Aiming for Half Gets You to the Top: Winning PowerTAC 2020

Stavros Orfanoudakis[1], Stefanos Kontos[1],
Charilaos Akasiadis[2], and Georgios Chalkiadakis[1]

[1] School of ECE, Technical University of Crete, Greece
[2] Institute of Informatics & Telecommunications, NCSR “Demokritos”, Greece

Abstract. The PowerTAC competition provides a multi-agent simulation platform for electricity markets, in which intelligent agents acting as electricity brokers compete with each other aiming to maximize their profits. Typically, the gains of agents increase as the number of their customers rises, but in parallel, costs also increase as a result of higher transmission fees that need to be paid by the electricity broker. Thus, agents that aim to take over a disproportionately high share of the market, often end up with losses due to being obliged to pay huge transmission capacity fees. In this paper, we present a novel trading strategy that, based on this observation, aims to balance gains against costs; and was utilized by the champion of the PowerTAC-2020 tournament, TUC-TAC. The approach also incorporates a wholesale market strategy that employs Monte Carlo Tree Search to determine TUC-TAC’s best course of action when participating in the market’s double auctions. The strategy is improved by making effective use of a forecasting module that seeks to predict upcoming peaks in demand, since in such intervals incurred costs significantly increase. A post-tournament analysis is also included in this paper, to help draw important lessons regarding the strengths and weaknesses of the various strategies used in the PowerTAC-2020 competition.

Keywords: electricity brokers · trading agents · bidding strategies

1 Introduction

The rise of renewable energy production in the residential market along with the latest popularization of electric vehicles is gradually creating needs for a “smarter” grid. The necessity of this new Grid is indisputable because of the unique features it will be consisted of. In Smart Grid settings, one of the main purposes is to reduce fossil fuel consumption. This is especially important since fossil fuels will be depleted at some point in the future, so alternative energy sources will be eventually required; and since the burning of fossil fuels has a major negative impact to the climate.

Thus, a feature of the new Smart Grid will be an energy market dealing largely in renewable energy, which will consist of a lot of more “prosumer” participants, with most of them being able to buy and sell energy at the same time.

Hence, researchers need tools and platforms that will help them to experiment in novel ways to make this new market viable. The Power Trading Agent Competition (PowerTAC) is a rich simulation platform that can provide researchers with efficient ways to try and test different strategies and approaches before actually deploying them in the future Smart Grid. PowerTAC already has most features a smart electricity grid can possess (e.g., interruptible consumption, electric vehicles, renewable energy) so the simulations can be as realistic as possible. Every year, since 2012, a PowerTAC competition is organized. Agents from research teams from around the globe are pitted against each other, and try to generate the highest profit by harnessing the energy supply and demand of the simulation environment. The agent that won PowerTAC 2020 was TUC-TAC.

Now, contemporary MAS research often builds on solid game-theoretic foundations, since game theory provides a compelling framework for strategic decision making in multi-agent environments [8]. TUC-TAC also gets inspiration from a known theoretical result in order to design the winning strategy of the PowerTAC 2020 competition. More specifically, TUC-TAC's basic goal is to get half of the available retail market share, leaving the rest to the others. By so doing, TUC-TAC expects to always have the highest income, while sharing the fees with the other agents. This basic principle underlying TUC-TAC's strategy has certain analogies to the equilibrium strategy used by the winning agent of the 2010 Lemonade Stand Game tournament [10], as will be explained below. TUC-TAC also employs *Monte Carlo Tree Search (MCTS)* for bidding in the double auctions of the wholesale market, similarly to the approach of Chowdhury et al. [1]. Moreover, TUC-TAC's post-competition strategy is enhanced by a consumption forecasting module (using linear regression and neural networks) to predict demand peaks.

In what follows, we first provide the necessary background for the problem domain; next we present TUC-TAC's architecture and strategy in detail; and then proceed to provide an extensive post-tournament analysis, along with an evaluation of the forecasting module built after the PowerTAC 2020 competition.

2 Background and Related Work

In this section we discuss PowerTAC and some past agent approaches.

2.1 The Power Trading Agent Competition

PowerTAC [3] is a rich competitive economic simulation of future energy markets, featuring several Smart Grid components. With the help of this simulator, researchers are able to better understand the behavior of future customer models as well as experiment with retail and wholesale market decision-making, by creating competitive agents and benchmarking their strategies against each other. In this way, a host of useful information is extracted which can be used by policymakers and industries in order to prepare for the upcoming market changes.

2.2 Past Agent Strategies

In this section, some of the most significant broker-agent strategies will be introduced. Every agent design, in these many years of competition, can be broadly separated into two different, almost distinct, parts. The first part is the Retail Market Module which tries to find the best tariff strategy, i.e. to decide which tariffs to offer to retail customers and to what price; and the second is the wholesale market module, whose main responsibilities are to submit bid and asks in periodic double auctions. Specifically, this module is very important, because the wholesale market is the main place that brokers can buy and sell energy.

Many agents in the past, like COLD [7] used reinforcement learning [9] in order to find the best tariff strategy. Recent agents tried similar strategies. For instance, Mertacor2020 employs Q-Learning techniques in order to maximize the profits from the retail market. The VidyutVanika [2] agent also used a combination of dynamic programming and Q-learning assisted by a Deep Neural Network predictor. However, AgentUDE [5], one of the most successful agents in PowerTAC history, which won the tournaments of 2014, 2017, 2018 and was in the top three brokers in 2016, and 2019, used a much simpler tariff strategy. Specifically, its strategy was based mainly on decision trees and it was being enhanced with some general principles. In addition to that, AgentUDE2017 [6] had a genetic algorithm module to further improve its tariff generation. TUC-TAC 2020 too, uses decision trees in order to find the best tariffs to offer in the retail market; but enhances them with some unique heuristics.

The complexity of the wholesale market actions space, requiring as it does participation in multiple auctions with agent preferences changing dynamically, calls for a very careful design of an agent's respective strategy, in order for it to be profitable. One of the first and most important works in this field was that of TacTex [11] agent in 2013. That team used an MDP price predictor which is the foundation of almost all modern brokers in PowerTAC. Specifically, SPOT [1] agent further improved the previous strategy using MCTS to find the best bids and asks at the best possible times. Another especially efficient wholesale market agent was VidyutVanika [2], which also uses the MDP based price predictor which was firstly implemented by TacTex 2013. Another interesting work, among many, is that of Nativdad et al.(2016) [4], which was using machine learning techniques to reduce the complexity of the wholesale market action space.

3 TUC-TAC's Architecture

TUC-TAC 2020 is an autonomous agent developed to compete in the 2020 Power Trading Agents Competition (PowerTAC-2020). Its main strategy—more specifically, the part of TUC-TAC's strategy that is used in the key for the game retail market—is based on the principle that, acquiring half of the market share will give TUC-TAC half of the total profits, but also only half of the inevitable *transmission capacity fees* (a notion we will explain later) will have to be paid by our agent. Early on in TUC-TAC's development it was realized that greedy strategies would not work in the competitive PowerTAC environment; and the main inspiration for the aforementioned principle came from an interesting equilibrium strategy employed in the context of the "Lemonade Game" competition, and which is briefly presented in Sec. 3.1 below. In order to achieve that, TUC-TAC uses decision trees enhanced with many heuristics and non-heuristics functions that help in the evaluation of the game state. It also employs MCTS for bidding in the double auctions of the wholesale market, adapting it to this setting. In this chapter, we will break down the agent into modules to easier understand how it was designed. Fig. 1 below depicts the main components of the agent; these will be analyzed in turn later in this section.

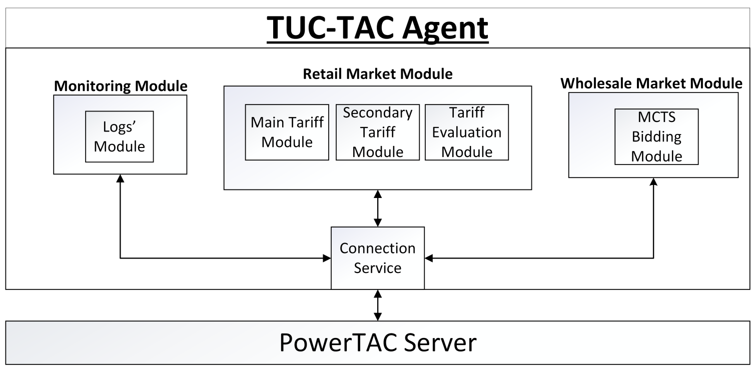

Fig. 1. TUC-TAC's architecture

3.1 An Interesting Equilibrium Strategy for Repeated Multiagent Zero-Sum Game Settings

The *Lemonade Stand Game (LSG)* is a game-theoretic setting with important real-life applications. Specifically, it is a game that can provide important intuitions regarding the choices facing online advertisers, regarding which spot to bid for when participating in real-time online auctions for slots showing up in sponsored search results. In its simplest form, LSG involves N lemonade vendors choosing a location to place their counter at, on the perimeter of a circular island. The utility of each vendor is determined by the distance between her, the neighbor vendors, and the defined space boundaries. In 2010 the first LSG tournament involving artificial intelligent agents took place, sponsored by *Yahoo! Research*, and the strategy of the winning team was shown to be the LSG equilibrium strategy [10]. In short, the strategy demands that one should *always* sit opposite *some* opponent, with the purpose of ensuring that they both maximize their profits via exploiting the third one in this specific iteration; over time, the agent employing this strategy gains most of the profits (across all game iterations), leaving the other two to share the rest of the "pie".

In general, a lesson learned from this equilibrium strategy is that in such settings we should seek to always, at each iteration, claim a *large enough* slice of the pie available, but without being too greedy. This strategy will ensure that any other player will be over time getting lower payoffs than ourselves. In our setting, we are inspired by this equilibrium strategy and develop a strategy for the retail market that seeks to control a high portion of the market share by subscribing a large number of consumers to our services, but also restrain our "greediness" to avoid suffering huge penalties due to transmission capacity fees.

3.2 The Retail Market Module

This component's main responsibility is to publish and revoke tariffs in a way that would be profitable for the agent. Publishing and revoking tariffs alone might sound simple, but there are many aspects of the game that have to be considered before even taking any of these actions. In the following subsections, all these different aspects of the TUC-TAC agent are described in detail.

Preferred Tariff Types A PowerTAC game has a specified amount of different types of power consumers, thus some distinct types of tariffs should be offered. In TUC-TAC's case, strategies for only 4 different tariff types are implemented. These tariffs are about *Consumption*, *Thermal Storage Consumption*, *Solar Production* and *Wind Production* costumers.

To summarize, simple *Consumption* tariffs were selected to be implemented and offered by our agent, because they provided TUC-TAC with an amount of profit that was in expectation significantly higher than that of other tariff types. Also, the two different sustainable energy production tariffs were selected, not for their potential of making a profit, but for their ability to reduce the transmission capacity fees. This will be further explained later. Finally, Thermal Storage consumption tariffs were selected because they provided a considerable stable income. Specifically, the income from these tariffs were several thousand "euros" from customers and the balancing market. Moreover, these tariffs were considered, because it was necessary to prohibit TUC-TAC's competitors from having an advantage by being uncontested in these non-*Consumption* tariffs.

Objective value of a Tariff One of the main problems a PowerTAC agent has to solve, is the evaluation of its opponents' tariffs, with the purpose to offer better ones. (Note that when a broker publishes a new tariff all customers and brokers are notified about its parameters.)

The difficulty of this problem derives from the complexity of the customers' evaluation model itself. A tariff has many parameters to consider while evaluating its objective value. For example some of these parameters are periodic payment, rates, early withdrawal penalties, sign up bonuses and so on.

The average value of rates was calculated using three different methods.[3] The first method tries to find the average value of the rates with the help of the weights which were produced from the time-of-use-technique [5]. In that publication the authors tried to shift the net demand peaks by offering time of use tariffs, so by using their formulas TUC-TAC tried to calculate the objective value of a tariff. The second method calculates the average directly by using the values of the rates without any normalization. The third method calculates the average after normalizing the values of every rate in the tariff. In the end, the second method was selected for the final version of TUC-TAC 2020, because it attracted more customers with the current settings.

Main Tariff Strategy Since the basics of the game and some"peripheral" strategy aspects have been explained, it is time now to describe in detail the strategy which was responsible for TUC-TAC's success. As mentioned earlier, the basic principle that was applied has certain analogies to the equilibrium strategy used by the winning agent of the 2010 LSG tournament [10]. TUC-TAC's strategy is quite similar to that since its basic goal is to get half the available market share leaving the rest to the others. So by doing that, TUC-TAC expects to always have the highest income, while it shares all the fees with the other agents. Fig. 2 below outlines the main TUC-TAC strategy components.

In the beginning, TUC-TAC publishes the initial tariffs and then waits for the assessment timeslots. When it is time for a reassessment of the market state, that agent first checks if any of its current tariffs are exceeding some specific *dynamic* bounds. The tariffs that are out of bounds get revoked, the others remain. Then it checks the number of customers that are subscribed in the total of a tariff type. If the number of subscribed customers is higher than some *MIDDLE-BOUND* value, it instantly revokes its cheapest tariff and creates a new one with the purpose to share the customers with the other brokers. If the amount of the subscribed customers is not higher than *MIDDLE-BOUND*, TUC-TAC checks a *LOWER-BOUND*. The purpose of having a *LOWER-BOUND*, is to remain competitive throughout a game, so, if the amount of the subscribed customers is lower than the *LOWER-BOUND*, it tries to create and publish a tariff that is more attractive than that of its opponents. The aforementioned bounds change dynamically during the games according to weather, time of year, and game state; however, the value of *MIDDLE-BOUND* remains rather close to 50% of the available customers base (more specifically, it ranges between 50% and 62.5%). The whole process is repeated until the game ends.

3.3 The Wholesale Market Module

The second but equally important module of TUC-TAC agent is the Wholesale Market one. Its main responsibilities are to buy and sell energy in the double

[3] It has to be clarified that opponent tariffs with unusual features were considered as "baits" and were not evaluated. Such features could be very high early withdrawal penalties, unusually high periodic payments, or values of rates.

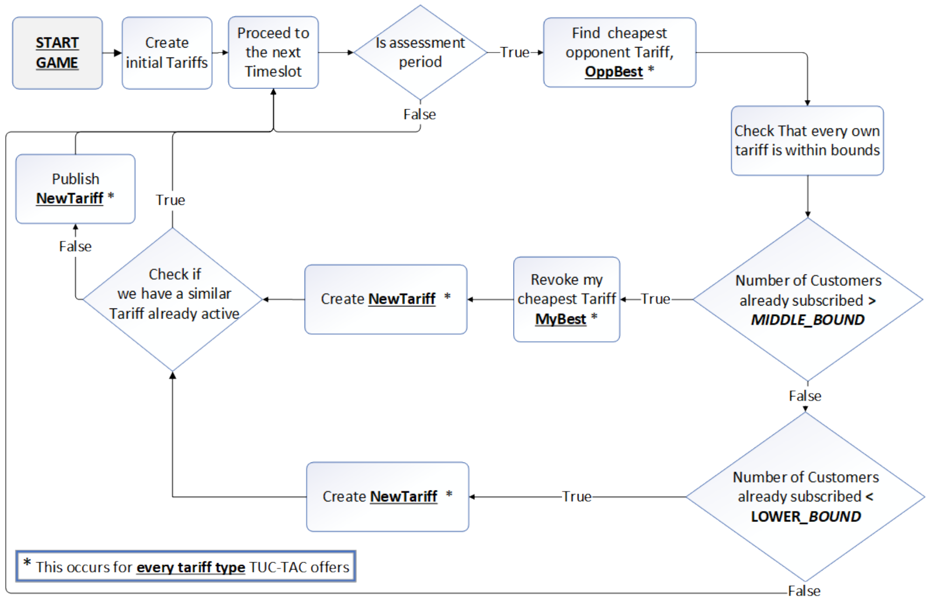

Fig. 2. Main Consumption Tariff strategy flowchart

auctions of the wholesale market. In order to be effective though, it requires finding the best bids so the customers would not have to resort to Balancing Utility to get their energy; in which case TUC-TAC would be charged higher for every single KWh that was not reserved by the agent.

The main algorithm implemented in this module was a variation of the Monte Carlo Tree search method previously developed by Chowdbury et al. [1]. In PowerTAC's case, the double auctions of the wholesale market constitute a complex action-space which requires fast and precise actions in order to be profitable. So, MCTS was selected for its ability to rapidly traverse through huge decision trees and find the best action. Though the concept of this algorithm is indeed suitable for this setting, and can be especially useful, judging by the results of Chowdhury et al [1], the lack of a proper predictor in TUC-TAC's case makes the current wholesale market approach completely naive. For this reason, we are already working towards creating a limit price predictor for PowerTAC 2021.

3.4 Net Demand Predictor Module

At the beginning of each simulated day, TUC-TAC must decide the amount of energy it has to buy for its customers. However, demand changes dynamically and this information is not available in advance. TUC-TAC's net demand predictor estimates the net demand of the customers for the upcoming 24 timeslots, based on the given weather forecast and the past net demand values. We tested two predictors: a classic linear regression method, and a deep learning one.

Dataset Construction In order for TUC-TAC to make a good prediction, a dataset with features that are relevant to the target value is required. The features we used according to the information available to the agent during the competition, are *(i)* the hour of day, *(ii)* the day of the week, *(iii)* the month, *(iv)* the year, *(v)* the temperature, *(vi)* the wind speed, *(vii)* the wind direction, and *(viii)* the cloud coverage. Also, we chose to include lags of the target variable (i.e., net demand value) for the previous time slot, as is common in time-series regression tasks; this was shown to improve our results. Formally, the input vector is represented as

$$x = [h_{t+1}, d_{t+1}, m_{t+1}, y_{t+1}, temp_{t+1}, wSpeed_{t+1}, wDirct_{t+1}, cloudCov_{t+1}, mwh_t]$$

denoting the hour, date, and weather forecasts and demand at timeslot t.

All data used for the training datasets arise from the log files of the 2020 PowerTAC final's games. This data was divided into different datasets, one for each geographical area from which weather data originate from, i.e. Denver, Minneapolis, and Phoenix. Note, however, that the particular area is not known to the agent during the game, and thus, data from every site was selected randomly to form a fourth dataset without geographical distinctions.

Linear Regression (LR) Predictor Our first approach on constructing the predictor was linear regression because it is a classical method for modeling relationships, which in our case is the one between the features in the algorithm's input and the prediction values in the output. To evaluate the method we incorporated the implementation of the scikit learn library in Python.

Deep Learning Regression (DLR) Predictor This approach uses a neural network to discover hidden patterns between the features and the target value. The neural network consists of 2 hidden layers of 24 neurons each, with 10 epochs of training over the training data. The input dataset consists of the features above, and the target variable is net demand in MWh for the next interval.

4 Experiments and Results

This section presents (a) a post-tournament analysis of the PowerTAC 2020 finals; and (b) experiments evaluating the TUC-TAC's demand predictor module.

4.1 PowerTAC 2020 Post-Tournament Analysis

There were 8 agents participating in the 2020 PowerTAC finals. Each agent participated in 40 eight-player games, 105 five-player games, and 63 three-player games. The scoreboard can be seen in Fig. 3 below.[4] . The vertical axis shows

[4] The complete results of PowerTAC 2020 are in https://powertac.org/log_archive/PowerTAC_2020_finals.html. An executable version of the TUC-TAC 2020 agent can be retrieved from https://www.powertac.org/wiki/index.php/TUC_TAC_2020

the score, while the horizontal axis presents the name of the broker in each of the three different scenarios: games with three, five, and eight players. We now proceed with an overview and details of our post-tournament analysis.

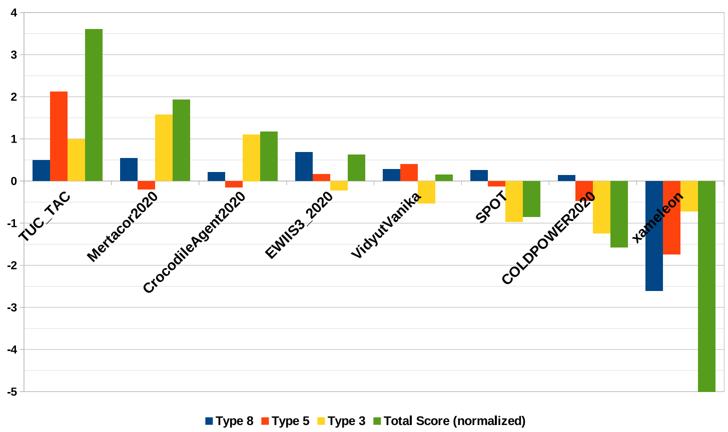

Fig. 3. Final Results of PowerTAC 2020 (Normalized *wrt.* average agent performance)

PowerTAC 2020 post-tournament analysis overview In PowerTAC 2020, most agents used similar technologies, like Markov Decision Processes or Q-Learning, but in the end they appear to have had strategies that corresponded to different "aggression" levels. [5]

TUC-TAC and Mertacor2020 were quite aggressive in the retail market, regardless of the number of players in the game. Specifically, Mertacor used an aggressive decision-making strategy informed by offline reinforcement learning. In the end, TUC-TAC's adherence to its central, though properly adjusted, strategy principle, its faster response times, and the offering of more attractive tariffs, allowed it to have an advantage over Mertacor, and thus TUC-TAC won most of the games it participated in. CrocodileAgent2020 was especially aggressive in 3-player games, but not so much in the more-than-three player categories. Moreover, ColdPower, Spot, VidyutVanika, and EWIIS3_2020 had a "conservative" behavior in the retail market, judging by their lower average scores in most games. It is also important to note here that apart from the retail market strategy failure of Spot and VidyutVanika, these agents were the ones performing best in the wholesale market, apparently having focused on that market, something that will not be further investigated in this work. On the other hand, Xameleon implemented a greedy strategy that did not perform well, probably because of the high fees it had to pay and some flaws in its design.

[5] Some specifics of their strategies were revealed during a post-tournament workshop.

TUC-TAC's stability with respect to transmission capacity fees Transmission capacity fees represent the amount of money a broker should pay for its customers' contribution to demand peaks. This means that when there is a demand peak each broker will have to pay for a portion of the exceeding energy(MWh). In the current PowerTAC competition these fees are the main problem an agent faces when it tries to dominate in the retail market. So after understanding how important those fees are and how these generally affect the game in theory, we present below the results of PowerTAC 2020 for each tier sorted by the total exceeding MWh paid by the brokers as transmission capacity fees. Figure 4 thus demonstrates that TUC-TAC's retail market strategy is the most stable, regardless of the amount of the fees imposed.[6]

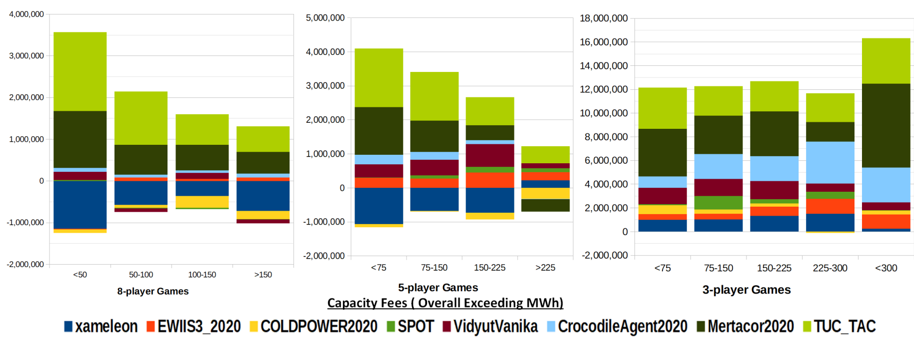

Fig. 4. Categorization by Transmission Capacity Fees: each color bar's area represents the average score for each agent in the games corresponding to different fee levels paid

Thus, we believe one can safely infer that TUC-TAC's retail market strategy achieved its goals. As mentioned earlier in the paper, this strategy was created to mitigate the costs of the transmission capacity fees across more than one agent while TUC-TAC could still take the highest share of the tariff profits. This strategy worked exceptionally well when the majority of agents were in the game, specifically in 8-player and 5-player games. At the same time, this strategy resulted in a very profitable stable performance throughout the 3-player games.

TUC-TAC's Tariff Profits There are 4 different tariff types offered by TUC-TAC, namely Consumption tariffs, Thermal Storage Consumption Tariffs, Solar Production Tariffs, and Wind Production Tariffs. The left plot of Fig. 5 demonstrates the net profits from Consumption and Thermal Storage Consumption tariffs, while the right plot of Fig. 5 shows the losses deriving from the use of Solar Production and Wind Production tariffs.

[6] This and subsequent figures (apart from Fig. 8) exclude the results of "Phoenix games" (see "categorization by balancing fees" below). The extraordinarily high fees paid by the agents in those games would just have added noise to the analysis.

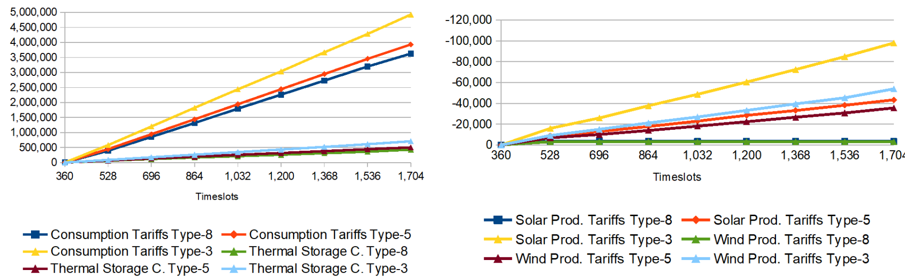

Fig. 5. TUC-TAC's tariffs profits and losses

As observed, the main source of income for TUC-TAC comes from the (electricity) Consumption tariffs, while a smaller but substantial income portion is the result of Thermal Storage Consumption tariffs.

Fig. 5 cannot show the real effect that renewable energy has on PowerTAC. The only thing visible is the amount of money spent in each case to acquire the useful effects of that power type; besides that, it is visible that the losses in each case are very small to be considered harmful. Some of these useful effects of Production Costumers, have to do with the transmission capacity fees. Specifically, when calculating the fees that each agent has to pay, the Balancing Utility of PowerTAC charges each agent according to its contribution to the net demand. So, if an agent has customers that produced energy in that timeslot, that will greatly reduce the transmission capacity fees that the agent will have to pay. In addition to that, it was necessary to compete and increase the tariff prices for production customers, especially in 3-player games, because other agents like EWIIS3_2020 had increased amounts of profits when they could get low-cost energy. Also, an agent can sell or provide immediately to its customers the produced energy, but this technique usually does not generate enough profit.

We believe it is clear how important the Consumption Tariffs are for the profitability of a PowerTAC agent. At the same time, there are other sources of income that are not equally important, but can be considered when deciding the tariff types to offer. We found that the Thermal Storage Consumption tariffs can also be key to making substantial profits in a game (Fig. 5).

Interactions among the main competitors Fig. 6 demonstrates how each of the three best (in 3-player games) agents perform when their main competitors are not part of a game. The most impressive graph is that of CrocodileAgent. As it seems when both TUC-TAC and Mertacor are absent, Crocodile's average score is over 1 million higher, while when only one of the main competitors is absent its average score is almost half a million higher. As such, it is fair to say that Crocodile performs better when the competition is weaker, thus that is the reason it got the second place in type 3 games (see Fig. 3).

In addition, Mertacor appears to have a better performance when TUC-TAC was *not* part of a 3 player game (half a million higher than the total average),

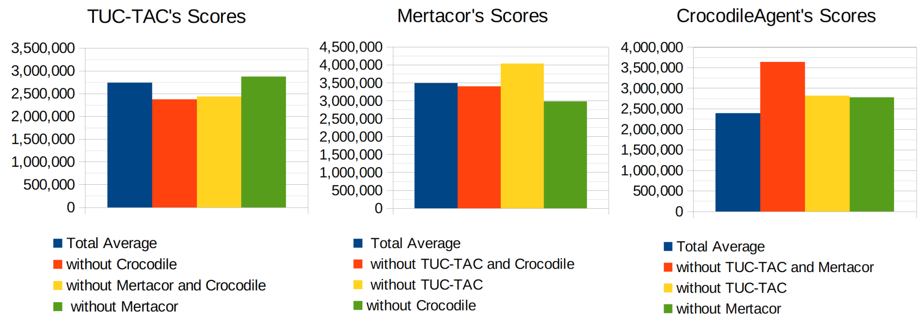

Fig. 6. Impact of TUC-TAC on 3-player games (rewards shown correspond to averages in "regular" games—see Fig. 8).

while its average score dropped by half a million when Crocodile was not part of the game. Still, Mertacor's average score in every case was higher than every other broker in the 3-player games.

However, TUC-TAC's performance was not that much affected by the absence of its main competitors. This of course is a plus in the sense the agent's performance is stable, but at the same time it signifies that TUC-TAC cannot exploit weaker agents that well, unlike CrocodileAgent and Mertacor. Nevertheless, TUC-TAC's stability allowed it to get third place in the 3 player games.

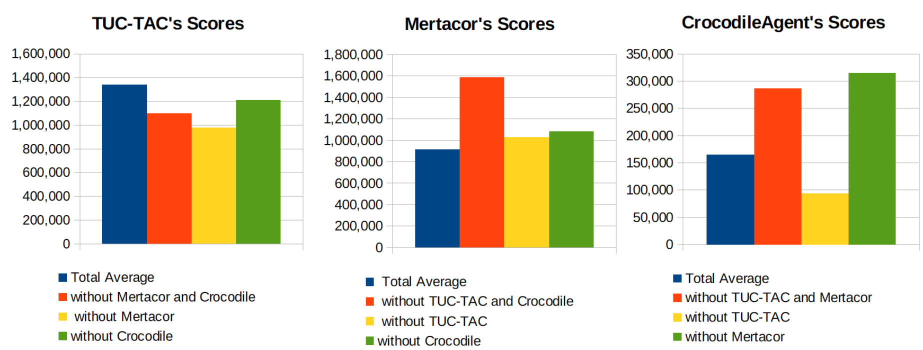

Fig. 7. Impact of TUC-TAC on 5-player games

Fig. 7 depicts the performance of the three best agents in the Regular 5-player games when their main competitors are not part of a game. At first, it is visible that TUC-TAC's performance is quite stable. In addition to that, it is inferred that TUC-TAC has better results when Mertacor is part of a game. This happens because the combination of these two highly competitive agents greatly reduces the market share of the other participating agents, thus resulting in higher profits for TUC-TAC and Mertacor. In Mertacor's case it is visible that the presence of TUC-TAC and CrocodileAgent greatly reduces his average income in a 5-player game. On the other hand, this fact shows that Mertacor is better at exploiting the rest of the agents when there is no direct competition (like when TUC-TAC and Crocodile are participating). Lastly, Crocodile seems to depend on TUC-TAC being present to generate profit in this game type. Its strategy in general is problematic in 5-player games, judging by the fact that its average score in every category (of 5-player games) is very low.

Categorization by balancing fees Balancing fees are the fees that are applied to the agents by the *balancing market* when they fail to procure the required energy. The most common reason a broker might fail to accumulate the required by its customers energy, is very high wholesale market prices.

There were two distinct types of games in this year's finals. The "regular games" and the "special Phoenix games". We term as "special Phoenix games" the games that have extended periods of timeslots with unusually high wholesale market prices. In such situations, agents that have not been careful to buy substantial amounts of energy early on, would have to buy energy in very high prices in the wholesale market. As it was observed in PowerTAC 2020, the leading agents were not prepared for this scenario, thus they could not obtain the required energy during these time periods, and resulting in very high balancing fees for each one of them. This phenomenon usually occurred during the summertime of games located in Phoenix.

	Type 8	**Type 5**	**Type 3**
"Regular" Games	28 / 34	81 /102	47 / 57
"Phoenix" Games	0 / 6	0 / 3	0 / 6
Total	**28 / 40 (70%)**	**81 / 105 (77%)**	**47 / 63 (74%)**

Table 1. Total wins of TUC-TAC

TUC-TAC won the majority of games in every "classic" category, though it had the best overall score in the 5-player games, while being third in the other two game types as we have already seen in Fig. 3. The number of TUC-TAC's wins in each game type, depicting performance in "Phoenix" games also, are shown in Table 1.

By comparing the two graphs in Fig. 8 it is clear how different "Phoenix" games are compared to regular ones, for almost all agents but especially for TUC-TAC and Mertacor. The main reason TUC-TAC was under-performing in these games was a flaw in the design of the wholesale module: as mentioned earlier, this flaw rendered TUC-TAC unable to buy enough energy from the wholesale market to provide to its customers, thus resulting in high penalties for the agent. As is apparent in Fig. 8, almost none of the other participants was prepared for these games as well. However, it seems like these exact same scenarios were very

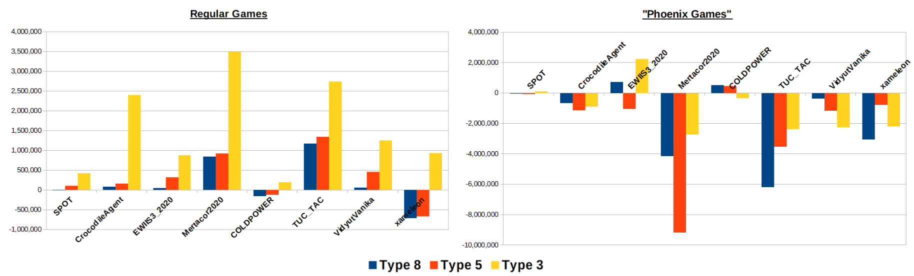

Fig. 8. Differences between Average Scores of “Regular” and “Phoenix” games

profitable for EWIIS3 2020. Even though EWIIS3 was apparently well-prepared for extreme situations, and its performance was generally stable across games, this was not enough to win the tournament. This shows that in order for an agent to win the tournament some aggressive actions should be taken.

4.2 Predictor Evaluation and Impact

In this section we describe experiments that compare our two net demand predictors, with the purpose of identifying the one with the best execution time to error rate ratio. To test the LR and DLR predictors, we perform three types of experiments. In the first one, we take the three geographically divided datasets, fit the 90 % of the datapoints to the linear regression model and predict the rest 10 % of them. In the second one, we combine datasets from 2 out of 3 areas, fit them to the model and predict the target values of the other area. In the third one, we combine all three datasets, shuffle the datapoints, fit 90 % of them to the model and predict the rest 10 % of it. All experiments were performed in a 5-fold validation scheme.

We report that in a preliminary evaluation, we examined the number of lag features that can be used to improve predicting performance. Results (not depicted here due to space restrictions) show that incorporating at least one lag feature leads to lower error rates, but adding more lag features does not help increasing profits significantly.

Now, comparing the two different predictor implementations in terms of execution time is very important, since TUC-TAC is obliged by PowerTAC rules to take decisions within the specific round duration. Results indicate that the execution of the linear regression predictor is about 20 times faster than DLR, which is a significant merit. Specifically, the LR execution time is only 0.69 seconds, as opposed to 655.55 seconds for DLR.

Moreover, the performance of LR in terms of accuracy is very close to that of DLR, as depicted in Fig. 9 which compares the methods in terms of *mean absolute error*, *root mean square error* and *coefficient of determination (R^2)* values. Thus, results indicate that the simpler and faster LR method is the most appropriate to

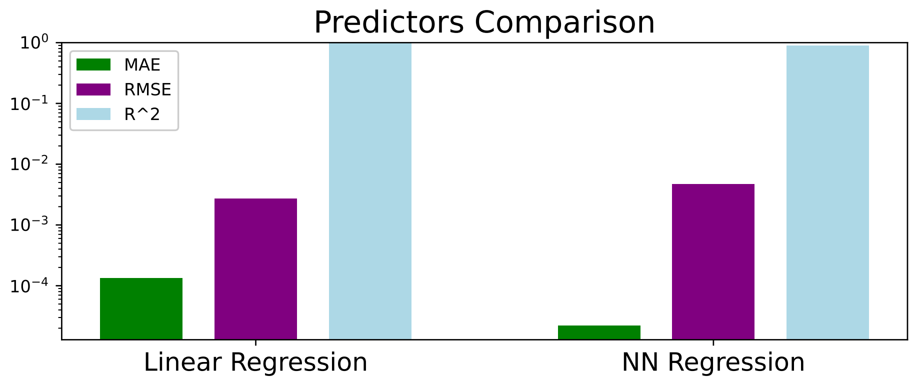

Fig. 9. Error rates comparison between Linear and Deep Learning Regression

incorporate in TUC-TAC. Further experimentation with alternative predictors is in order, and is currently under way.

5 Conclusions and Future Work

The importance of testing the aspects of the emerging Smart Grid before deploying is essential; this is why simulation environments such as PowerTAC are important. This paper presented the strategy of TUC-TAC 2020, the champion of The PowerTAC 2020 competition. A novelty which was arguably key to TUC-TAC's success, is the basic principle underlying its strategy. That principle resembles to some extent the winning, equilibrium strategy of the Lemonade Stand multiagent zero-sum repeated game, in which agents try to acquire approximately half of the market share, leaving the other half to their opponents. However, because of the nature and complexity of PowerTAC, it is next to impossible, in our view, to solve for an actual equilibrium strategy in this domain.

Though TUC-TAC won PowerTAC 2020, there is much room for improvement for the agent. For instance, the wholesale market module has certain drawbacks that need to be overcome. Our first priority is to add a wholesale market limit price predictor. In addition, the MCTS part of the wholesale module needs to be reworked to exploit information accumulated during the 2020 finals. At the same time, it is important to improve the retail market module. Thus, we will look for new ways to improve the agent in the retail market too, having as a first priority to reduce the transmission capacity fees as much as possible. Also, we are looking into ways to support more tariff types to increase profits.

More broadly, we anticipate that the techniques developed for TUC-TAC 2020, can also be applied in a multitude of other multi-agent domains as well. For instance, the generic "equilibrium" strategy for the Retail Market of this competition is conceivably a simple but powerful strategy to use in a host of alternative competitive domains as well.

References

1. Chowdhury, M.M.P., Kiekintveld, C., Tran, S., Yeoh, W.: Bidding in periodic double auctions using heuristics and dynamic monte carlo tree search. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18. pp. 166–172 (2018)
2. Ghosh, S., Subramanian, E., Bhat, S.P., Gujar, S., Paruchuri, P.: Vidyutvanika: A reinforcement learning based broker agent for a power trading competition. Proceedings of the AAAI Conference on Artificial Intelligence **33**(01), 914–921 (2019)
3. Ketter, W., Collins, J., Reddy, P.: Power tac: A competitive economic simulation of the smart grid. Energy Economics **39**, 262 – 270 (2013)
4. Natividad, F., Folk, R.Y., Yeoh, W., Cao, H.: On the use of off-the-shelf machine learning techniques to predict energy demands of power tac consumers. pp. 112–126. Springer International Publishing, Cham (2017)
5. Özdemir, S., Unland, R.: The Strategy and Architecture of a Winner Broker in a Renowned Agent-Based Smart Grid Competition. Web Intelligence p. 165–183 (2017)
6. Özdemir, S., Unland, R.: AgentUDE17: A genetic algorithm to optimize the parameters of an electricity tariff in a smart grid environment. In: Demazeau, Y., An, B., Bajo, J., Fernández-Caballero, A. (eds.) Advances in Practical Applications of Agents, Multi-Agent Systems, and Complexity: The PAAMS Collection. pp. 224–236. Springer International Publishing, Cham (2018)
7. Serrano Cuevas, J., Rodriguez-Gonzalez, A.Y., Munoz de Cote, E.: Fixed-Price Tariff Generation Using Reinforcement Learning, pp. 121–136. Springer International Publishing, Cham (2017)
8. Shoham, Y., Leyton-Brown, K.: Multiagent systems: Algorithmic, game-theoretic, and logical foundations. In: Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations (2008)
9. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. The MIT Press, second edn. (2018)
10. Sykulski, A., Chapman, A., Munoz de Cote, E., Jennings, N.: The winning strategy for the inaugural lemonade stand game tournament. Frontiers in Artificial Intelligence and Applications **215** (2010)
11. Urieli, D., Stone, P.: Tactex'13: A champion adaptive power trading agent. Proceedings of the AAAI Conference on Artificial Intelligence **28**(1) (2014)