SUPERGAMES IN ELECTRICITY MARKETS: BEYOND THE NASH EQUILIBRIUM CONCEPT

Pedro Correia, Thomas Overbye, Ian Hiskens
Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign
Urbana, USA
correia@students.uiuc.edu, overbye@ece.uiuc.edu, hiskens@ece.uiuc.edu

Abstract - The periodical repetition of market conditions over time leads to the repeated playing of similar games (or supergames) by participants in electricity markets. This repetition of games tempts participants to walk away from the best-response equilibrium strategies provided by Nash solutions. Although Nash solutions make theoretical sense in non-repeated games, their applicability in repeated games is weakened by the fact that these solutions are not, in general, Pareto optimal. This fact paves the way to more complex games where participants are driven by profit maximization in the long run and are, therefore, enticed to explore different solutions in the short term. Knowing that they will meet in similar games in the near future, makes the players adopt implicit cooperative behavior. The willingness to work to a common end may be modeled in automata or agents – which substitute the players – that incorporate collaborative profiles in their stochastical responses to the other automata strategic moves.

Keywords - Electricity markets, supergames, Pareto optimal, automata

1 INTRODUCTION

THE shift from cost-based to price-based trade solutions in re-regulated electricity markets gave rise to market gaming by participants seeking profit maximization. Strategic solutions, namely strategic equilibria, are of major interest in helping to understand trading – and trading outcomes – in the new competitive electricity markets. Multiple equilibria have been shown to exist in a Poolco model wherein some particular assumptions are made [1, 2]. In addition, a method to find these multiple equilibria under this model has been proposed [3]. The model we refer to assumes that the participants or players in the market are rational and attempt to maximize their individual profits by untruthfully revealing their costs in their bid curves. What they play is assumed to be characterized as a static, non-cooperative, continuous-kernel game under complete information, and the solutions prescribed by this game are Nash equilibria [4], either in pure or in mixed strategies.

When the game is non-repeated, the players are left with multiple Nash equilibria to choose from. A Nash equilibrium (or non-cooperative equilibrium) is a solution that is an indididual's best response to strategies actually played by his or her opponents. In other words, it has individual stability.

However, the market conditions have cycles corresponding to the natural and predictable swings of the load over daily, monthly, seasonally, and yearly periods. This makes the players meet again and again under similar scenarios and, more importantly, makes the players learn and collect information from these repeated games. Supergame is the term used to describe an infinite sequence of these ordinary games played repeatedly over time. In [5], supergame equilibria were characterized for games wherein a discount factor was applied to the players' revenues obtained for an infinite time horizon. This discount factor is what allows players to measure the temptation of deviating from the equilibrium solution for a particular stage of the game. If, however, the period of time over which the game will be repeatedly played is unknown, the discount factor loses its applicability. In our model we assume there is no such known parameter. Moreover, we assume that players do not take refuge in an equilibrium solution to avoid being hurt by other players' strategies.

In addition, evidence extracted from experimentation suggests that high revenue solutions can be maintained repeatedly when these do not necessarily correspond to near-equilibrium outcomes [6]. This paradoxical reality underlines the idea that the repetition of games combined with the fact that Nash equilibria are not, in general, Pareto optimal, drives the market participants to higher-revenue solutions. A decision vector is Pareto optimal if there does not exist another decision vector for which some individual objective function may be improved without deteriorating the remaining individual objective functions [7]. In addition to Pareto optimality we can also define weak Pareto optimality. A decision vector is weakly Pareto optimal if there does not exist any other decision vector for which all the individual objective functions are improved. The Pareto optimal set is a subset of the weakly Pareto optimal set. The following definitions and theorems, which may be found in [7, 4] along with the respective proofs, formalize the concepts of Nash equilibrium and Pareto optimality.

Definition 1 *A decision vector $\mathbf{x}^* \in S$ is called a non-cooperative equilibrium or Nash equilibrium if*
$f_i(\mathbf{x}^*) = \sup_{x_i \in S_i} f_i(x_i, \mathbf{x}^*_{-i})$ *for all* $i = 1, \ldots, k$,
where f_i are the individual utility functions, and $\mathbf{x}^*_{-i}$ denotes $\mathbf{x}^*$ with x_i removed.

Definition 2 *A decision vector $\mathbf{x}^* \in S$ is Pareto optimal if there does not exist another decision vector $\mathbf{x} \in S$ such that $f_i(\mathbf{x}) \geq f_i(\mathbf{x}^*)$ for all $i = 1, \ldots, k$ and $f_j(\mathbf{x}) > f_j(\mathbf{x}^*)$ for at least one index j.*

Definition 3 *A decision vector $\mathbf{x}^* \in S$ is weakly Pareto optimal if there does not exist another decision vector $\mathbf{x} \in S$ such that $f_i(\mathbf{x}) > f_i(\mathbf{x}^*)$ for all $i = 1, \ldots, k$.*

Theorem 1 *The solution of the weighting problem*

$$\begin{aligned} &\textit{maximize} \quad \sum_{i=1}^{k} w_i \cdot f_i(\mathbf{x}) \\ &\textit{subject to} \quad \mathbf{x} \in S, \end{aligned}$$

where $w_i \geq 0$ *for all* $i = 1, \ldots, k$ *and* $\sum_{i=1}^{k} w_i = 1$

is weakly Pareto optimal.

Theorem 2 *The solution of the weighting problem is Pareto optimal if the weighting coefficients are positive, that is $w_i > 0$ for all $i = 1, \ldots, k$.*

Because the Pareto optimal solution set and the Nash equilibria set are distinct, the players are caught between choosing higher revenue strategies and best-response strategies. In our point of view, none of them provides the strategic solutions to the game the players have to face repeatedly. Instead, they will put in place strategies that maximize their revenues in the long run, by reacting on each stage of the game according to the moves or readjustments of the other players. This approach brings in notions of "collaborative" and "non-collaborative" that indicate the player's profile or measure the degree by which the players feel attracted to contribute to a higher revenue joint solution. This comprises a paradoxical behavior, for it is clear that rationality dictates the choice of best-response solutions. However, in the long-run these might lead to lower revenue outcomes for all the players involved.

In [8], the authors address the problem of distributed decision-making by distributed agents and characterize the problem of chosing between attractors and Pareto optima as a question of local versus global optimization. They propose ways of changing their decision-attractors both by modifying the way the agents work and by modifying their environment.

As markets become more complex, players' strategies are likely to be represented by computer programs or automata that conduct strategies based on observation of past moves by its opponents. In [9], the authors explore the existence of Nash equilibria in repeated games where players use finite automata to implement their strategies. It is assumed that the players seek profit maximization while minimizing implementation costs. In our model we assume no such cost implementation. Instead, we propose a very simple, finite state machine that is capable of implementing a simple stochastic game.

So, rather than proposing any type of supergame equilibrium, we find a stochastic approach modeling the players' behavior to be more appropriate. Stochastic moves may be seen as a means to mask strategies and, as a result, to lead to non-equilibrium solutions. That is, the model we propose does not provide us with deterministic solutions for the market outcome. It is a stochastic model that indicates how, based on deterministic information provided by the Nash equilibria and the Pareto optimal sets, the supergames in electricity markets should be played by automata when pursuing long term profit maximization.

Furthermore, these automata should not only be able to implement a specified stochastic profile but also capable of updating their profiles – due to market changes as new participants enter – using estimates based on collected market data. Besides, if a collaborative outcome is the desirable solution, these automata should play clear moves that are not mistaken by its opponents and, therefore, do not undermine the market confidence in a collaborative result – which likely has the consequence of driving the market to an equilibrium.

The paper has three more sections. In Section 2 we present the readjustment algorithm used in the repeated game. In addition, the automata models are also explained. Our proposals are exemplified in Section 3 through a simple case. In Section 4 we draw some conclusions about the work presented in this paper.

2 THE READJUSTMENT STRATEGY

Under the Individual Welfare Maximization (IWM) algorithm, the players find the non-cooperative equilibrium solutions for the non-repeated game by independently solving a nested optimization problem for each of the bidding space regions defined by transmission line constraints [1, 2, 3]. The problem assumes the form

$$\begin{aligned} &\max_{\boldsymbol{\alpha}_p} \quad F_p = f_p(\mathbf{P}_p, \mathbf{D}_p, \boldsymbol{\lambda}_p) \ , \quad \forall p \in \mathcal{P} \\ &\text{s.t. } (\mathbf{P}_p, \mathbf{D}_p, \boldsymbol{\lambda}_p) \text{ are determined by} \\ &\qquad \left(\begin{array}{ll} \max\limits_{\mathbf{x},\mathbf{P},\mathbf{D}} & \sum\limits_{i \in \mathcal{D}} B_i(D_i, \alpha_{D,i}) - \sum\limits_{i \in \mathcal{G}} C_i(P_i, \alpha_{P,i}) \\ \text{s.t.} & \mathbf{h}(\mathbf{x}, \mathbf{P}, \mathbf{D}) = \mathbf{0} \\ & \mathbf{g}(\mathbf{x}, \mathbf{P}, \mathbf{D}) \leq \mathbf{0} \end{array} \right) \end{aligned} \tag{1}$$

where $f_p(\cdot)$ denotes the utility function of player p in the set of players, $\mathcal{P}$. Also, $f_p(\cdot)$ is defined as the difference between the sum of benefits minus charges and the sum of payments minus costs from the set of his or her controlled generators and loads. The vectors of generation and load controlled by player p are denoted by $\mathbf{P}_p$ and $\mathbf{D}_p$, respectively. Each player p controls a vector of reported variables that is represented by $\boldsymbol{\alpha}_p$. The nodal prices applied to the generation and load controlled by player p are a byproduct of the OPF and appear as $\boldsymbol{\lambda}_p$. The cost and benefit functions of each generator and load are denoted by C_i and B_i, respectively. $\mathcal{G}$ represents the set of generators and $\mathcal{D}$ represents the set of loads. The cost and benefit function are assumed to be well described by quadratic functions

$$\begin{aligned} C_i(P_i) &= a_{P,i}.P_i^2 + b_{P,i}.P_i + c_{P,i}\ ,\ i \in \mathcal{G} \quad (2)\\ B_i(D_i) &= a_{D,i}.D_i^2 + b_{D,i}.D_i + c_{D,i}\ ,\ i \in \mathcal{D} \quad (3)\end{aligned}$$

where $\boldsymbol{\alpha}_p$, the untruthfully reported parameter or parameters, substitutes one or more of the true cost coefficients in the quadratic function. The equality and inequality constraints are represented by $\mathbf{g}(\cdot)$ and $\mathbf{h}(\cdot)$, respectively, where $\mathbf{P}$ is the vector of all generated power, $\mathbf{D}$ is the vector of all loads, and $\mathbf{x}$ represents the vector of state variables.

A Pareto optimal solution can be found, according to theorems 1 and 2, by substituting the objective function by the weighted sum of all individual utility functions.

$$\begin{aligned} &\max_{\boldsymbol{\alpha}_p} \quad F_p = \sum_{q\in\mathcal{P}} w_{p,q} \cdot f_q(\mathbf{P}_q, \mathbf{D}_q, \boldsymbol{\lambda}_q)\ , \quad \forall p \in \mathcal{P}\\ &\text{s.t. } (\mathbf{P}_q, \mathbf{D}_q, \boldsymbol{\lambda}_q) \text{ are determined by}\\ &\left(\begin{array}{ll} \max\limits_{\mathbf{x},\mathbf{P},\mathbf{D}} & \sum\limits_{i\in\mathcal{D}} B_i(D_i, \alpha_{D,i}) - \sum\limits_{i\in\mathcal{G}} C_i(P_i, \alpha_{P,i}) \\ \text{s.t.} & \mathbf{h}(\mathbf{x},\mathbf{P},\mathbf{D}) = \mathbf{0} \\ & \mathbf{g}(\mathbf{x},\mathbf{P},\mathbf{D}) \le \mathbf{0} \end{array} \right)\\ &\text{where} \quad w_{p,q} \ge 0,\ \forall q \in \mathcal{P} \quad \text{and} \quad \sum_{q\in\mathcal{P}} w_{p,q} = 1. \end{aligned} \tag{4}$$

This method is called the Weighting Method and the weights $w_{p,q}$ measure the relative importance that player p gives to the objective function of player q. The problem as given in (1) may be seen as a special case of problem (4) when $w_{p,q} = 0, \forall q \in \mathcal{P}$ such that $p \neq q$. In a non-cooperative competitive environment it is, however, difficult to agree upon the ranking or relative merit among functions of the multiobjective problem. Two cases appear, nevertheless, as being particularly relevant: the first, when the objective functions are all weighted equally, corresponding to the equivalent situation where there is a unique decision maker; the second, when the profits at the Pareto optimal point are all made equal, which corresponds to a somewhat 'fair' outcome. We will discuss later the consequences of not agreeing upon the same Pareto optimum.

The automata readjust $\boldsymbol{\alpha}_p$ for each stage of the repeated game by choosing the objective function to be either an equilibrium or a Pareto optimal solution and then, using a local readjustment technique, move one step in the desired direction. Newton's method

$$\boldsymbol{\alpha}_p^{(k+1)} = \boldsymbol{\alpha}_p^{(k)} - \epsilon \cdot (\nabla^2_{\boldsymbol{\alpha}_p} F_p)^{-1}\big|_k \cdot \nabla_{\boldsymbol{\alpha}_p} F_p)\big|_k\ , \quad \forall p \in \mathcal{P} \tag{5}$$

is a good choice for the readjustment scheme because it is locally optimal, it is easily implemented by an automaton, and it is a method less prone to misinterpretation by other automata. The ability to interpret the moves by the other automata is of crucial importance since the game model we propose preserves non-cooperation among players.

Given that the automata should not play using short-term rationality, they require a representation of the players' preferences or profiles. In absence of short-term rational – therefore deterministic – choices, the profile representation has to be necessarily stochastic. One simple representation would be a matrix where a player writes his or her willingness to collaborate at the next stage of the game conditioned on his or her perception of the other players' moves during the previous stage. These could have been either non-collaborative (NC) or collaborative (C). Such a table would look like Table 1, where p_1, p_2, p_3, and p_4 represent conditional probabilities.

		Other players' last move	
		C	NC
My last move	C	p_1	p_2
	NC	p_3	p_4

Table 1: Players' profiles

These simple automata that move stochastically in reaction to their opponents' moves may be also represented by finite states machines, as illustrated by Figure 1. In this case we assume that the players limit their choices to this particular type of machine with this set of possible states, which means the players will move either towards the equilibrium solution (to state NC) or towards a particular Pareto optimum (to state C). Were either of the players' strategies more complicated or the set of possible desirable outcomes augmented – multiple Nash equilibria, multiple Pareto optima – the machines would necessarily have to include more finite states.

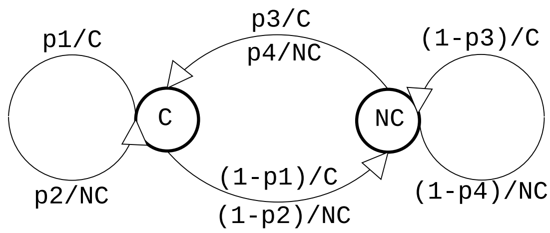

Figure 1: Automaton as a finite state machine

Any representation of the players' profiles that relies on the interpretation of the other players' reactions becomes more and more challenged as the number of players increases. For an Oligopoly with a reduced number of players, it is easy to estimate the individual moves for each stage of the game. Moreover, a good estimate of the profiles may also be constructed. If the number of players is high, clear estimates of the individual moves and profiles may be difficult to obtain and may require the use of complicated metrics. When this is the case, an estimate of the market willingness to pursue collaborative outcomes may be better obtained by looking at aggregated variables such as the market-clearing prices and quantities.

3 EXAMPLE

In order to illustrate the merits of our proposals we ran some experiments on a simple, lossless, unconstrained system, which includes only two competing generators and one independent load. The coefficient values for the linear prices of the generators and value of the load are given, respectively, in Tables 2 and 3.

Generator (i)	$a_{P,i}$ ($\$/MW^2h$)	$b_{P,i}$ ($\$/MWh$)
1	0.01	10.0
2	0.02	10.0

Table 2: Price coefficients

Load (i)	$a_{D,i}$ ($\$/MW^2h$)	$b_{D,i}$ ($\$/MWh$)
1	-0.04	30.0

Table 3: Value coefficients

In these experiments we assume that each of the generators uses an automaton to implement his or her strategies during the supergame that consists of several stages of a game of similar market conditions defined by the unchanging load. We represent the players' profiles by similar transition tables as Table 1 and we assume, given the reduced number of players, that the moves by the opponent automaton are always correctly interpreted. So, once the automata make their moves, they become common knowledge. In addition, and for a matter of simplicity, the automata use only one parameter (the linear coefficient of their reported price curve) to game the system. The profiles' conditional probabilities are uniformly distributed and the automata readjust their gaming parameter using Newton's method (5).

The representative solutions, whose reported parameters are showed on Table 4, for this simple system are: the Nash equilibrium; the Pareto optimal solution where both individual profit functions are equally weighted (Pareto 1 or P1); and the Pareto optimal solution where the profits are made equal for both players (Pareto 2 or P2).

	$\hat{a}_1$ ($\$/MW^2h$)	$\hat{a}_2$ ($\$/MW^2h$)	w_1	w_2
Nash	0.0292	0.0369	–	–
P1	0.0700	0.1400	0.5	0.5
P2	0.0988	0.0928	0.4661	0.5339

Table 4: Reported parameters for specific solutions

The profits attained at each of these three solutions are given in Table 5.

	f_1 (\$/h)	f_2 (\$/h)	$f_1 + f_2$ (\$/h)
Nash Eq.	475.6	331.1	806.7
Pareto 1	769.2	384.6	1153.8
Pareto 2	570.4	570.4	1140.8

Table 5: Profits for specific solutions

Figures 2 and 3 depict, respectively, the Pareto optimal decision vector set and the Pareto optimal profit set when the weighting factors $w_1 = w$ and $w_2 = 1 - w$ vary within the given interval.

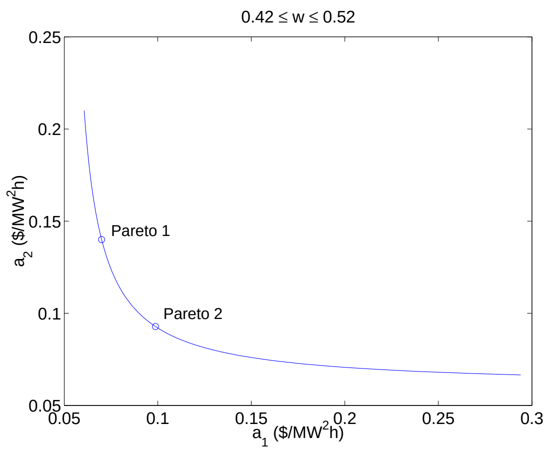

Figure 2: Pareto optimal decision vectors

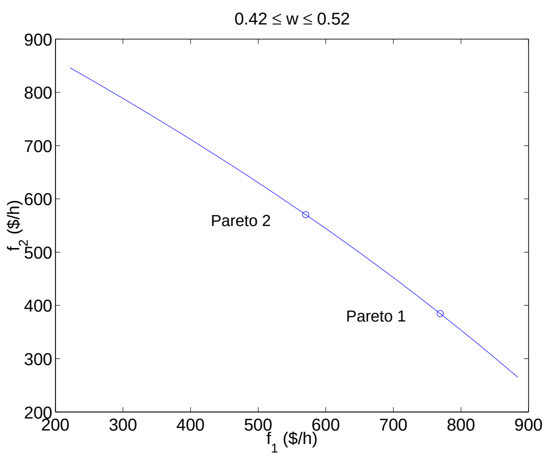

Figure 3: Pareto optimal profits

From the point of view of collaboration among players, it is desirable to maintain the solutions at every stage of the game inside a region where the profits for every player involved in the game are higher than the profits obtained at the Nash equilibrium. Moreover, if the players are seeking the maximization of their revenues, they should never play beyond the curve defined by the Pareto set, e.g., to the right of the curve in Figure 2. If the three conditions – for the two players of our example – are met, the result is the region of mutual benefit as depicted in Figure 4. The points defining this area are, besides the Nash equilibrium, the extreme points $(\hat{a}_1, \hat{a}_2) = (0.0653, 0.1648)$ and $(\hat{a}_1, \hat{a}_2) = (0.1226, 0.0820)$.

Limiting the game to this region brings the game to a new level of rationality, but does not provide the players with a deterministic solution. This is so because of the infinite number of solutions on the Pareto optimal set. In the absence of agreements or side payments, the players are still compelled to keep playing stochastically as a means to keep them away from the non-cooperative equilibrium.

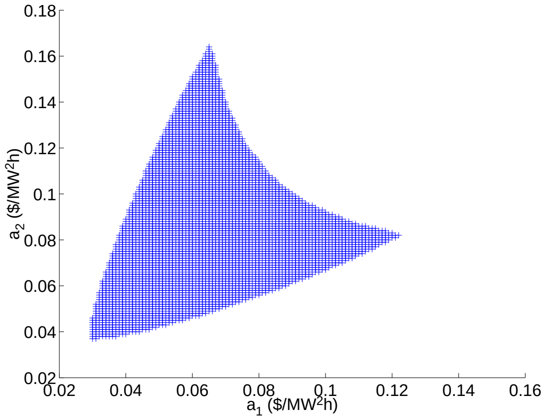

Figure 4: Mutually beneficial region

However, keeping the moves inside the region of mutual benefit creates its own challenges. Each player controls only its own units and therefore can only move along specific coordinates. Without any type of coordination, the players have to divise a strategy if they want to keep the solutions in a desirable region. The only way to overcome the lack of coordination is for the players to assume some principle or non-enforcible rule that dictates the strategy. One of those rules could be the players self imposing equal maximum deviation from the current solution in all coordinates. This strategy would translate into quadrangular gaming areas of variable size for the two-player example. It can be seen in Figure 5 that this strategy effectively keeps the game inside the mutually beneficial region.

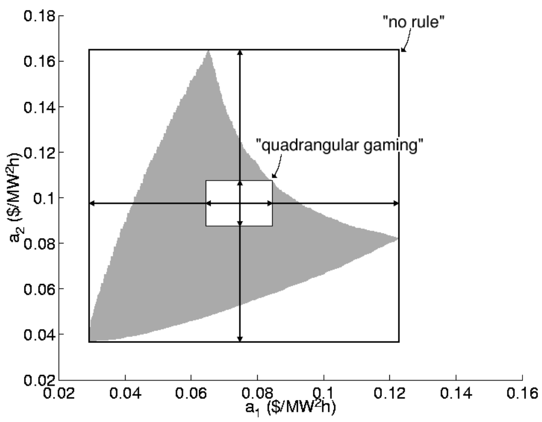

Figure 5: Deviation strategies

Without adopting some kind of rule as the one proposed, the best the players can do is to restrict their moves to the rectangular area enclosing the mutually beneficial region. In the experiments of this section we assume this is the case.

All the experiments that follow were run for 50 periods, each period including 50 stages of the game. At the beginning of each period, the initial values for the players' reported parameters are randomly chosen to lie on the rectangular region defined by the Nash solution and the elected Pareto solution – given by proper choice of weights $w_{p,q}$ in (4) – with a uniform distribution. We picked $\epsilon = 0.1$ for the Newton's method stepsize to avoid large jumps in the values of the reported parameters caused by the steepness of the function when these parameters have small values. From a theoretical standpoint, the choice of the stepsize is not important. What matters is that the players' moves may be correctly interpreted by their opponents.

3.1 Collaborative players

In this first experiment we choose two collaborative players whose profiles are chosen to be the same and equal to $\{p_1 = 0.9; p_2 = 0.5; p_3 = 0.7; p_4 = 0.5\}$. We assume that the players agree on the weakly Pareto optimal solution to be the Pareto 2. The stochastic solutions obtained for these profiles are illustrated on Figure 6.

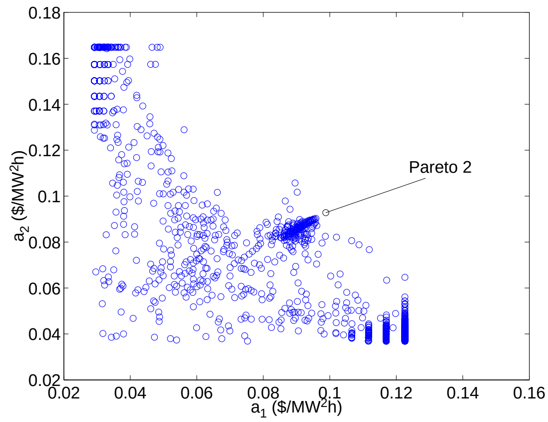

Figure 6: Stochastic solutions for two collaborative players

The straight trajectories that may be observed at the upper-left and lower-right edges on Figure 6 are the result of restricting the moves of the automata to the rectangular area defining the region of mutual benefit. We observe that the players tend to concentrate their moves close to the Pareto 2 point, which is the desirable solution for most of the time. This game grants the players the approximate average profits as showed on row 1 of Table 6 for any particular 50-stage period.

Experiment	$\tilde{f}_1$ $(\$/h)$	$\tilde{f}_2$ $(\$/h)$
1	578	437
2	586	468
3	492	549
4(a)	496	549
4(b)	570	493
4(c)	498	561
Nash	476	331

Table 6: Average profits

These values are obtained by running several periods of the game until the average profits converge. The difference in profits results from the asymmetry of the profit functions. Anyway, both players achieve an average profit that is well above the Nash equilibrium profits.

3.2 Non-collaborative players

This experiment is similar to the previous experiment except for the fact that the players have a non-collaborative profile described by $\{p_1 = 0.5; p_2 = 0.2; p_3 = 0.5; p_4 = 0.2\}$. The results of this experiment are illustrated on Figure 7. The average profits are given on row 2 of Table 6. Surprisingly, this game grants an average profit that is higher for both players than for the previous game. It shows that the solutions of the game do not always concide with what the players antecipate. Instead, the outcomes are dependent on the particular profit functions. In this second experiment the stochastic solutions end up concentrating in more favorable zones, contrary to what was expected, due to the particular shape of the profit functions.

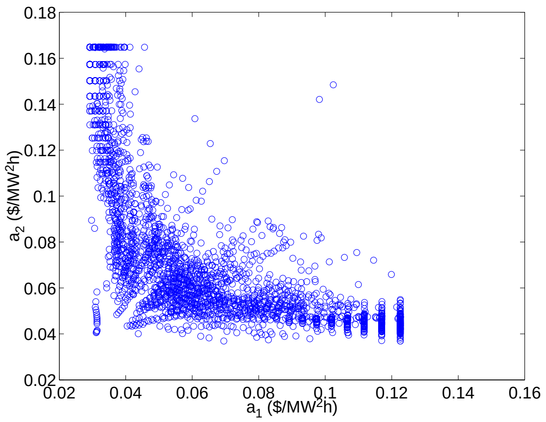

Figure 7: Stochastic solutions for two non-collaborative players

3.3 One collaborative player and one non-collaborative player

In this experiment we are able to evaluate the risk of engaging in collaborative behavior when the other players are not willing to do so. We choose player 1 to be collaborative and player 2 to be non-collaborative. The profile tables have the same probabilities as used previously. Figure 8 shows a period of this stochastic game.

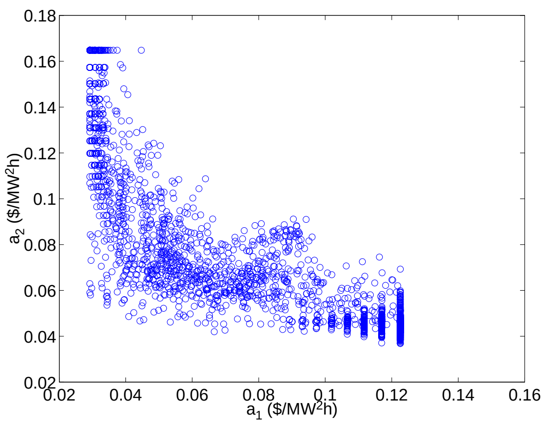

Figure 8: Stochastic solutions for one collaborative player and one non-collaborative player

In comparison to the previous experiments, player 2 is better off and player 1 is effectively "double-crossed", seeing his or her profits coming down (see row 3 of Table 6). However, the average outcome is still advantageous for both players when compared to the Nash equilibrium.

3.4 Disagreement over the Pareto optimal solution

In this experiment we want to evaluate the impact of the disagreement on the desirable Pareto optimal solution. We assume that player 1 chooses the Pareto 1 as his or her desirable collaborative solution and that player 2 chooses the Pareto 2 solution instead – reflected in assumed values of $w_{p,q}$ in (4). The previous three experiments are run as before without any further change. We may observe in Figure 9 a period of the game when both players are collaborative.

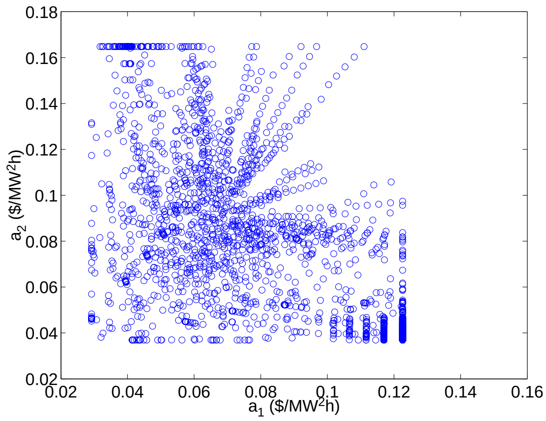

Figure 9: Stochastic solutions for two collaborative players disagreeing on the Pareto optimal solution

The results on rows 4(a), 4(b), and 4(c) of Table 6 show that, again, contrary to what was expected, player 1's profits decrease for experiment 4(a) and 4(b) and increase for experiment 4(c) when compared to the corresponding three previous experiences. Player 2's profits increase in all three experiments. When comparing the results of experiment 4(b) with those of 4(a) we now conclude that for the non-collaborative game, only the revenues of player 2 decrease. For experiment 4(c) we obtain similar results as to those of the first experiment.

3.5 Comments

We adopted this small system in these experiments for illustrative purposes. The major drawback of using this small system is the steepness of the profit functions which lead to some unexpected results. If more players were used in these experiments, the profit functions would be better behaved, also allowing the adoption of less conservative stepsizes in Newton's method.

It becomes clear after these simple experiments that the revenues obtained by the players depend not only upon the choice of one's profile but also upon the opponent's choice of profile. In addition, the choice of the desirable Pareto optimal solution has an impact on everyone's revenues as well. More importantly, independently of their choices, the players maintain in all experiments an average profit that is higher than the equilibrium profit, as may be observed on Table 6.

4 CONCLUSIONS

The work presented here proposes a stochastic approach to the problem of finding trading solutions in a market where trading conditions repeat over time and where players' strategies take into account the existence of both Nash and Pareto optimal solutions. By playing randomly through the use of automata, the players disguise their individual moves to their opponents, since only the profile probabilities are predetermined and not the moves themeselves.

Through the example in this paper we showed that there exists a region between the Pareto optimal set and the Nash solution that always grants the players higher revenues than the Nash solution. If the players restrict their moves to this mutually beneficial region they are indeed increasing their revenues over the game period. One key idea is that only random moves allow them to reap those benefits since deterministic choices of best-response moves drive the players to equilibrium solutions.

The ideas conveyed in the paper may be generalized to any number of players. The two-player example used in the paper was adopted for illustrative purposes. For a n-player scenario the Pareto set would be $n-1$ dimensional and the mutually beneficial region would take n dimensions.

One interesting feature of Pareto optimal solutions that is explored in the paper is that once a system constraint is hit, while the players game the system, at least one of the players involved in the game will see his or her gains decrease or stay constant. Therefore, the Pareto optimal set is bounded by constraints. So, although Nash equilibria may be sustained for different sets of active constraints [3], only one equilibrium may exist for the set of constraints active for the Pareto optimal set. This and other implications of bringing together the Nash equilibria set (multiple equilibria) and the Pareto optimal set have to be addressed in future research.

The paper assumes automata with static profiles. Future research should address the possibility of dynamically changing profiles. Although it is clear that some combinations of profiles lead to better outcomes than other combinations, the calculation of the expected profits with respect to the continuously changing profiles is beyond normal computation power in a reasonable amount of time. What the automata may do is start with a randomly generated set of profiles and, as the game evolves through subsequent stages, select the ones that lead to a better outcome. The dynamic update will likely lead players to implicitly adopt trading solutions that are almost always close to the Pareto optimal set.

ACKNOWLEDGMENTS

This work was supported by the Science and Technology Foundation of the Portuguese Ministry of Science and Technology, under the scholarship PRAXIS XX1/BD/13801/97, and by the U.S. National Science Foundation through grants ECS 00-80279 and DMI 01-09461.

REFERENCES

[1] James D. Weber, "Individual Welfare Maximization in Electricity Markets Including Consumer and Full Transmission System Modeling", Ph.D. dissertation, University of Illinois at Urbana-Champaign, Department of Electrical and Computer Engineering, October 1999.

[2] J. D. Weber and T. J. Overbye, "A Two-level Optimization Problem for Analysis of Market Bidding Strategies", in *Proc. of 1999 IEEE Power Engineering Society Summer Meeting.*, pp. 682-687.

[3] Pedro F. Correia, James D. Weber, Thomas J. Overbye, and Ian A. Hiskens, "Strategic Equilibria in Centralized Electricity Markets", in *Proc. of IEEE Power Tech*. Porto, Portugal: September 10-13th, 2001, Volume I.

[4] Jean-Pierre Aubin, *Optima and Equilibria*. Norwell, MA: 2nd ed., Graduate Texts in Mathematics, Springer, 1998.

[5] James W. Friedman, "A Non-cooperative Equilibrium for Supergames", The Review of Economic Studies, Volume 38, Issue 1 (Jan., 1971), 1-12.

[6] S. Ede, T. Mount, W. Schulze, R. Thomas, and R. Zimmerman, "Experimental Tests of Competitive Markets for Electric Power", in *Proc. of the Thirty-Fourth Annual Hawaii International Conference on System Sciences*, 2001.

[7] Kaisa M. Miettinen, *Nonlinear Multiobjective Optimization*, Norwell, MA: Kluwer Academic Publishers, 1999.

[8] Sarosh Talukdar and Eduardo Camponogara, "Network Control as a Distributed, Dynamic Game", in *Proc. of the Thirty-Fourth Annual Hawaii International Conference on System Sciences*, 2001.

[9] Dilip Abreu, and Ariel Rubinstein, "The Structure of Nash Equilibrium in Repeated Games with Finite Automata', *Econometrica*, Volume 56, Issue 6 (Nov., 1988), 1259-1281.

[10] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. Philadelphia: 2nd ed., SIAM's Classics in Applied Mathematics, SIAM, 1999.