Maxon16: A Successful Power TAC Broker

Tobias Urban[1] and Wolfram Conen[2]

[1] Institute for Internet Security
Westphalian University of Applied Sciences
`urban@internet-sicherheit.de`
[2] Westphalian University of Applied Sciences
Department of Computer Science and Communication
`wolfram.conen@w-hs.de`

Abstract. Renewable and sustainable energy production by many small and distributed producers is revolutionizing the energy landscape as we know it. Consumers produce energy, making them to prosumers in the smart grid. The interaction between prosumers and other entities in the grid and the optimal utilization of new smart grid components (electric cars, freezers, solar panels, etc.) are crucial for the success of the smart grid. The Power Trading Agent Competition is an open simulation platform that allows researchers to conduct low risk studies in this new energy market. In this work we present `Maxon16`, an autonomous energy broker and champion of the 2016's Power Trading Agent Competition. We present the strategies the broker used in the final round and evaluate the effectiveness of the strategies by analyzing the tournaments results.

Keywords: Autonomous Agents, Smart Grid, Artificial Intelligence, Multi-Agent System, Machine Learning

1 Introduction

The shift from large energy producers to small and distributed producers is leading to new challenges that need to be addressed. An important part of this energy revolution is the production of energy from renewable and sustainable sources (for example wind or water). The biggest challenges in this context are that on the one hand, most producers cannot produce energy on demand, and on the other hand it is desirable to store the produced energy efficiently. This leads to the transformation of the traditional energy grid to a 'smart grid'. In Europe [13] and other parts of the world [10][4], governments are implementing legislation to promote the extension of smart grids. Smart grids offer opportunities for new business models (e.g. [7]) but also come with some risks [6]. An important feature of the smart grid is that small, local producers (e.g. solar panels on a roof top) can sell their energy using brokers as intermediairies [17]. An interesting research topic for the AI community is to implement such brokers as autonomous agents [14].

The Power Trading Agent Competition [8] (Power TAC) (see Section 2) is a detailed and realistic simulation of a smart grid environment. Within Power TAC, autonomous agents compete against each other trying to maximize their own profits. One of those brokers, `Maxon16` the champion of the 2016's tournament (see Section 3, is introduced in Section 4. We conclude the paper by giving a brief overview of existing work in the Power TAC environment (see Section 6) and by discussing possible future work (see Section 5).
The main contributions of this paper are:

– We give a detailed overview of `Maxon16`, the winning agent of 2016's Power TAC. We analyze the performance of the broker in the final round of the tournament and describe the decision making of the broker.
– We provide a brief overview of the final round of the 2016's Power TAC.

2 Power TAC

The Power TAC is a competitive simulation of a modern energy market. Within the competition, different brokers try to maximize their profits by buying energy on the wholesale market and selling energy on the retail market (by offering tariffs). A typical simulation runs for approximately two simulated months. Each day in the simulation is represented by 24 different timeslots which represent the hours of the day. Each timeslot is five seconds long ('real world time') which results in a total simulation time of almost two hours. Each simulation takes place in a different city and uses its real weather data. Different customer models are used to represent a realistic set of entities that could occur in the future energy market. These customers include devices to store energy (e.g. electric cars and thermal storages), small consumers (e.g. households and small offices), large consumers (e.g. medical centers and office complexes), and local producers (e.g. solar and wind). Each customer supports hourly metering of their consumption and production. Brokers can buy energy at any time on a wholesale market, which is modeled as a simple call market. Divergent from real world energy markets, the wholesale market is implemented for a single region only. On the wholesale market, large and small producers (e.g. wind parks) offer energy to the competing brokers. The market is modeled as day-ahead market which means that brokers can buy or sell energy at most one day ahead (24h) of the actual time when the energy is needed. In an energy grid the demand and supply must be in balance at all times. In Power TAC, this is guaranteed by the Distribution Unit (the owner of the local grid). If a broker's demand and supply are not in balance the Distribution Unit will charge the broker, if his balance has negative impact on the overall balance (e.g. the broker has a negative balance and the overall balance is also negative). A broker is paid if its imbalance has a positive impact on the overall imbalance in the grid (e.g. the broker has a positive imbalance and the overall imbalance is negative). The capacity of modern transmission networks is driven by the peak demands. Due to this fact, it seems plausible to charge brokers for their contribution to these peak demands. In Power TAC, peak demands have to be paid in retrospect for the previous two weeks. Thus,

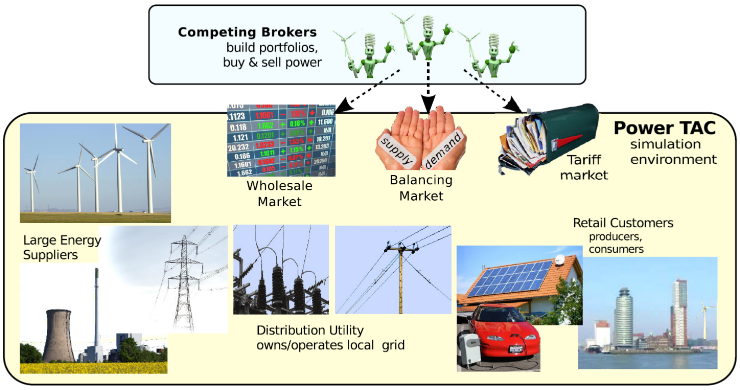

Fig. 1. Main components in Power TAC (borrowed from [8]).

Table 1. Final score of the 2016's Power TAC final round.

	Broker	Game Size 7	Game Size 5	Game Size 3	Norm. Score
1	`Maxon16`	1.42	1.42	0.63	3.47
2	`COLDPower`	0.67	0.68	-0.1	1.26
3	`AgentUDE`	0.54	0.30	-0.1	0.75
4	`SPOT`	0.08	0.185	-0.32	-0.07
5	`Mertacor`	0.08	-0.04	-0.79	-0.75
6	`AgentCU`	-0.96	-0.61	-0.07	-1.64
7	`CrocodileAgent`	-1.82	-2.01	-2.41	-6.24

once every two weeks the broker has to pay for its contribution to the peaks in the last two weeks. On the retail market, brokers can offer different tariffs to which the customers can subscribe to. The main components of the simulation are displayed in Fig. 1.
In general, each broker has three main tasks: (1) Buying and selling energy on the wholesale market, (2) balancing the demand and supply for their customers, and (3) offering tariffs to the customers.

3 Power TAC 2016's final round

The Power TAC 2016 final round was played in three different game sizes: 7 brokers (29 games), 5 brokers (64 games), and 3 brokers (104 games). The final scoreboard of the 2016's final round is shown in Table 3. The scores are the normalized profit (z-scores) each broker made in each game size. In total, 8 brokers competed in the final round. `Maxon16` outperformed every other broker, in every game size, by a distinct margin. The broker `Mertacor` had connection problems and disconnected from most of his games within the first minutes of the simulation.

All statistic provided in this paper are extracted from the official log files of the final round[3].

[3] `http://ts.powertac.org/log_archive/finals_2016_06/`

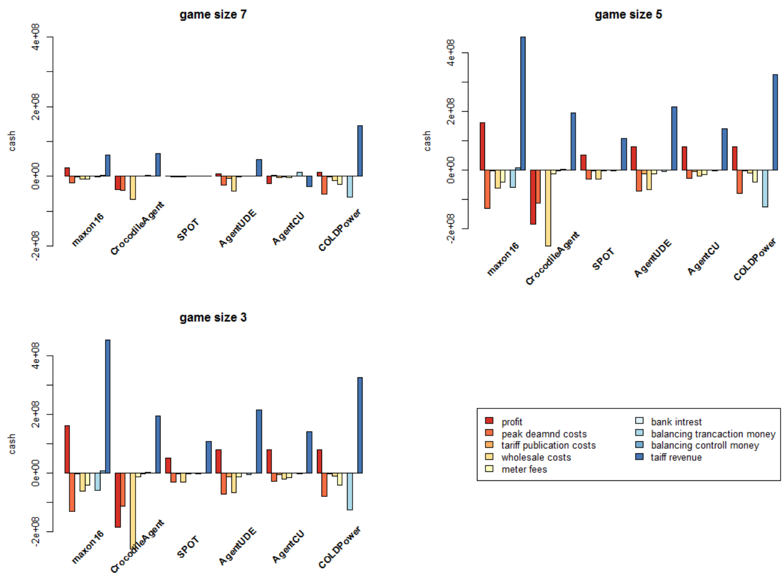

Fig. 2. Overall results of the 2016's Power TAC for different games sizes.

Fig. 2 gives a more detailed overview of the tournament. In the largest games (game size 7) the agent `COLDPower` made the most revenue on the retail market but only made little overall profit in this game size. It is notable that `Maxon16`, `CrocodileAgent`, and `AgentUDE` made almost the same revenue on the retail market but only `Maxon16` was able to make a decent profit from this revenue. The overall costs of `Maxon16`, in this game size, were small compared to the other brokers. In smaller game sizes (game size 3 and 5) `Maxon16` made the most revenue on the tariff market while keeping the costs to produce this higher revenue almost equal to other broker who made less revenue. Thus, `Maxon16` made the most profit in these game sizes.

4 Maxon16

In this section, we give a detailed description of `Maxon16`, the winning agent of the Power TAC 2016 tournament. In the subsections of this section, we focus on the three main tasks the broker has to perform (retail market trading, wholesale market trading, and balancing).

4.1 Retail Market

In the retail market, we use four different consumption tariff types: (1) time-of-use tariffs (TOU), (2) tiered tariffs, (3) flat tariffs, and (4) interruptible consumption tariffs. TOU tariffs have different pricing periods during the day (see

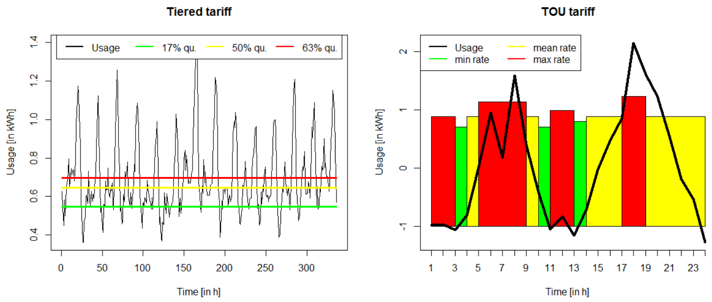

Fig. 3. Tiered (on the left) and TOU tariffs (on the right) used by `Maxon16` in the 2016's Power TAC final round. Tiered tariffs use different usage thresholds (the green, yellow, and red lines) to determine different prices if customers use more than a certain amount of energy. TOU tariffs offer different prices at different times of a day.

Fig. 3). Tiered tariffs consist of different tiers which charge different pricings based on the customer's usage (see Fig. 3). A flat tariff consists of a uniform rate for each used kWh. The last tariff type, an interruptible consumption tariff, is offered to devices that can, to some extent, be controlled by the broker or, on behalf of the broker, by the Distribution Unit.

`Maxon16` uses tiered and TOU tariffs to flatten peak demands by rewarding customers if they alter their usage behavior. TOU tariffs, published by `Maxon16`, offer cheaper rates at times when no peak is expected and tiered tariffs reward customers if they distribute their consumption over the day. Flat tariffs are offered to customer models who cannot, or do not want to, alter their behavior (e.g. a medical center). Controllable devices are used to avoid participation in high demand peaks by down regulating those devices when a demand peak is expected.

`Maxon16` uses the bootstrap usage data (usage data from all customers two weeks prior to the start of the simulation when prices for energy are relatively high and only flat tariffs are offered by a default agent) to identify the customer's usual behavior. We assume that customers will not change their usual behavior during the game if they are not rewarded for doing so, implying that users will change their behavior if they can expect a reward (e.g. a lower price). We design our TOU tariffs not directly by prices that the broker has to pay for the needed energy (like approaches in [2] or [3]), but rather by the times of peak demands that we learned from the bootstrap usage data. Alg. 1 shows the procedure that is used to generate the TOU tariffs. The algorithm uses the aggregated bootstrap usage data of all customer models (B), the baseline price p_b and two scaling factors which scale the prices of the maximum and minimum demands. We use a simple hill climbing algorithm to determine the local maxima and minima of the usage data. Each extremum μ consists of two values μ_v, the usage value of the extremum, and μ_t, the time of day when the extremum occurred. For all the extrema, `Maxon16` computes a price that is used at that time of the day. The price

Algorithm 1: Time-of-use tariff generation

input : The bootstrap data B, the baseline price p_b, and two scaling factors (α_{max}, α_{min})
output: The TOU tariff
localMinima = `HillClimbingAlgorithmMin`(B);
localMaxima = `HillClimbingAlgorithmMax`(B);
globalMax = `max`(B); globalMin = `min`(B);
Tariff t;
forall the *localMaxima* μ **do**
  $p_\mu = p_b \cdot \frac{\mu_v \cdot \alpha_{max}}{globalMax}$; timeSpan $=[\mu_{t-1}, \mu_{t+1}]$;
  t.addRate(p_μ, timeSpan);
forall the *localMinima* μ **do**
  $p_\mu = p_b \cdot \frac{\mu_v \cdot \alpha_{min}}{globalMin}$; timeSpan $= [\mu_{t-1}, \mu_{t+1}]$;
  t.addRate(p_μ, timeSpan);
forall the *remaining times t* **do**
  timeSpan = {t};
  t.addRate(p_b, timeSpan);
`cleanRates`(*t, localMaxima, localMinima*);
return *t*

depends on the actual value of the extrema so that large extrema offer a higher reward or penalty respectively. On one hand, `Maxon16` does this because the hill climbing algorithm may identify insignificant local extrema (e.g. at 2am), and on the other hand, because a higher penalty/reward will increase the motivation of customers to change the behavior. Note that $p < p_b$ for minima and $p_\mu > p_b$ for maxima apply to a rate during the span of one hour both before and after the occurrence of an extremum (μ) (e.g. $\mu_t = 12 \rightarrow$ span= $[11, 13]$). This is done so that customers actually have to shift their usage away from the demand peaks, rather than just move it by one hour. All times of the day that are not adjacent to an extremum use the baseline price p_b. In a last step, we clean up the overlapping rates which occur if there is a local minimum close to a local maximum. In this case, we use the higher rate (local maximum) and shorten the applicability of the cheaper rate (local minimum). The baseline price is computed by using the costs the broker had to pay for each kWh during the last two weeks (168 timeslots):

$$p_b = \text{Production}_{168} + \text{Clearing}_{168} + Imbalance_{168} + \frac{\text{meterFee}}{12} + \text{margin} \quad (1)$$

In the first two weeks of a game, `Maxon16` uses heuristic values (for the imbalance, production, and clearing prices) we computed from the seeding and qualification round prior to the final round of 2016's Power TAC tournament.
We intentionally do not use the costs for peak demands to compute the broker's costs per kWh because they vary drastically and are not controllable by the broker. For example if the customers of other brokers cause a big peak demand our

broker's customers also contribute to this peak. Even if it is a small contribution there will be some peak demand costs for our broker. These costs are covered by the margin we add to our price per kWh.

In order to compute the thresholds for the tiered tariffs, we utilize the usage data of one specific customer group which represents 20,000 different households. In total there are 57,512 customers. We use this customer group because they can adjust their usage and they have a similar usage profile (they use slightly more energy) to another big group of customers (30,000 customers) who can adjust their usage profile as well. We assume again that the customers won't change their behavior (that is captured in the bootstrap data) if they do not expect any reward. Therefore, we never update these values because we expect the updated values to be biased based on the tariffs we offer. The threshold of each tier is computed using different quantiles of the usage data. In total, we use four different tiers: (1) 0-17%, (2) 17-50%, (3) 50-63%, and (4) 63-100%. We use our baseline price p_b for the second tier. Tier one, three and four prices were 7% below, 7% above and 14% above, respectively.

In addition to these two tariff types, we also use a simple flat tariff which uses a rate that is slightly above our computed baseline price. We used this tariff because there are some customers that will not change their behavior and therefore our other tariffs are not attractive to them. The different parameters (α_{min}, α_{max} and the quantile thresholds) which are used to compute the tariffs have been chosen heuristically playing 100 games against the sample broker. We tested how many customers choose the TOU or tiered tariff over the simple flat tariff. We've chosen the parameters in a way that around 70% of the customers subscribe either to the TOU or tiered tariff.

Additionally, we offer tariffs for controllable devices (e.g. the battery of an electric car). The main purpose of this tariff is to offer balancing capacities to the Distribution Unit. For example, the broker can offer to up-regulate a device, resulting in the usage increase of the device, which can be utilized if there is an overall surplus of energy in the grid. If needed, the Distribution Unit can use these offered capacities to keep the grid in balance, and pays the broker or gets paid by the broker respectively. Some controllable devices also allow the broker to economically control their devices (e.g. decreasing the cooling in a freezer). By using this the broker can avoid participation in predicted peak demands since energy usage of the broker's customers is reduced. `Maxon16` only executes economic controls if a demand peak is expected.

Controllable devices split up into two types: (1) storage devices and (2) interruptible consumption. We offer simple flat tariffs to both of these types. For storage devices, we use a price that is only a fraction (33.33%) of our computed baseline price. We did this because we aggressively wanted to get those customers in order to use the balancing control features of these devices even if we lose some money for each kWh these devices use. Note that the usage of these devices is strongly influenced by the issued balancing orders (the device doesn't have to charge its battery if it was charged previously by a balancing order). Interruptible consumption tariffs are, compared to storage customers, less valuable

Algorithm 2: `Maxon16`'s tariff improvement strategy

input : The usage data D, a price baseline p_b, all competing consumption tariffs CT

output: The computed new price baseline p_n

bestCompetingUtility = `min(utilityOfAll(`CT, D`))`;
maxonUtil = ∞;
$\lambda = 2$;
while $maxonUtil \cdot 0.95 \geq bestCompetingUtility$ **do**
  $p_n = \lambda \cdot p_b$;
  maxonUtil = `utility(`p_n, D`)`;
  $\lambda = \lambda - 0.001$;
if $maxonUtil \cdot 0.95 \leq bestCompetingUtility$ **then**
  return p_n

since they cannot be used to store energy. As the name states, the consumption of those devices can only be curtailed for a given timeslot. Thus, we do not fight very aggressively for these customers and offer tariffs with rates that are only slightly under our computed baseline price, around 90% of the computed value. The developed broker also publishes production tariffs, which aim at small and medium sized customers that produce energy (e.g. solar panels on the rooftop of houses). The energy bought on the retail market can either be sold on the wholesale market or used to balance the demand of the broker's other customers. We used our predicted price of the wholesale market (see Section 4.3) as upper price limit in those tariffs.

Tariff improvement The broker's general strategy is to start with relatively expensive initial tariffs (about 20% above the computed baseline) and get cheaper if these tariffs do not perform well.
We measured the success of our tariff set by our market share on the energy market. This share (of the previous timeslot) is computed as follows:

$$\text{share}(t-1) = \frac{E_{\text{bought}}(t-1) - E_{\text{balance}}(t-1) + E_{\text{produced}}(t-1)}{\sum C(t-1) - \sum B(t-1)} \tag{2}$$

with E_{bought} representing the energy bought for a specific timeslot, E_{balance} the energy balance of our broker, E_{produced} the energy produced by the customers, $\sum C$ the total consumption in the grid and $\sum B$ the overall balance of the grid. All these values are taken from the previous timeslot. We compute a threshold t which is based on the number of brokers in the current game. If the current market share is below this threshold, we test if we can publish tariffs that might be more successful than our current tariff set.
We compute the costs of a tariff from the customers' point of view (i.e. in this case the price that the customer has to pay if the customer subscribes to a tariff). In order to compute the costs of a tariff we compute the total price a customer

has to pay for a given tariff and the customers' usage data. Since we compute the costs from the customers' point of view, lower costs are desirable.
The improvement of our tariffs is shown in Alg. 2. First we compute the best (lowest) costs of all competing consumption tariffs. We then continuously lower our price until they are at least 5% lower than the best competing tariff. Finally, we check if our costs are only 5% above the best competing tariff at most (notice the switched relation sign); if so, we return the determined price and publish new tariffs with this base price. We would publish a tariff that is slightly more expensive than the best competing tariff because we hope that at least some customers will subscribe to our newly published tariffs. The broker always publishes three tariffs (one TOU, tiered and flat tariff) based on the computed price whenever `Maxon16` publishes tariffs.
At this stage the broker's `utility()` function only considers flat tariffs. We only consider flat tariffs because the cost computation is straight forward. When customers evaluate TOU and tier tariffs, in the Power TAC environment, they use a risk-adjusted estimate of the expected costs which influences the possibility of them subscribing to TOU and tiered tariffs negatively (from the broker's point of view). Since all customers use different, non-public 'risk adjustment' values we do not take these into consideration. We use a similar approach to improve our production tariffs.
`Maxon16` and `ColdPower` were the only brokers that offered non-flat tariffs (tiered and TOU tariffs). `Maxon16`'s tariffs seem to be attractive to medium sized customers (in terms of energy usage - note the 'meter fee' in Fig. 1). Other brokers (e.g.`COLDPower`) get more customers (in terms of meters) but these customers seem to be less profitable. The medium sized customers seem to be attracted to the possible benefits of shifting their behavior while smaller households (e.g. households) have doubts towards these kind of tariffs. We believe that `Maxon16`'s tariff strategy is the main reason for the broker's success.

4.2 Balance Market

Predicting needed energy is an important part of each brokers tasks because it determines the amount of energy the broker needs for his customers and how much the production customers will produce. At the same time, in Power TAC, this is extremely challenging due to the high number of uncontrollable factors and events that affect the actual usage of the customers. Some of these factors are the weather, new subscriptions or withdraws from tariffs, or the fact that customers change their behavior based on the tariff they are subscribed to. In addition, only very little information about the customers usage is known to the broker: the bootstrap data only consist of two weeks of usage data for each customer. Thus, only two usage values are known to the broker at the beginning of the game for each time of day.
`Maxon16` uses a multiple linear regression model to predict the amount of needed energy. We do not compute the usage for all our customers at the same time but

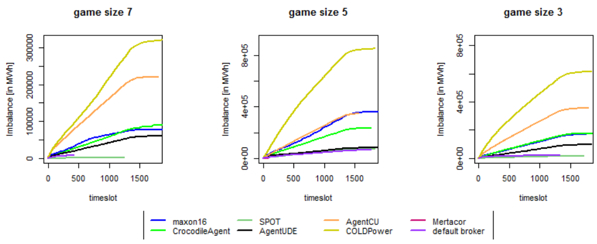

Fig. 4. Overall summed absolute imbalance for the three different game sizes.

for each customer group. We compute the needed energy as follows:

$$u(t+n) := \underbrace{E(t+n)}_{\text{Arith. mean}} + \underbrace{\Lambda(v, v-1, v-2)}_{\text{Trend at time of day}} + \underbrace{\Psi(t, t-1, t-2)}_{\text{Trend for this day}} \tag{3}$$

with $E(t+n)$ representing the arithmetical mean value for a specific time of day. $\Lambda(v, v-1, v-2)$ a linear function that uses difference of the actual usage and the arithmetical mean of the most recent three usage values prior to this time of the day. $\Psi(t, t-1, t-2)$ a linear function that uses the weighted difference between the most recent three predictions and the actual usage of the customer at those times. The parameters of the linear functions (Λ and Ψ) were optimized using usage values from 100 games played. In essence, $u(t+n)$ tries to compute the deviation between the arithmetical mean value and actual value based on the last seen usages.

Even though our approach is not complex it seems to perform quite well in comparison to the forecast of other brokers (if the judgment is based on the observable imbalance stats of the brokers). Fig. 4 shows the absolute imbalances, summed over all games in the final round, for all brokers in each game size. `Maxon16` performs significantly better than two other brokers (`COLDPower` and `AgentCU`). The brokers `SPOT` and default broker' seem to have nearly zero imbalance which indicates that they may have traded very little energy throughout the tournament. `CrocodileAgent` performs almost equally to `Maxon16`. Note that the imbalance is of course not only influenced by the prediction performance, but also by the success of the broker when he tries to buy and/or sell needed or superfluous energy on the wholesale market (and, to an extent neglectable for the 2016 competition, by executing economic controls).

Another aspect on the balancing market are balancing orders. A broker can offer a balancing order to the Distribution Unit allowing the Distribution Unit to use controllable devices contracted by the broker to balance the grid. Fig. 5 shows the use of controllable devices aggregated over all games. Only `Maxon16` and `AgentUDE` offered, successfully, balancing orders. Due to our aggressive strategy (mentioned in Section 4.1) we were able to successfully create a monopole for these devices, only the agent `AgentUDE` got very few controllable devices at the beginning of some games. Note that these values do not affect the broker's imbalance stats; rather, balancing orders are used if there is an imbalance in the

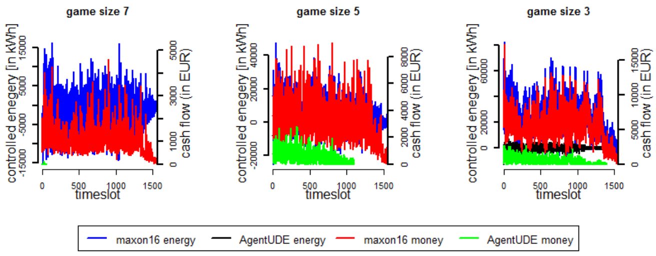

Fig. 5. Balancing orders across the three games sizes of the 2016's final round

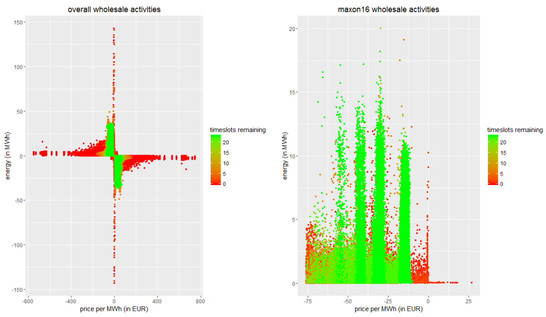

Fig. 6. Cleared bids and asks on the overall wholesale market and for `Maxon16`.

grid. Over the course of all games, `Maxon16` made 15mio EUR revenue by issuing balancing orders.

4.3 Wholesale Market

On the wholesale market, `Maxon16` played a defensive strategy. The idea is to buy the needed energy as early as possible (24 hours ahead). Rather than submitting one huge order, `Maxon16` splits the order into multiple smaller orders at different price levels. The intention is to buy the energy at a low price, if the market allows it, but to still get a significant amount of the needed energy if the market prices, for the given time slot, are high. If the broker does not get the needed energy the price limit is increased, based on the last clearing prices, up to a level until it is cheaper to run into balance. The broker only sells energy if he is absolutely confident that he doesn't need it. This case occurs basically never.
In retrospective, it was wise to play a defensive strategy on the wholesale market because at least one broker tried to monopolize the market by buying huge amounts of energy. Because of that monopole some brokers had to buy energy at very high prices. The left hand graph of Fig. 6 shows the activities on the wholesale market over the course of the tournament. Note that the prices of the traded energy increases while the amounts of traded energy decrease. The right hand graph of Fig. 6 displays the activities of `Maxon16` on the wholesale market.

The defensive strategy of splitting the needed energy worked quite well. Most of the times the broker was able to buy the needed energy early at moderate prices.
This defensive strategy would probably work well against most strategies on the wholesale market (in terms of getting most of the energy `Maxon16` needs at a reasonable price). `Maxon16` does not make any money on the wholesale market and does not perform substantially better than other brokers (in terms of money paid for energy). Since `Maxon16` buys the needed energy early, in multiple chunks and at different price levels, the opponents strategy needs to be very sophisticated to block `Maxon16` from getting the needed energy (e.g. by buying huge amounts at high prices and reselling it to `Maxon16` at higher prices that `Maxon16` is still willing to pay). Note that if a broker buys a huge amount of energy, the expected balance costs will be low because there is probably a surplus of energy in the grid. In that case the Distribution Unit charges brokers that contribute to this surplus and pays brokers that do not. If there is no surplus of energy, `Maxon16` would most likely get most of the needed energy because the market usually provided energy at a price that is reasonable enough.

5 Future Work & Conclusion

`Maxon16` is an efficient and successful energy broker, nevertheless further improvement is needed prior to the next Power TAC. The highest priority is to develop an adaptive wholesale strategy that can dynamically adjust to the market situation. Another important issue is the usage of more controllable devices in combination with the prediction of needed energy. We have shown that the usage of controllable devices has a significant impact in the grid. However, our current approach is relatively simple and can be further tuned in various ways. One approach could be to predict the expected imbalance for a timeslot, issue expensive balance orders and buy more (or less) energy based on the prediction. In general, our energy forecasting mechanism can be improved greatly. But given the limited amount of data at the beginning of the game and the many factors that lie out of the brokers hands, it is difficult to find a more accurate and efficient way to predict the needed energy. A critical part is to take weather conditions into account since the weather has a significant influence on the actual usage and production of customers.
In this paper, we presented `Maxon16`, the champion of the Power TAC 2016 tournament. The broker takes an empirical, practical and largely heuristic approach to the problems brokers will face in future energy markets. We explained in detail some important aspects of the decision making process of the broker and gave an overview of results achieved in the final round.
We also gave some practical suggestions that can be used by new and existing teams to improve their brokers. We published the binaries of our broker in the Power TAC broker repository[4], allowing other researchers to use it for their own research in the Power TAC environment.

[4] `http://www.PowerTAC.org/wiki/index.php/Form:Broker`

6 Related Work

Most Power TAC related papers describe brokers that participate in the competition. The champion of the 2013's tournament and the runner-up of the 2015's tournament, `TacTex`, is describe in [16] [15]. `TacTex` is a very sophisticated broker that formalized all tasks that an autonomous broker faces in Power TAC utilizing high-dimensional MDPs.

In [3] D. Urieli et al. present an approach how to use TOU tariffs in competitive energy markets. The approach differs from our approach since it is driven by the costs the broker has at specific times of a day rather than by the occurrence of peak demands. However, D. Urieli et al. approach reduced peak-demands in their experimental setting by 15%.

`AgentUDE` [12], the champion of 14's competition, uses an empirical approach to predict prices on the wholesale market by predicting a base price from past data and differentiated these predictions to the final price. On the retail market `AgentUDE` publishes tariffs with cheap rates, a sign up bonus (the user gets paid for subscribing to the tariff), and a high early withdraw payments (the customer has to pay a fee if he withdraws early from his contract). The main idea behind this strategy is that customers will leave the tariff if other brokers offer new cheaper tariffs, and by doing so `AgentUDE` can collect the high withdraw payments.

`COLDPower` [2] also uses MDPs and Q-Learning approaches in order to create tariffs and to determine the prices on the wholesale market. This is the only other broker, besides `Maxon16`, that used TOU tariffs in the 2016 competition. The rates of the TOU tariffs are based on the user's average consumption profile.

Another successful broker of past tournaments, `cwiBroker` [9], tried to establish equilibria (price and energy) on the wholesale market. The broker tries to buy all needed energy (for all customers) and then resell it at higher price. On the retail market `cwiBroker` used a strategy that is quite similar to Tit-For-Tat.

In [11] decision trees are compared, in relation to the performance of the agent `SPOT`, to MDPs in combination with Q-Learning to predict the prices on the wholesale market.

The problem of predicting the energy consumption of customers in the Power TAC environment is discussed by F. Natividad et al. in [5] using off-the-shelf machine learning techniques. They achieved 'reasonably accurate' predictions for the consumers. In their work F. Natividad et al. used k-means, k-medoids, and DBSCAN algorithms to predict the demands.

A detailed analysis of the Power TAC 2014 competition can be found in [1].

Acknowledgments We would like to thank the organizers and developers of Power TAC for their efforts. We are thankful for Christopher Huber's contributions to our broker. Last but not least we would like to thank all reviewers for their insightful comments on the paper, as these comments led us to an improvement of the work.

References

1. Babic, J., Podobnik, V.: An Analysis of Power Trading Agent Competition 2014. In: Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets (2014)
2. Cuevas, J.S., Gonzalez, A.Y.R., Alonso, M.P., de Cote, E.M., Sucar, L.E.: Distributed energy procurement and management in smart environments. In: 2015 IEEE First International Smart Cities Conference (ISC2). pp. 1–6 (2015)
3. Daniel Urieli, Peter Stone: Autonomous Electricity Trading Using Time-of-Use Tariffs in a Competitive Market. In: Proceedings of the 30th Conference on Artificial Intelligence (2016)
4. Federation, R.: Energy Strategy of Russia for the Period up to 2030 (2010)
5. Francisco Natividad, Russell Y. Folk, William Yeoh, Huiping Cao: On the Use of Off-the-Shelf Machine Learning Techniques to Predict Energy Demands of Power TAC Consumers. In: Proceedings of the International Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA) (2016)
6. Joskow, P.L.: Lessons Learned from the Electricity Market Liberalization. In: The Energy Journal, Special Issue on the Future of Electricity. pp. 9–42 (2008)
7. Karnouskos, S., Silva, P.G.D., Ilic, D.: Energy services for the smart grid city. In: Proceedings of the IEEE International Conference on Digital Ecosystems and Technologies. pp. 1–6 (2012), `http://ieeexplore.ieee.org/iel5/6222390/6227884/06227925.pdf?arnumber=6227925`
8. Ketter, W., Collins, J., Weerdt, M.d.: The 2016 Power Trading Agent Competition. SSRN Electronic Journal (2016)
9. Liefers, B., Hoogland, J., La Poutré, H.: A Successful Broker Agent for Power TAC. In: Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets (2014)
10. Ministry of Economic Development: New Zealand energy strategy to 2050: Powering our future : towards a sustainable low emissions energy system (2007)
11. Moinul Morshed Porag Chowdhury, Russell Y. Folk, Ferdinando Fioretto, Christopher Kiekintveld, William Yeoh: Investigation on Learning Strategies for the SPOT Broker in Power TAC. In: Proceedings of the International Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA) (2016)
12. Ozdemir, S., Unland, R.: A Winner Agent in a Smart Grid Simulation Platform. In: 2015 IEEE / WIC / ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT). pp. 206–213 (2015)
13. Publications Office of the European Union: Energy roadmap 2050 (2012)
14. Ramchurn, S.D., Vytelingum, P., Rogers, A., Jennings, N.R.: Putting the 'smarts' into the smart grid. Communications of the ACM 55(4), 86 (2012)
15. Urieli, D., Stone, P.: TacTex'13: A Champion Adaptive Power Trading Agent. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems (2014)
16. Urieli, D., Stone, P.: An MDP-Based Winning Approach to Autonomous Power Trading: Formalization and Empirical Analysis. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (2016)
17. Wissner, M.: The Smart Grid – A saucerful of secrets? Applied Energy 88(7), 2509–2518 (2011)