCO	Energy Services Company
GA	Genetic Algorithm
GENCO	Generation company
GP	Genetic Programming
GUI	Graphical User Interface
ICA	Independent Contract Administrator
ISO	Independent System Operator
LISP	Interpretive Programming Language used traditionally in AI research
MC	Marginal Cost
MR	Marginal Revenue
PGC	Power Generation Company
PJM	Pennsylvania, New Jersey, Maryland (Control Area)
TRANSCO	Transmission Company

REFERENCES CITED

1. M. T. Bisat, C. W. Richter, G. B. Sheblé, Using Adaptive Agents to Study Bilateral Contracts and Trade Networks, Iowa State University, IA 50011 (1996).

2. A. Wood, B. Wollenberg, Power Generation, Operation and Control, John Wiley & Sons, New York, NY (1984).

3. C. W. Richter, D. Ashlock, G. Sheblé, Effects of Tree Size and State Number on GP-Automata Bidding Strategies, Iowa State University, IA 50011 (1995).

4. D. Cliff and J. Bruten, Zero is Not Enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Market, Hewlett-Packard Laboratories, Bristol, UK (1996).

5. Gerald B. Sheblé, Computational Auction Mechanisms for Restructured Power Industry Operation, Kluwer Academic Publishers, Boston, MA (1999).

6. Gode, D. K. and Sunder, S. Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality, Journal of Political Economy, 101 (1), pp. 119-137 (1993).

7. J. A. Kagel and A. E. Roth, A Handbook of Experimental Economics, Princeton University Press, Princeton, NJ (1995).

8. Alaywan, Ziad, "Facilitating the Congestion Management Market in California," California Independent System Operator, [www.caiso.com, accessed 1998] (1998).

9. Cliff, D. and Bruten, J., "Zero is Not Enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Market," Hewlett-Packard Laboratories, Bristol, UK (1996).

10. Gode, D. K. and Sunder, S. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," Journal of Political Economy, 101 (1), pp. 119-137 (1996).

11. Kagel , J. A. and Roth, A. E., A Handbook of Experimental Economics, Princeton University Press, Princeton, NJ (1995).

12. Richter, C. W., Ashlock, D., Sheblé, G., "Effects of Tree Size and State Number on GP-Automata Bidding Strategies," Proceedings of the Third Annual Genetic Programming Conference, pp. 329-337, Morgan Kaufmann Publishers, San Francisco, CA (1998).

13. Richter, C. W., Sheblé, G. B., and Ashlock, D., "Comprehensive Bidding Strategy with Genetic Programming/Finite State Automata," PE-030-PWSR-0-10-1998 IEEE Transactions on Power Systems (1998).

14. Sheblé, Gerald B., Computational Auction Mechanisms for Restructured Power Industry Operation, Kluwer Academic Publishers, Boston, MA (1999).

15. Wood, A., Wollenberg, B., Power Generation, Operation and Control, John Wiley & Sons, New York, NY (1984).

16. Alaywan, Ziad, "Facilitating the Congestion Management Market in California," California Independent System Operator, [www.caiso.com, accessed 1998] (1998).

17. Cliff, D. and Bruten, J., "Zero is Not Enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Market," Hewlett-Packard Laboratories, Bristol, UK, (1996).

18. Gode, D. K. and Sunder, S. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," Journal of Political Economy, 101 (1), pp. 119-137 (1996).

19. Kagel , J. A. and Roth, A. E., A Handbook of Experimental Economics, Princeton University Press, Princeton, NJ (1995).

20. Richter, C. W., Ashlock, D., Sheblé, G., "Effects of Tree Size and State Number on GP-Automata Bidding Strategies," Proceedings of the Third Annual Genetic Programming Conference, pp. 329-337, Morgan Kaufmann Publishers, San Francisco, CA (1998).

21. Richter, C. W., Sheblé, G. B., and Ashlock, D., "Comprehensive Bidding Strategy with Genetic Programming/Finite State Automata," PE-030-PWSR-0-10-1998 IEEE Transactions on Power Systems (1998).

22. Sheblé, Gerald B., Computational Auction Mechanisms for Restructured Power Industry Operation, Kluwer Academic Publishers, Boston, MA (1999).

23. Wood, A., Wollenberg, B., Power Generation, Operation and Control, John Wiley & Sons, New York, NY (1984).

24. Alaywan, Ziad, "Facilitating the Congestion Management Market in California," California Independent System Operator, [www.caiso.com, accessed 1998] (1998).

25. V. Petrov, C. W. Richter, G. B. Sheblé, Economically Destabilizing Electric Power Markets for Profit, Department of Electrical and Computer Engineering, Iowa State University, Ames, Iowa 50011 (1999).

26. Gode, D. K. and Sunder, S. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," Journal of Political Economy, 101 (1), pp. 119-137 (1996).

27. Kagel , J. A. and Roth, A. E., A Handbook of Experimental Economics, Princeton University Press, Princeton, NJ (1995).

28. Richter, C. W., Ashlock, D., Sheblé, G., "Effects of Tree Size and State Number on GP-Automata Bidding Strategies," Proceedings of the Third Annual Genetic Programming Conference, pp. 329-337, Morgan Kaufmann Publishers, San Francisco, CA (1998).

29. Richter, C. W., Sheblé, G. B., and Ashlock, D., "Comprehensive Bidding Strategy with Genetic Programming/Finite State Automata," PE-030-PWSR-0-10-1998 IEEE Transactions on Power Systems (1998).

30. Sheblé, Gerald B., Computational Auction Mechanisms for Restructured Power Industry Operation, Kluwer Academic Publishers, Boston, MA (1999).

31. Wood, A., Wollenberg, B., Power Generation, Operation and Control, John Wiley & Sons, New York, NY (1984).

32. John R. Koza, Genetic Programming, The MIT Press, Cambridge, MA (1992).

33. D. Zongker, Bill Punch, Lil-GP 1.0 User's Manual, Michigan State University, Ann Arbor, MI (1995).

34. Sheblé, G. B., Computational Auction Mechanisms for Restructured Power Industry Operation, Kluwer Academic Publishers, Boston, MA (1999)

35. Newbery, D. M., Privatization, Restructuring, and Regulation of Network Utilities, The MIT Press, Cambridge, MA (1999).

36. J. Nicolaisen, V. Petrov, M. Smith, L. Tesfatsion, Concentration, Capacity, and Market Power in a Computational Electricity Market with Discriminatory Double-Auction Pricing, CEC2000, San Diego, CA, July 16-21 (2000).

37. D. Lane, A. Kroujiline, V. Petrov, G. Sheblé, Electricity Market Power: Marginal Cost and Relative Capacity Effects, CEC2000, San Diego, CA, July 16-21 (2000).

38. I. Erev and A.E. Roth "On the need for low rationality, cognitive game theory: Reinforcement learning in experimental games with unique, mixed strategy equilibria," mimeo, University of Pittsburgh, Pittsburgh, PA (July 1995).

39. A. E. Roth and I. Erev "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," Games and Economic Behavior, Special Issue: Nobel Symposium, vol. 8, pp. 164-212 (January 1995).

40. D. M. Newbery, Privatization, Restructuring, and Regulation of Network Utilities, The MIT Press, Cambridge, MA (1999).

41. P. Milgrom, "Auctions and Bidding: A Primer," Journal of Economic Perspectives 3, pp. 3-22 (1989).

42. D. Friedman and J. Rust, The Double Auction Market: Institutions, Theories, and Evidence, Volume XIV, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, New York, NY (1993).

43. G. B. Sheblé, Computational Auction Mechanisms for Restructured Power Industry Operation, Kluwer Academic Publishers, Boston, MA (1999).

44. J. Nicolaisen, V. Petrov, M. Smith, and L. Tesfatsion, "Concentration and Capacity Effects on Electricity Market Power," Proceedings of the 2000 Conference on Evolutionary Computation (CEC), pp. 1041-1047 (2000).

45. D. Lane, A. Kroujiline, V. Petrov, G. Sheblé, "Electricity Market Power: Marginal Cost and Relative Capacity Effects," Proceedings of the 2000 Conference on Evolutionary Computation (CEC), pp. 1048-1055 (2000).

ACKNOWLEDGMENTS

Many thanks to my major professor, Dr. Gerald Sheblé, for introducing me to the field of electric power and power markets in general. His help and guidance extended not only over the course of my study but my professional life as well. May his glass be always full.

I would also like to thank my committee members, Dr. Daniel Berleant and Dr. David Hennessy for the effort they put in helping me defend and graduate. They were great as professors, too, and I wish them all the best.

In addition, thanks to my family, and my wife in particular, for their support.

Last, but not the least, I would like to thank the professors and staff of the Electrical and Computer Engineering Department of Iowa State University. In particular I would like to extend my thanks to Pam Myers, for her enduring patience in putting up with me.

Thank you all.