VidyutVanika21: An Autonomous Intelligent Broker for Smart-grids

Sanjay Chandlekar[1*], **Bala Suraj Pedasingu**[2*], **Easwar Subramanian**[2], **Sanjay Bhat**[2], **Praveen Paruchuri**[1] and **Sujit Gujar**[1]
[1]International Institute of Information Technology (IIIT), Hyderabad, India
[2]TCS Innovation Labs, Hyderabad, India
sanjay.chandlekar@research.iiit.ac.in, {sujit.gujar, praveen.p}@iiit.ac.in,
{balasuraj.p, easwar.subramanian, sanjay.bhat}@tcs.com

Abstract

An autonomous broker that liaises between retail customers and power-generating companies (GenCos) is essential for the smart grid ecosystem. The efficiency brought in by such brokers to the smart grid setup can be studied through a well-developed simulation environment. In this paper, we describe the design of one such intelligent energy broker called *VidyutVanika21* (VV21) and analyze its performance using a simulation platform called PowerTAC (Power Trading Agent Competition). Specifically, we discuss the retail (VV21–RM) and wholesale market (VV21–WM) modules of VV21 that help the broker achieve high net profits in a competitive setup. Supported by game-theoretic analysis, the VV21–RM designs tariff contracts that a) maintain a balanced portfolio of different types of customers; b) sustain an appropriate level of market share, and c) introduce surcharges on customers to reduce energy usage during peak demand times. The VV21–WM aims to reduce the cost of procurement by following the supply curve of the GenCo to identify its lowest ask for each auction which is then used to generate suitable bids. We further demonstrate the efficacy of the retail and wholesale strategies of VV21 in PowerTAC 2021 finals and through several controlled experiments.

1 Introduction

A *smart grid* is an electricity network used to supply energy to consumers via two-way digital communication. It allows for monitoring, analysis, control and communication between participants to improve efficiency, transparency and reliability of the energy supply chain [Techopedia.com, 2021]. A smart grid environment consists of a retail market involving various types of customers, a wholesale market involving power generating companies, distribution utility and energy *brokers* who interact between retail and wholesale markets. There are potential opportunities to improve smart grid operations, especially when renewable energy generation is getting bigger by the day, to efficiently manage fluctuating supply-demand scenarios and grid imbalance situations.

To investigate potential solutions to improve smart grids, the Power Trading Agent Competition (PowerTAC) [Ketter *et al.*, 2020] provides an efficient simulation of real-world smart grids, with energy brokers playing a pivotal role. In PowerTAC, a broker serves its customer base in the retail market by purchasing the required energy from the wholesale market. The broker handles supply-demand imbalances via balancing market. The primary goal of the broker is to maximize its profit by operating in these three markets. Hence, PowerTAC organizes an annual tournament where multiple teams design autonomous broker agents to compete against each other across several games in multiple player configurations. Brokers get revenue from subscriptions in the retail market and incur costs for procuring energy in the wholesale market. Often, brokers also pay penalties for contribution to grid imbalance and peak time usages. In PowerTAC, peak time usage charges are relatively high and are called as *capacity transaction charges*. The brokers aim to develop retail strategies to maximize revenue and lessen penalties; moreover, to design wholesale strategies to reduce procurement costs.

Brokers in previous PowerTAC tournaments used gradient-based optimization strategies [Urieli and Stone, 2016], MDP based strategies [Ghosh *et al.*, 2020], genetic algorithm based approaches [Özdemir and Unland, 2018] to offer tariffs in the retail market. The seemingly optimal class of strategies of monopolizing the retail market have been found to suffer from high capacity transactions charges. To overcome this, broker agent TUC_TAC proposed an interesting strategy aimed at acquiring half the retail market share [Orfanoudakis *et al.*, 2021] in PowerTAC 2020 tournament. To the best of our knowledge, most retail strategies proposed thus far have been generic and are not effectively specialized for different player configurations and therefore fail to maintain performance across different player configurations. Thus, a careful design of the broker's retail strategy is needed.

Similarly, in the wholesale market, past brokers used MDP-based approaches [Urieli and Stone, 2014; Ghosh *et al.*, 2020] or heuristics [Chowdhury *et al.*, 2018] to minimize energy procurement costs without paying much attention to the bidding behavior of sellers. This often results in higher costs to brokers, and sub-optimal bidding patterns by brokers will push energy purchases to the balancing market where the

*Authors contributed equally to this work costs are typically high. Hence, a bidding strategy that considers the bidding behavior of sellers is warranted.

Our Approach and Contributions: To address the shortcomings discussed above, we present our broker *VidyutVanika21* (VV21), which is the winner of the PowerTAC 2021 tournament. The retail strategy of VV21 selects the market share of its retail portfolio by ingeniously following a game-theoretic analysis specialized to player configuration. In addition, tariffs contain surcharges for peak hours, resulting in customers shifting load to off-peak hours apart from providing a cushion to handle peak demand charges. The wholesale strategy aims to place bids by identifying the lowest ask of the largest seller by following its supply curve. This ask price is used as an upper bound to procure as much power from other sellers present in the wholesale market at cheaper prices. Further, to reduce imbalance costs, VV21 attempts to procure all the required energy from the wholesale market, thus enabling effective energy procurement. To summarise:

- We present the design of an autonomous energy broker, VV21, the winner of the PowerTAC 2021 tournament.
- We describe the design of VV21–RM, the retail module of VV21, which uses game-theoretic analysis to offer optimal tariffs with the aim of maximizing net revenue.
- We detail the working of VV21–WM, the wholesale module of VV21, that recognizes the bidding pattern of sellers and uses it for procuring power at cheaper prices.
- We analyze the performance of VV21 in the PowerTAC 2021 tournament and showcase its efficacy through several controlled offline experiments.

2 Related Work

Research work on smart grids has focused on various sub-problems in energy management with approaches backed by theory and methods driven by data analysis. For example, the work of [Chung *et al.*, 2020] considers the problem of load scheduling by modeling it as a stochastic game and uses deep RL to find the Nash equilibrium (NE) of the game. In [Jain and Gujar, 2020], the authors proposed bandit-based mechanisms to achieve demand response in the smart grids. However, such strategies are impractical in PowerTAC as brokers need to respond in real-time.

Among the strategies deployed in past tournaments, MDP-based strategies are popular and have been used both in retail and wholesale market operations. For example, Reddy & Veloso, COLD Energy, and VidyutVanika18 modeled the decision process in the retail market as an MDP to generate tariff contracts [Reddy and Veloso, 2011; Serrano *et al.*, 2017; Ghosh *et al.*, 2019]. Whereas TacTex'13 used a gradient-based optimization method to generate tariffs and AgentUDE17 implemented a genetic algorithm-based tariff strategy [Urieli and Stone, 2016; Özdemir and Unland, 2018]. In the wholesale market, TacTex'13, VidyutVanika18, and AstonTAC employed an MDP-based strategy [Urieli and Stone, 2014; Ghosh *et al.*, 2020; Kuate *et al.*, 2013]. AgentUDE17 employed adaptive Q-learning, while SPOT and TUC_TAC adopted MCTS to design their bidding strategies [Chowdhury *et al.*, 2018; Orfanoudakis *et al.*, 2021]. Moreover, a DDPG-based bidding strategy is employed by [Chandlekar *et al.*, 2022], which works in continuous state and action spaces. However, none of the previous works present an equilibrium-based strategy for the retail market; furthermore, none of the past brokers have attempted to learn seller's behavior in the wholesale market. Our novelty lies in designing a game-theory-inspired strategy in the retail market and adopting a bidding strategy based on the supply curve of the prominent seller in the wholesale market.

3 The PowerTAC Environment

PowerTAC is a simulation platform that replicates crucial elements of a smart-grid ecosystem like retail, wholesale, and balancing markets along with a distribution utility (DU). The retail market contains state-of-the-art customer models that replicate a smart grid's real-world users like consumers, producers, and storage users. Consumers vary from households, offices, and villages to hospitals. Producers are modeled to generate energy from renewable sources like solar or wind. Storage customers use batteries or electric vehicles to store energy which can be supplied to the grid on a need basis. The wholesale market consists of a major power generating company (GenCo) which sells energy to buyers using day-ahead periodic double auctions (PDAs). The balancing market manages the real-time balance of supply and demand using economic incentives. The distribution utility in PowerTAC owns and manages the distribution infrastructure. A detailed description of the simulator is provided in [Ketter *et al.*, 2020].

A crucial feature of PowerTAC is that the simulator provides a way for several autonomous energy brokers to compete to make profits by operating in retail, wholesale, and balancing markets. An annual PowerTAC tournament is organized wherein multiple teams deploy autonomous energy brokers that compete with each other in all three markets. The tournament consists of several games organized between brokers in different player configurations and varying weather conditions. The duration of a game is around 60 simulation days, and the simulation time is discretized into time slots corresponding to every hour of the day. During the game, a broker agent aims to develop a subscriber base in the retail market by offering competitive tariffs, which could be fixed price (FPT), tiered, or time-of-use (ToU). Brokers satisfy the energy requirement of their subscriber base by buying power in the wholesale market through day-ahead PDAs. A broker can participate in 24 double auctions corresponding to consumption slots that could be one to twenty-four hours away at any simulation time. Through subscriptions to storage customers, brokers can also manage grid imbalances.

The simulation environment broadcasts various information, including wholesale market clearing prices, anonymized uncleared bids and asks, retail market briefs consisting of new and revoked tariffs and weather information to help brokers in decision-making. Additionally, before starting a game, a two-week bootstrap simulation exercise is arranged in which the DU acts as the default broker. The output of this boot period, which includes identities and consumption data of all retail customers, average clearing prices in the wholesale market, weather and calendar information, is made known to the bro-

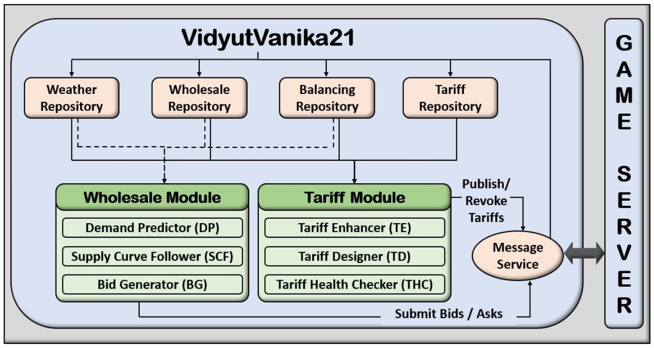

Figure 1: Architecture of VV21

kers. Equipped with all this data, a broker agent acts in various markets to attain a healthy cash position in each game compared to its peers. The final cash position of all brokers is aggregated across games to arrive at a normalized score which is used to determine the winner of the tournament.

4 Design and Overview of VV21

The system design of VV21 as shown in Fig. 1 consists of two main modules, namely, the retail (or tariff) module called VV21–RM and the wholesale module denoted as VV21–WM. VV21–RM handles tariff publication and revocation in the retail market, while VV21–WM is responsible for energy procurement in the wholesale market. Additionally, VV21 maintains repositories to store data received from the game server about the weather, wholesale, balancing and retail market, which is then used in the decision-making process.

The strategy used by VV21–RM in the retail market is inspired by game theory and aims to maintain an appropriate level of market share in the retail market. Furthermore, the retail module also targets to achieve a balanced portfolio consisting of customers of different types (consumption, production and storage). Tariffs usually contain surcharges during peak hours to further mitigate the effect of capacity transaction charges. For effective bidding in the wholesale market, VV21–WM follows the supply curve of the GenCo to recognize its lowest ask for a particular auction. It uses this lowest ask as an upper bound on its limit price and tries to procure as much power from other sellers present in the wholesale market at cheaper prices than what is demanded by GenCo. Finally, our design of VV21–RM and VV21–WM helps in handling balancing market activities as well. Precisely, VV21–WM procures a large share of energy from the wholesale market, thus reducing purchases from the balancing market, which results in low balancing costs. During high-demand scenarios, when GenCo cannot fulfill all of the market demand, VV21–RM supplies energy via producers and storage customers in its portfolio to make up for the market deficit and, in the process, generates revenue.

4.1 Retail Module of VV21 (VV21–RM)

As shown in Fig. 1, VV21–RM has three sub-modules, namely, Tariff Enhancer (TE), Tariff Designer (TD) and Tariff Health Checker (THC). The TE sub-module calculates mean tariff rates to maintain retail market share within predefined bounds. The TD sub-module generates weekly ToU tariffs based on the mean rates suggested by TE with surcharges during peak hours. The THC sub-module periodically checks for active tariffs' profitability and subscription rates and removes loss-making and under-subscribed tariffs. Algorithm 1 shows the working of VV21–RM. Below We detail TE and TD.

Tariff Enhancer (TE): TE is in charge of maintaining an appropriate level of market share in a game by improving VV21's active tariffs by considering all active tariffs present in the market. Recall that aiming to monopolize the retail market will result in high penalties for peak time usage, a fact that is recognized by brokers like TUC_TAC [Orfanoudakis *et al.*, 2021] in past tournaments. Therefore, a broker needs to carefully choose market share bounds and publish tariffs to maintain market share within the chosen bounds. More precisely, we maintain three hyperparameters HB, LB and MB which correspond to higher, lower and middle bounds. TE aims to maintain retail market share between MB and HB, which helps it attain higher profits through sufficient retail market revenue without being penalized heavily for peak time usage. Based on the game-theoretical analysis, we choose values for these hyperparameters for each player configuration. We include the details of these calculations in Sec. 5.

Algorithm 2 shows the working of TE, which suggests average tariff rates for all customer types (referred to as powerType). First, TE computes the cumulative market share of all its active tariffs using the *EstimateMarketShare()* method for each customer type. The *EstimateMarketShare()* method relies on the tariff repository shown in Fig. 1 to calculate market share of each of its active tariff. Thereafter, if VV21's current market share is higher than HB, to reduce market share, TE removes VV21's cheapest tariff and suggests costlier tariff rates by calling the method *GenerateCostlierTariffs()*. Similarly, if the market share is lower than LB, TE calls *GenerateCheaperTariffs()* that suggests cheaper tariff rates to help increase the market share. In the case when the market share is between LB and MB, TE calls *FineTuneTariffs()* to finetune the VV21's cheapest tariff by making it slightly cheaper to achieve the desired market share bounds. TE checks the viability of new tariff rates to ensure that the mean tariff rates are higher than the average energy procurement costs before calling the TD sub-module to generate weekly ToU tariffs.

Tariff Designer (TD): TD is responsible for designing a weekly ToU tariff based on the average input price (*avgPrice*) received from TE as shown in Algorithm 3. First, TD calls *DefineWeeklyTariffPattern()* method that generates a binary weekly tariff pattern indicating peak and non-peak hours. This pattern is designed by analyzing the historical net market demand values retrieved from the logs of past PowerTAC tournaments. For example, we found that morning and evening hours of a day have high demand, and these hours are marked as peak hours along with several other peak hours identified from the analysis. Then TD calls *DefineSurplusMultipliers()* method to define surplus multipliers s_i for each of the 168 hours in a week. These multipliers are greater than 1 for peak hours and are set to 1 for non-peak hours. Specifically, for peak hours, the value of s_i would depend on the magnitude of the peak as observed from market demand data.

Algorithm 1 VV21–RM

activeTariffs ← GenerateInitialTariffs()
if assesmentInterval **then**
  revokedTariffs ← TariffHealthChecker(activeTariffs)
  RevokeTariffsFromMarket(revokedTariffs)
  activeTariffs ← activeTariffs \ revokedTariffs
  newTariffs ← TariffEnhancer(activeTariffs)
  PublishTariffs(newTariffs)
  activeTariffs ← activeTariffs ∪ newTariffs
end if

Algorithm 2 TariffEnhancer(activeTariffs)

for powerType *in* offeredPowerTypes **do**
  tariffs ← getTariffs(activeTariffs, powerType)
  marketShare ← EstimateMarketShare(tariffs)
  if marketShare $\geq$ HB **then**
    RevokeCheapestTariff()
    newTariffs ← GenerateCostlierTariffs(tariffs)
  else if marketShare $\leq$ LB **then**
    newTariffs ← GenerateCheaperTariffs(tariffs)
  else if marketShare $\leq$ MB **then**
    newTariffs ← FineTuneTariffs(tariffs)
  end if
  newToUTariffs = []
  for tariff *in* newTariffs **do**
    if isTariffViable(tariff) **then**
      ToUTariff ← TariffDesigner(tariff, powerType)
      newToUTariffs ← newToUTariffs ∪ ToUTariff
    end if
  end for
end for
return newToUTariffs

Algorithm 3 TariffDesigner(avgPrice, powerType)

pattern ← DefineWeeklyTariffPattern().
$s[]$ ← DefineSurplusMultipliers(pattern)
find $normRate$: $\frac{\sum_{i=1}^{168} s_i * normRate}{168}$ = avgPrice
rate[i] ← $s_i * normRate$, for $i \in \{1, 2, ..., 168\}$
ToUTariff ← CreateTariff(rate, powerType)
return ToUTariff

We observed that such tariffs help in mitigating the effect of capacity transaction penalties, and this is shown in Sec. 7.

4.2 Wholesale Module of VV21 (VV21–WM)

The wholesale module of VV21 (VV21–WM), as explained in Algorithm 4, aims to procure the energy needed for its subscriber base at the lowest possible cost. To this end, first, VV21–WM uses the demand predictor (DP) sub-module to estimate the expected demand (Q_T) of its subscriber base for a delivery time slot T. After that, the bid generator (BG) sub-module participates in periodic double auctions (PDA) by placing bids of the form (p, q) where q is the amount of energy sought to be procured at no more than price p. Recall that in the PowerTAC setting, for a delivery time slot T, a broker can start procuring energy from time slot $T - 24$ at hourly intervals. Hence, BG has 24 opportunities to buy the required energy. This also implies that at any bidding time slot t, BG participates simultaneously in 24 auctions corresponding to 24 future (delivery) time slots denoted by $(t+1, t+2, \cdots, t+24)$. BG can submit multiple bids for a single auction. The supply curve follower (SCF) sub-module aids BG in determining suitable limit prices for bids that are to be placed in an auction. Any shortfall or excess in energy procurement by BG after exhausting all 24 opportunities is adjusted in the balancing market, which is typically expensive for the broker. We describe SCF and BG sub-modules here and use the PowerTAC sample broker's predictor as a DP sub-module. Note that, although the wholesale module is described from a buyer's perspective, it is straightforward to design a similar behavior for brokers who wish to sell energy using production customers in their portfolio.

Supply Curve Follower (SCF): The SCF sub-module helps in determining suitable limit prices for generating bids. Recall that, in the PowerTAC wholesale market, a GenCo acts as the leading supplier. The asks placed by this GenCo follow a quadratic price curve. The core idea behind SCF is to follow this quadratic price curve through uncleared asks made available to the broker, which helps in identifying the least ask that the GenCo could exhort for a particular auction. Since the wholesale auctions are PDAs, each auction can be marked using the tuple (t, T) where $t \in \{T-24, \cdots, T-1\}$ are the bidding times, and T is the delivery time. VV21–WM does not place any bids at the first opportunity (i.e., at $t = T-24$). After that, at any bidding time $t = T-k$ for $k \in \{23, \cdots, 1\}$, the least price that the GenCo could ask is identified using the uncleared asks of the previous auction corresponding to bidding time $t = T - k - 1$. This least ask price would also depend on the outstanding energy requirement of VV21 and its competitors for the delivery time T. When the uncleared asks are empty, the least ask price is taken as the maximum of the market-clearing prices of past auctions where the difference between the delivery time and bidding time is $T - t$. Algorithm 5 outlines the details of SCF.

Bid Generator (BG): The BG sub-module is responsible for generating multiple bids (M bids in total) on behalf of VV21–WM using the outputs of DP and SCF. Based on empirical observation, BG recognizes that it is possible to procure energy at a cheaper price than the least ask of the GenCo as there are other smaller sellers (like opponent brokers) who sell their excess procurement in the wholesale market. However, it is difficult to predict the exact time and ask prices of such sales. Hence, when the bidding time t is far from the delivery time T, BG generates bids with limit prices less than the predicted least ask of the GenCo (using hyper parameters α_f, β_f), hoping to capitalize on asks from smaller sellers. Closer to the delivery slot, BG generates bids with limit prices closer or greater than the least ask of the GenCo (using hyper parameters α_c, β_c). This ensures procurement of the outstanding energy from GenCo and thereby avoids going to the balancing market where energy costs are typically high. Algorithm 6 describes the working of BG.

5 Optimizing VV21–RM Parameters

In this section, we present a game-theoretical analysis to determine the optimal market share bounds determined by parameters HB, MB and LB used in Algorithm 2 specific to player configuration. We show the analysis for the 5-player configuration and follow the same methodology for other configurations as well. We model the PowerTAC game as

Algorithm 4 VV21–WM(currentTime)

```
mktUsage[] = NetDemandPredictor(currentTime)
askPrices[] = SupplyCurveFollower(currentTime, mktUsage)
bids[][] = BidGenerator(currentTime, askPrices)
for hour in [1,2...23]) do
    futureTime = currentTime + hour
    Fetch M bids from bids[futureTime]
    Submit M bids to Auction(currentTime,futureTime)
end for
```

Algorithm 5 SupplyCurveFollower(currentTime, netDemand)

```
for hour in [1 . . . 23] do
    futureTime = currentTime + hour
    Get UnclearedAsks of Auction(currentTime-1,futureTime)
    if UnclearedAsks is empty then
        Get ClearedPrices of Auction(t,t+hour) ∀ t<currentTime
        askPrices[futureTime] = max [ClearedPrices]
    else
        Sort UnclearedAsks (p_i, q_i) such that p_i < p_{i+1}
        brokerBalance = energyBought(futureTime)
        q̂ = netDemand[futureTime] - brokerBalance
        askPrices[futureTime] = min p_r s.t q̂ ≤ Σ_{i=0}^{r} q_i
    end if
end for
return askPrices[]
```

Algorithm 6 BidGenerator(currentTime, askPrices)

```
estUsage[] = DemandPredictor(currentTime)
for hour in [1,2...23] do
    futureTime = currentTime + hour
    if (currentTime is far from futureTime) then
        minPrice = α_f * askPrices[futureTime]
        maxPrice = β_f * askPrices[futureTime]
    else
        minPrice = α_c * askPrices[futureTime]
        maxPrice = β_c * askPrices[futureTime]
    end if
    Sample M prices in the range [minPrice,maxPrice]
    Distribute estUsage[futureTime] uniformly across M prices
    Create bids (q_i, p_i) ∀i ∈ [1, · · · , M] into bids[futureTime]
end for
return bids[][]
```

a *two-player zero-sum game* with utility u_1 for VV21 defined as the difference between its average cash position and the average cash position of the opponents. Our modeling choice helps maximize the difference between VV21's cash balance and the cash balance of opponents and thereby helps VV21 generate higher profits than opponents.

To perform analysis, we let VV21 act as a row player and all the opponents act as a (single) column player. We discretized the higher bound (HB) on market share to the strategy set $S_1 = \{0\%, 15\%, 30\%, 45\%, 60\%, 75\%, 100\%\}$. To generate the column player's strategy set S_2, we consider broker agents from past PowerTAC tournaments. Specifically, we use TUC_TAC (TT), VidyutVanika18 ($VV18$), VidyutVanika20 ($VV20$), CrocodileAgent (C) and AgentUDE (A) as opponent brokers. To aid the reader, we introduce a few game theory definitions before proceeding further.

Definition 1 (Mixed Strategy)**.** *For player i, its mixed strategy σ_i is a probability distribution over the strategy set S_i, i.e., $\sigma_i(s_i), s_i \in S_i$ indicates the probability with which player i plays s_i.*

VV21 / Opp	(TT, VV18, VV20,C)	(TT, VV18, VV20,A)	(TT, VV18, A,C)	(TT, A, VV20,C)	(A, VV18, VV20,C)
0%	-0.893	-0.298	-0.169	-0.156	1.737
15%	-0.199	-0.017	-0.205	-0.146	1.581
30%	0.112	-0.049	0.106	0.044	1.898
45%	-0.083	0.041	0.159	0.143	1.808
60%	-0.312	0.027	-0.288	-0.102	1.741
75%	-0.493	-0.228	-0.373	-0.409	1.025
100%	-0.498	-0.561	-0.188	-0.188	0.996

Table 1: Five player game analysis (with utility values in millions)

Definition 2 (Mixed Strategy Nash Equilibrium (MSNE))**.** *Given a N player game $\Gamma =< N, (S_i), (u_i) >$, a mixed strategy profile $(\sigma_1^*, ..., \sigma_n^*)$ is called a mixed strategy* Nash equilibrium *if,* $\forall i \in N, \; u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*), \forall \sigma_i \in \Delta(S_i)$. *$\sigma_{-i}^*$ denotes mixed strategies of all players except i.*

The utility values of the game shown in Table 1 are calculated by playing a set of m games for each combination of VV21's and column player's strategies. With the help of Gambit [McKelveya *et al.*, 2014], we found that the above game has a unique Nash equilibrium, where VV21's MSNE suggests to randomize between 30% and 45% market shares with probabilities 0.43 and 0.57, respectively, which results in equilibrium market share of 38.55% $(0.43 * 30 + 0.57 * 45)$. Similar analysis reveals that for 7 and 3 player games, equilibrium market shares turn out to be 30% and 48%, respectively.

The above analysis suggests how we should randomize in targeting market share. On top of this randomization, due to the stochasticity of the PowerTAC simulation and customer models, it is difficult to maintain one particular market share across different games. Hence, we aim to maintain market share within specific bounds such that VV21's average market share is close to the equilibrium market share. Based on the above analysis, for PowerTAC 2021, we decided to keep HB around 45% for 7 and 5 player configurations and 60% for 3 player configurations. The values of MB and LB are set to 60% and 40% of the HB. PowerTAC 2021 finals analysis shows that VV21 maintained an average market share of approximately 38%, 40%, and 55% in 7, 5, and 3 player configurations, respectively, which are in line with the equilibrium market shares. Through ablation experiments detailed in Sec. 7, one can see that VV21's retail strategy of achieving equilibrium market share contributed immensely to its success.

6 Performance in PowerTAC 2021 Finals

We now present the performance of VV21 in the finals of PowerTAC 2021. The 2021 edition of the tournament had 7 finalists; the other six competing brokers were TUC_TAC (TB1), CrocodileAgent2020 (TB2), IS3 (TB3), COLDPOWER2021 (TB4), UTA_PTA2021 (TB5) and Mertacor2021 (TB6), as shown in Leaderboard in Table 2. The tournament had a total of 386 games played across 3,5 and 7 player configurations in three real-world locations denoted as L1, L2, and L3. Each broker had to play 50 7-player, 90 5-player, and 90 3-player games.

Table 2 shows the leader-board of 2021 finals having the normalized scores of each broker across each game config-

Brokers	7-Player	5-Player	3-Player	Total
VV21	1.32 (100)	1.48 (100)	1.77 (100)	4.58 (100)
TB1	0.72 (26)	0.82 (48)	0.54 (55)	2.08 (53)
TB2	0.38 (-13)	0.52 (24)	0.91 (68)	1.81 (59)
TB3	0.64 (18)	0.42 (17)	-0.50 (17)	0.57 (17)
TB4	-0.38 (-105)	-0.87 (-83)	-0.86 (4)	-2.11 (-13)
TB5	-0.87 (-164)	-1.43 (-127)	-0.95 (1)	-3.25 (-25)
TB6	-1.81 (-277)	-0.94 (-88)	-0.92 (2)	-3.67 (-20)

Table 2: Leader board of PowerTAC 2021 Finals

Broker	L1	L2	L3	3-Player	5-Player	7-Player
TB4	-103.1	-113.1	-318.4	-148.5	-252.2	-134.1
TB2	-85.44	-99.32	-123.4	-97.42	-97.47	-115.9
TB6	-69.42	-89.03	-219.4	-127.8	-143.6	-182.4
TB1	-82.32	-92.94	-214.2	-126.2	-126.9	-110.2
VV21	-55.72	-60.78	-39.17	-59.76	-46.63	-55.11

Table 3: Robustness of VV21–WM across locations and player-configuration. We omit the unnormalized final cash positions for brevity but indicate the percentage cash position with respect to the broker VV21. The leader board numbers establish that our autonomous broker VV21 is the winner of the 2021 finals by a huge margin. Our modeling of the PowerTAC game as a zero-sum game to maximize the difference between the cash position of VV21 and its opponents, along with the ability of VV21–WM to procure energy in wholesale and balancing markets for cheaper prices, helped the broker to achieve approximately double the cash position than the second-placed broker in each of the 3, 5 and 7 player configurations. Further analysis of the finals data establishes that VV21–WM has the lowest mean market price in power procurement from the wholesale and balancing market with respect to different player configurations and different weather locations, as shown in Table 3. The overall mean market price of VV21–WM in the tournament is $-54.87\$/MW$, which is relatively 79% cheaper than the price of the next cheapest broker.

Accounting analysis of the final data is presented in Fig. 2 provides the break up of broker revenue in individual markets, namely, retail (tariff), wholesale, and balancing market, along with capacity transaction charges, which are peak usage charges and the final cash position of each broker. Handsome profits in the retail market and lower capacity transaction charges are due to an effective retail strategy. Low balancing costs are due to a strong wholesale and retail strategy. Furthermore, VV21 achieved the best income-to-cost ratio of 1.67 among the revenue-making brokers of the tournament.

Upon analyzing the tournament data further, we noticed that both VV21 and TUC_TAC won more than half of their respective games. In fact, TUC_TAC won slightly more number of games than VV21, but made considerable losses in the games it lost, whereas VV21 managed to curtail its losses in the games where it could not win. In fact, VV21 ended up in negative profits in fewer games than any other broker in the tournament. The ability to curtail losses also played an essential role in the success of VV21 in the 2021 tournament.

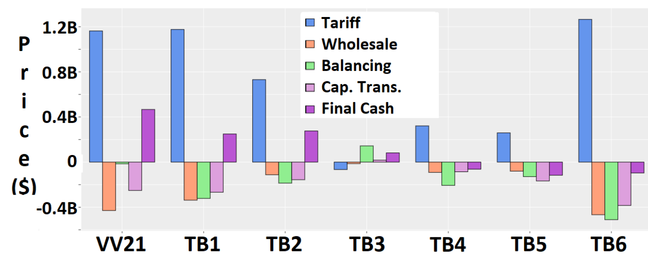

Figure 2: Accounting Analysis of Each Broker in Major Markets

Broker	Profit (%)	Broker	Profit (%)
VV21_RM	50.54	VV21_TD	79.70
VV21_WM	60.16	VV21_CONS	91.92

Table 4: Ablation analysis of VV21

7 Ablation Experiments

In this section, we perform controlled ablation experiments to analyze the efficacy of each constituent module of VV21. To this end, we disable critical modules of VV21 individually to create new brokers and report the performance drop observed by playing 50 two-player games between full VV21 and the curtailed brokers. $VV21\_RM$ and $VV21\_WM$ are agents created by substituting the retail and wholesale module of VV21 with the corresponding strategies of the PowerTAC sample broker. The agent $VV21\_TD$ is generated by disabling the tariff designer sub-module of VV21. To study the effect of maintaining a balanced portfolio, we consider an agent ($VV21\_CONS$) that offers only consumption tariffs.

The performance drop figures are reported in Table 4 which is the percentage profit the curtailed broker generates compared to full VV21. $VV21\_RM$ and $VV21\_WM$ were only able to generate 50.54% and 60.16% profits, respectively, indicating the importance of VV21–RM and VV21–WM. $VV21\_TD$ manages to achieve 79.70% profit, thus reinforcing the utility of weekly ToU tariffs with peak hour surcharges to mitigate capacity transition charges. The performance drop of $VV21\_CONS$ to 91.92% indicates the usefulness of having a balanced portfolio with production and storage customers.

8 Conclusion

In this exposition, we discussed the design of an autonomous intelligent energy broker VV21, the winner of the PowerTAC 2021 tournament. In particular, we detailed the working of retail (VV21–RM) and wholesale (VV21–WM) modules of VV21. The novelty of VV21–RM lies in utilizing game theory to determine optimal retail market share bounds specific to player configuration. Furthermore, VV21–WM uses the knowledge of the seller's bidding pattern in the wholesale market to procure power at cheaper prices and avoid making purchases in the balancing market. We presented a detailed analysis of VV21's performance in the PowerTAC 2021 tournament. Finally, ablation studies established the efficacy of each component of VV21 and showed that both VV21–RM and VV21–WM played a pivotal role in VV21's dominant performance.

References

[Chandlekar *et al.*, 2022] Sanjay Chandlekar, Easwar Subramanian, Sanjay Bhat, Praveen Paruchuri, and Sujit Gujar. Multi-unit double auctions: Equilibrium analysis and bidding strategy using ddpg in smart-grids. *arXiv,2201.10127,cs.GT*, 2022.

[Chowdhury *et al.*, 2018] Moinul Morshed Porag Chowdhury, Christopher Kiekintveld, Tran Cao Son, and William Yeoh. 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)*, 2018.

[Chung *et al.*, 2020] Hwei-Ming Chung, Sabita Maharjan, Yan Zhang, and Frank Eliassen. Distributed deep reinforcement learning for intelligent load scheduling in residential smart grids. *IEEE Transactions on Industrial Informatics*, 17(4):2752–2763, 2020.

[Ghosh *et al.*, 2019] Susobhan Ghosh, Easwar Subramanian, Sanjay P. Bhat, Sujit Gujar, and Praveen Paruchuri. Vidyutvanika: A reinforcement learning based broker agent for a power trading competition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:914–921, 2019.

[Ghosh *et al.*, 2020] Susobhan Ghosh, Sujit Gujar, Praveen Paruchuri, Easwar Subramanian, and Sanjay Bhat. Bidding in smart grid pdas: Theory, analysis and strategy. *In Proceedings of the AAAI Conference on Artificial Intelligence AAAI'2020*, pages 1974—-1981, 2020.

[Jain and Gujar, 2020] Shweta Jain and Sujit Gujar. A multiarmed bandit based incentive mechanism for a subset selection of customers for demand response in smart grids. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(02):2046–2053, Apr. 2020.

[Ketter *et al.*, 2020] Wolfgang Ketter, John Collins, and Mathijs de Weerdt. The 2020 power trading agent competition (march 30, 2020). *ERIM Report Series Reference No. 2020-002, http://dx.doi.org/10.2139/ssrn.3564107*, 2020.

[Kuate *et al.*, 2013] Rodrigue T. Kuate, Minghua He, Maria Chli, and Hai H. Wang. An intelligent broker agent for energy trading: An mdp approach. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.

[McKelveya *et al.*, 2014] Richard D. McKelveya, Andrew M. McLennan, and Theodore L. Turocy. Gambit: Software Tools for Game Theory, Version 16.0.1. http://www.gambit-project.org, 2014. [Online; accessed 27-December-2021].

[Orfanoudakis *et al.*, 2021] Stavros Orfanoudakis, Stefanos Kontos, Charilaos Akasiadis, and Georgios Chalkiadakis. Aiming for half gets you to the top: Winning powertac 2020. In *Multi-Agent Systems, 18th European Conference, EUMAS 2021*, pages 144–159, 07 2021.

[Reddy and Veloso, 2011] Prashant P. Reddy and Manuela M. Veloso. Strategy learning for autonomous agents in smart grid markets. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Two, IJCAI'11*, pages 1446—-1451. AAAI Press, 2011.

[Serrano *et al.*, 2017] Jonathan Serrano, , Ansel Y. Rodríguez González, and Munoz de Cote. Fixed-price tariff generation using reinforcement learning. *Fujita K. et al. (eds) Modern Approaches to Agent-based Complex Automated Negotiation. Studies in Computational Intelligence, Springer, Cham*, 674, 2017.

[Techopedia.com, 2021] Techopedia.com. Smart Grid. https://www.techopedia.com/definition/692/smart-grid, 2021. [Online; accessed 27-December-2021].

[Urieli and Stone, 2014] Daniel Urieli and Peter Stone. Tactex'13: A champion adaptive power trading agent. In *Proceedings of the Twenty Eighth AAAI Conference on Artificial Intelligence*, pages 465–471. Association for the Advancement of Artificial Intelligence, 2014.

[Urieli and Stone, 2016] Daniel Urieli and Peter Stone. Autonomous electricity trading using time-of-use tariffs in a competitive market. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*. Association for the Advancement of Artificial Intelligence, 2016.

[Özdemir and Unland, 2018] Serkan Özdemir and Rainer Unland. Agentude17: A genetic algorithm to optimize the parameters of an electricity tariff in a smart grid environment. *Advances in Practical Applications of Agents, Multi-Agent Systems, and Complexity: The PAAMS Collection*, pages 224–236, 06 2018.