



The 2018 Power Trading Agent Competition

Wolfgang Ketter^{ab}, John Collins^c,
Mathijs de Weerd^d

December 2017
ERS-2017-016-LIS

^aInstitute of Energy Economics, University of Cologne, Germany

^bRotterdam School of Management, Erasmus University, The Netherlands

^cComputer Science Department, University of Minnesota, USA

^dDelft University of Technology, Netherlands

Department of Technology and Operations Management
Rotterdam School of Management
Erasmus University
3062 PA Rotterdam, The Netherlands

Keywords: Autonomous Agents, Electronic Commerce, Energy, Preferences, Portfolio Management, Power, Policy Guidance, Sustainability, Trading Agent Competition

Abstract

This is the specification for the Power Trading Agent Competition for 2018 (Power TAC 2018). Power TAC is a competitive simulation that models a “liberalized” retail electrical energy market, where competing business entities or “brokers” offer energy services to customers through tariff contracts, and must then serve those customers by trading in a wholesale market. Brokers are challenged to maximize their profits by buying and selling energy in the wholesale and retail markets, subject to fixed costs and constraints; the winner of an individual “game” is the broker with the highest bank balance at the end of a simulation run. Costs include fees for publication and withdrawal of tariffs, and distribution fees for transporting energy to their contracted customers. Costs are also incurred whenever there is an imbalance between a broker’s total contracted energy supply and demand within a given time slot.

The simulation environment models a wholesale market, a regulated distribution utility, and a population of energy customers, situated in a real location on Earth during a specific period for which weather data is available. The wholesale market is a relatively simple call market, similar to many existing wholesale electric power markets, such as Nord Pool in Scandinavia or FERC markets in North America, but unlike the FERC markets we are modeling a single region, and therefore we approximate locational-marginal pricing through a simple manipulation of the wholesale supply curve. Customer models include households, electric vehicles, and a variety of commercial and industrial entities, many of which have production capacity such as solar panels or wind turbines. All have “real-time” metering to support allocation of their hourly supply and demand to their subscribed brokers, and all are approximate utility maximizers with respect to tariff selection, although the factors making up their utility functions may include aversion to change and complexity that can retard uptake of marginally better tariff offers. The distribution utility models the regulated natural monopoly that owns the regional distribution network, and is responsible for maintenance of its infrastructure. Real-time balancing of supply and demand is managed by a market-based mechanism that uses economic incentives to encourage brokers to achieve balance within their portfolios of tariff subscribers and wholesale market positions, in the face of stochastic customer behaviors and weather-dependent renewable energy sources.

Changes for 2018 are focused on stability, on simplifying interaction with the balancing market, and on encouraging more vigorous competition, and are highlighted by change bars in the margins. See Sections 3.1.2, 6, and 8.5 for details.

Contents

| | | |
|----------|--|-----------|
| 1 | Background and motivation | 1 |
| 2 | Competition overview | 1 |
| 2.1 | Simulation time | 3 |
| 2.2 | Customer market | 3 |
| 2.3 | Wholesale market | 4 |
| 2.4 | Distribution Utility | 4 |
| 2.5 | Balancing Market | 5 |
| 2.6 | Accounting | 5 |
| 2.7 | Weather reports | 6 |
| 3 | Brokers | 6 |
| 3.1 | Tariffs | 6 |
| 3.1.1 | Dynamic pricing | 10 |
| 3.1.2 | Capacity controls | 10 |
| 3.1.3 | Revoke and supersede | 11 |
| 3.2 | Portfolio management | 11 |
| 3.3 | Information available to brokers | 12 |
| 4 | Customers | 13 |
| 4.1 | Choose tariffs | 15 |
| 4.1.1 | Tariff utility | 16 |
| 4.1.2 | Choice based on tariff utility | 18 |
| 4.1.3 | Revoked and superseded tariffs | 19 |
| 4.2 | Provide balancing capacity | 19 |
| 4.3 | Generate meter readings | 21 |
| 5 | Wholesale market | 21 |
| 5.1 | Trading and time slots available for trade | 22 |
| 5.2 | Market clearing | 22 |
| 5.3 | Wholesale suppliers and buyers | 23 |
| 6 | Market-based balancing | 24 |
| 7 | Distribution utility | 29 |
| 7.1 | Distribution fees | 29 |
| 7.2 | Transmission capacity fees | 29 |
| 8 | Competition format and interaction | 30 |
| 8.1 | Initialization and Default Broker | 31 |
| 8.2 | Simulation duration | 32 |
| 8.3 | External metrics and game logs | 33 |
| 8.4 | Winner determination | 33 |
| 8.4.1 | Performance criteria | 33 |
| 8.4.2 | Final ranking algorithm | 34 |

| | | |
|----------|---------------------------------|-----------|
| 8.5 | Tournament structure | 34 |
| 8.6 | Competition rules | 35 |
| 9 | System architecture | 35 |
| 9.1 | Tournament deployment | 35 |
| 9.2 | Research deployment | 36 |
| A | Assumptions | 39 |
| B | Acknowledgements | 40 |

1 Background and motivation

We know how to build “smart grid” [2] components that can record energy usage in real time and help consumers better manage their energy usage. However, this is only the technical foundation. Energy prices that truly reflect energy surplus and scarcity can motivate consumers to shift their loads to minimize cost, and producers to better dispatch their capacities [11]. There is a significant reservoir of capacity for grid management and balancing among customer populations [17], such as water heaters, electric vehicle batteries, and cold-storage warehouses. Effective use of pricing and demand response will be critical to the effort to develop a more sustainable energy infrastructure based on increasing proportions of variable-output sources, such as wind and solar energy. Unfortunately, serious market breakdowns such as the California energy crisis in 2000 [4] have made policy makers justifiably wary of setting up new retail-level energy markets.

The performance of markets depends on economically motivated behavior of the participants, but proposed retail energy markets are too complex for straightforward game-theoretic analysis. Agent-based simulation environments have been used to study the operation of wholesale energy markets [18], but these studies are not able to explore the full range of unanticipated self-interested or destructive behaviors of the participants. Smart grid pilot projects, on the other hand, are limited in their ability to test system dynamics for extreme situations. They also lack the competitiveness of open markets, because a single project consortium typically controls and optimizes the interaction of all parts of the pilot regions. Therefore, we offer Power TAC, an open, *competitive* market simulation platform that will address the need for policy guidance based on robust research results on the structure and operation of retail energy markets [12, 14]. These results will help policy makers create institutions that produce the intended incentives for energy producers and consumers. They will also help develop and validate intelligent automation technologies that will allow effective management of retail entities in these institutions.

Organized competitions along with many related computational tools are driving research into a range of interesting and complex domains that are both socially and economically important [14, 3]. The *Power Trading Agent Competition*¹ is an example of a platform that supports the *Competitive Benchmarking* [13] research paradigm. Earlier successful examples of Trading Agent Competitions include the Trading Agent Competition for Supply-Chain Management (TAC SCM) [7] and the Trading Agent Competition for Ad Auctions (TAC AA) [10].

2 Competition overview

The major elements of the Power TAC scenario are shown in Figure 1. Competing teams will construct trading agents to act as self-interested “brokers” that aggregate energy supply and demand with the intent of earning a profit. In the real world, brokers could be energy retailers, commercial or municipal utilities, or cooperatives. Brokers will buy and sell energy through contracts with retail customers (households, small and medium enterprises, owners of electric vehicles), and by trading in a wholesale market that models a real-world market such as the European or North American wholesale energy markets [5]. Brokers compete with each other to attract customers by offering *tariff* contracts to a population of customers (households, businesses, industrial facilities). Contract terms may include fixed or varying prices for both consumption and production of energy, along with other incentives such as rebates for energy conservation, or even sign-up bonuses or

¹For up-to-date information see the project website at <http://www.powertac.org>

early-withdrawal penalties. Separate contracts may be offered for charging electric vehicles, which could limit charging during high-demand periods, or even offer to pay the customer for feeding energy back into the grid at certain times. Variable prices may follow a fixed schedule (day/night pricing, for example), or they may be fully dynamic, possibly with a specified advance notice of price changes. Dynamic pricing could motivate some customers to invest in “smart” appliances that can receive price signals and adjust energy use to control costs.

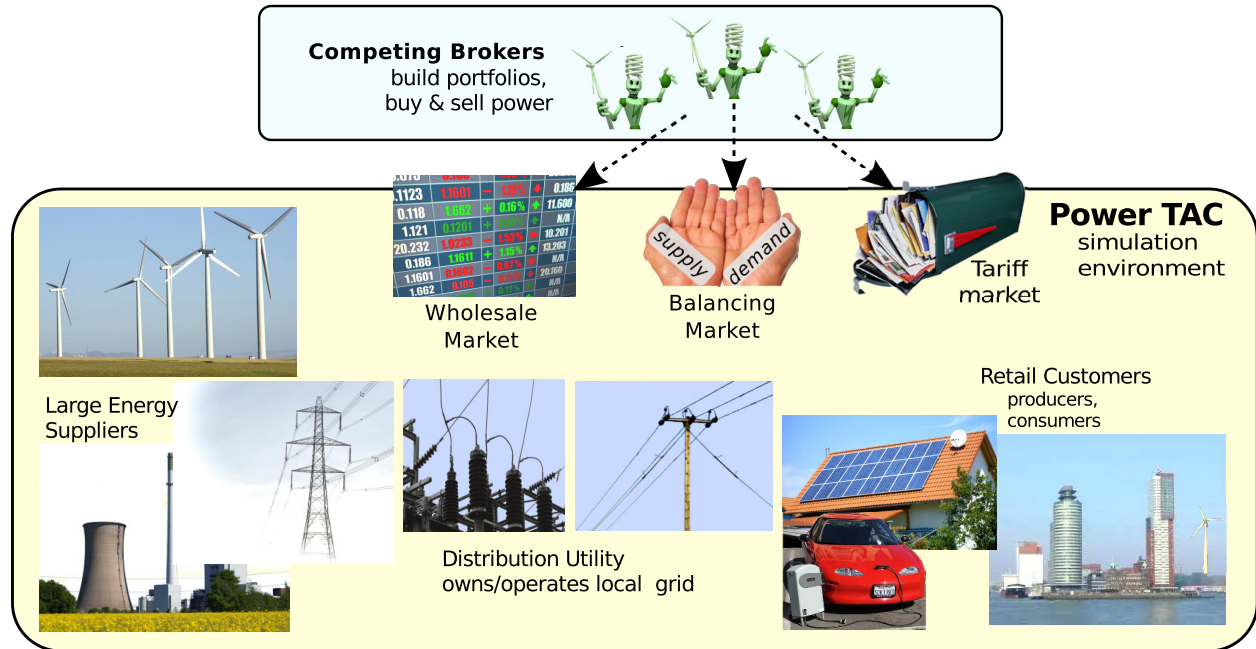


Figure 1: Major elements of the Power TAC scenario.

The simulation is designed to model energy markets primarily from an economic rather than from a technical viewpoint, and therefore we currently do not simulate the physical infrastructure (see Appendix A for a list of assumptions). In the future, it would be possible to integrate the Power TAC market simulation with a physical simulation in order to be able to evaluate the technical feasibility of the market’s energy allocation over time.

Broker agents are challenged to operate profitably by planning and executing activities over multiple timescales in three markets, a customer market, a wholesale market, and a market for balancing resources. Over a planning horizon from weeks to months, brokers build portfolios of consumer, producer, and electric vehicle customers by offering tariff contracts. At the operational level, over a time horizon of 24 hours, brokers must balance the fluctuating energy demands of their contracted energy consumers against the actual output of their contracted energy producers. Projected differences between supply and demand must be accommodated by influencing the levels of supply and demand among customers using price signals (demand response), by exercising controls on customer capacity (demand management), and by purchasing or selling energy in the wholesale market. Residual short-term imbalances between supply and demand in a broker’s portfolio are resolved in the balancing market, by exercising control over customer resources and, if necessary, by wholesale regulation resources. Retail market dynamics thus influence the wholesale market and

vice versa.

A broker’s primary goal in portfolio development is to develop a good-quality set of tariff subscriptions and individual contracts with customers who will sell or purchase energy. The ideal portfolio is profitable and can be balanced, at least in expectation, over a range of environmental conditions. A secondary goal is to manage financial and supply/demand imbalance risks. For example, an broker will benefit from having reasonably-priced energy sources that can be expected to produce energy when demand is expected to be highest within its load portfolio. Predictability is also important, and will generally improve both with volume and with a balanced portfolio of uncorrelated generation capacities and loads. Risk can be managed by acquiring uncorrelated sources and loads that can be expected to balance each other in real time, by acquiring storage capacity, by acquiring flexible consumption and generation capacities (balancing capacity), by selling variable-price contracts, and by trading future energy supply contracts on the wholesale market.

We summarize major features of the simulation in the remainder of this section. We then examine brokers, customers, and the wholesale market more closely, followed by discussion of competition rules and format, and the architecture of the Power TAC software infrastructure.

2.1 Simulation time

In the Power TAC simulation, time proceeds in discrete blocks or “time slots,” each one hour in simulated time. Each time slot takes nominally 5 seconds of real time. A typical simulation runs for roughly 60 simulated days, or 1440 time slots, over approximately 2 hours of real time. At any given time, there is a “current” time slot, and a set of “enabled” future time slots for which the wholesale market is open for trading. A primary goal of a broker is to achieve balance between energy supply and demand in each future time slot, primarily through interactions in the customer market and through trading energy delivery commitments for enabled time slots in the wholesale market.

The simulation environment depends on clock synchronization between the simulation server and the brokers. For this to work correctly, the server and brokers must be installed on machines that synchronize their clocks using `ntp`, the Network Time Protocol [15]. Synchronization of simulation time is initialized by the `SimStart` message, sent to brokers at the start of a simulation. In rare cases where the server cannot complete its processing on time, it pauses the clock by issuing a `SimPause` message to signal that the clock is stopped, and a `SimResume` message with a revised clock offset to restart the clock. In the tournament configuration, the clock is paused whenever less than 2 seconds remains between sending the `TimeslotComplete` message (the last message sent in each timeslot) and the start of the next timeslot.

2.2 Customer market

In the customer market, broker agents try to acquire energy generation capacity from local producers, and load capacity from local energy consumers. Brokers buy and sell energy in the customer market by offering tariff contracts that specify pricing and other terms, and customers choose among the tariffs on offer whenever they decide to evaluate tariffs. New tariffs and certain types of tariff modifications may be posted by a broker at any time, without regard to the daily and hourly cycle of the simulation. However, tariffs will be published to retail customers and to competing brokers in batches, once every six simulated hours.

Power TAC supports rich tariff specifications modeled on current developments in real-world electricity markets. Brokers can specify periodic payments, time-of-use tariffs with hourly or daily intervals, tiered rates, sign-up bonuses and early withdrawal fees, as well as dynamic pricing where the rate can be continuously adjusted by the broker. These tariff design elements allow brokers to shape and control their portfolios.

Tariff contracts are able to specify

Time: including points in time, time intervals, and periodicity. These terms can be used to specify contract duration as well as other time-related contract terms.

Energy: including amounts of energy produced or consumed, and rate of production or consumption (power). Contracts or tariffs may also specify amounts of energy that can be remotely controlled or curtailed, for example by shutting off a domestic water during peak demand periods, or by discharging a battery into the grid. Such remotely-controllable sources or loads are called “controllable capacity.”

Money: Agreements may specify payments to or from the customer based on time (one-time sign-up fee or bonus, fixed monthly distribution fees), or time and energy (fixed or variable prices for a kilowatt-hour).

Communication: contract award and termination, notification of price changes, exercising capacity controls, etc.

To develop their portfolios, brokers will need to estimate and reason about consumer and producer preferences as well as actions of competitors in order to design appropriate tariffs.

2.3 Wholesale market

The wholesale market allows brokers to buy and sell quantities of energy for future delivery, typically between 1 and 24 hours in the future. For this reason, it is often called a “day-ahead market”. The Power TAC wholesale market is a periodic double auction [21], clearing once every simulated hour [21]. Participants include the brokers and a set of wholesale participants that provide bulk energy and liquidity to the market, simulating the demand of a regional market that is substantially larger than the demand represented by the Power TAC simulation scenario.

2.4 Distribution Utility

The Distribution Utility (or simply DU) represents the regulated electric utility entity that owns and operates the distribution grid. It plays three roles in the Power TAC simulation:

1. It distributes energy through its distribution grid to customers. In this role it is a natural monopoly, and in the real world may be a cooperative, a for-profit regulated corporation, or a government entity. Each broker must pay distribution fees for the use of the distribution grid in proportion to the number of customers it serves.
2. It provides the connection to the transmission grid, where energy is imported from and exported to the wholesale market. Because the capacity and costs of both the transmission and distribution grids and associated services are based on maximum expected demand, these costs are assessed to brokers in proportion to their contributions to peaks in net demand as viewed from the transmission grid.

3. It acts as the broker of last resort, offering simple “default” tariffs for energy consumption and production. In this role it is called the “default broker,” simulating the electric utility in a non-competitive regulated customer market that typically exists prior to market liberalization. The default tariffs also form a “ceiling” that constrains the potential profitability of brokers, because customers are always free to choose the default tariffs over competing broker offerings. The default broker role is an essential element of the simulation, because customers must always have access to energy, and therefore at the beginning of a simulation all customers are subscribed to the default tariffs. Brokers must lure them away using more attractive terms.

2.5 Balancing Market

The Balancing Market is responsible for real-time balance of supply and demand (see Section 6) on the distribution grid. The market creates an incentive for brokers to balance their own portfolios of energy supply and demand in each time slot by ensuring that they would be better off balancing their portfolios than relying on the balancing market to do it. Note that in the real world, this function is typically handled higher in the grid hierarchy, through ISO/TSO organizations [5] and their “ancillary services” markets; Since Power TAC does not model the full hierarchy, the balancing market provides a simplified version of the reserve and regulating capacity markets and associated controls normally operated by an ISO/TSO. Brokers can exploit the balancing market in two ways: first, by having a net imbalance that reduces the overall imbalance; and second, by offering storage and flexible load resources from customers who have agreed to external controls in exchange for appropriate compensation.

2.6 Accounting

To ensure consistency and fairness, the Power TAC simulator keeps track of broker cash accounts, customer subscriptions, and wholesale market positions. Cash accounting records customer transactions for tariff subscription and withdrawal, and power consumption and production. Other transactions include tariff publication fees, distribution fees, wholesale market settlements, balancing market settlements, and interest on debt. Market position accounting tracks commitments in the wholesale market for each broker in each time slot. This information is needed by the Balancing Market to run its balancing process.

Each broker has a cash account in the central bank, and starts the game with a balance of zero in the account. Credits and debits from the various transactions are added to the account during each time slot. Brokers are allowed to carry a negative balance during the course of the game.

When the broker’s balance is negative, the broker is charged interest on a daily basis. The balance is updated daily (once every 24 hours) as

$$b_{d+1} = (1 + \beta)b_d + \text{credits}_d - \text{debits}_d \quad (1)$$

Where b_d is the balance for day d , β is the daily loan interest rate. A typical daily loan interest rate is $\beta = 10\%/365$.

When the broker’s balance is positive, the broker is paid a daily interest. This is done by updating the daily balance as

$$b_{d+1} = (1 + \beta')b_d + \text{credits}_d - \text{debits}_d \quad (2)$$

Typical daily savings interest is $\beta' = 5\%/365$.

Values for β and β' are provided to the broker at the beginning of the game (see Table 2 on page 32 for standard tournament values).

2.7 Weather reports

Weather forecasts and current-hour weather conditions are sent to brokers in each time slot. Some customer models will use this information to influence energy consumption (temperature, for example), and production (wind speed, cloud cover). Brokers with weather-sensitive customers will also need this data to predict production and consumption. Weather reports and forecasts will be drawn from real-world weather and forecast history data for some real-world location. The specific location and date range for the weather dataset is privileged information, not revealed to brokers. However, the latitude and time-of-year are given, because these variables affect the output of solar producers.

3 Brokers

Figure 2 provides a simplified overview of the timeline and information exchange between a broker and the simulation environment in each time slot. While the sequence of major processes in the simulation environment is fixed, brokers can send messages at any time, as long as they arrive before the server needs them.

In each time slot, a broker may initiate any of the following actions.

Offer new tariffs (customer market): Design and submit new tariffs for publication to customers.

Modify tariffs (customer market): Change tariff terms for existing customers by replacing a superseded tariff with a new one.

Adjust prices (customers): Adjust prices for existing tariffs, if tariff terms allow it.

Curtail demand (customers:) For those customers who have subscribed to tariffs that allow for curtailment, brokers may exercise curtailment to manage overall demand or to reduce demand when wholesale prices are high.

Submit balancing order (balancing market): Offer controllable capacities for real-time balancing, to the extent allowed by tariff terms. See Section 4.2 for details.

Submit asks and bids (wholesale market): Create asks and bids to sell or procure energy for future time slots. See Section 5 for details.

In the remainder of this section we describe the broker’s view of the simulation in more detail.

3.1 Tariffs

Brokers design and offer tariffs, and may also modify existing tariffs by superseding them with a new ones, then revoking the original tariffs. Each tariff applies to a specific `PowerType`, such as general consumption, interruptible consumption, general production, solar production, electric vehicle, thermal storage, etc. The detailed structure of a tariff offering is shown in Figure 3. This




Figure 2: Overview of Power TAC activities within one time slot. A broker interacts with the wholesale, customer, and balancing markets, and receives information from the weather service, customers, the balancing market, and the accounting service.

structure supports a number of features within a simple, compact object graph. Many concepts are represented in the `TariffSpecification` type (payments, energy-type), but the `Rate` structure is specified separately, allowing for a range of rate structures without requiring space (memory and bandwidth) for unused features. This also allows a simple convention of empty references for unused features.

Quantities of money and energy in `TariffSpecifications` and associated structures are represented from the viewpoint of a Customer. For money, this means that a positive value represents payment from the Broker to the Customer, while a negative value represents payment from the Customer to the Broker. Similarly, a positive quantity of energy represents energy delivered to the Customer, and a negative quantity represents energy delivered to the Broker. In all communications with

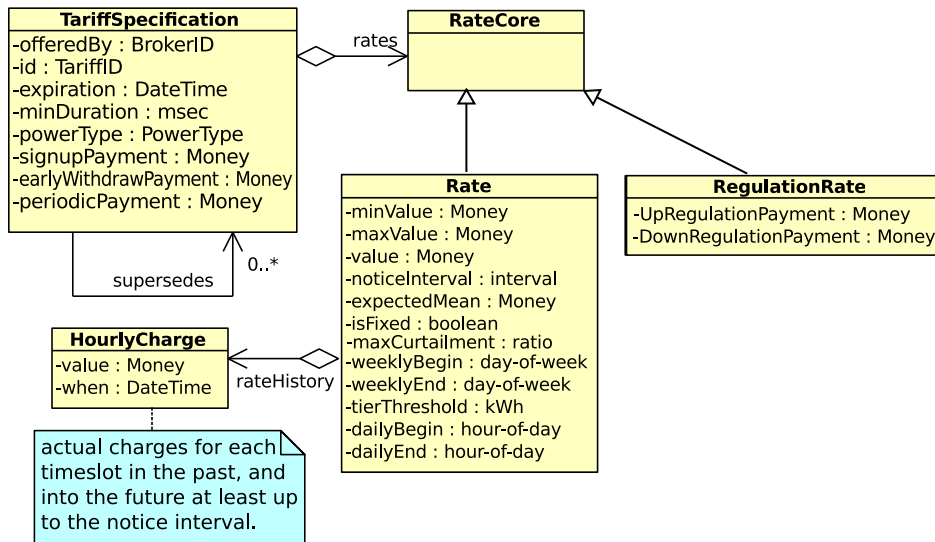


Figure 3: Tariff structure. Details are available in the software documentation.

customers, quantities of energy are represented in kWh.

Here are some common tariff features that can be represented with this structure:

- Tiered rates are specified by providing multiple Rates with different values for tierThreshold. For example, if a Tariff has Rate r1(tierThreshold=0, value=-.10) and Rate r2(tierThreshold=20, value=-.15), customers would pay 0.1/kWh for the first 20 kWh in a day, and 0.15 for any additional usage during the day (a “day” is midnight-to-midnight).
- Time-of-use rates are specified by some combination of dailyBegin/dailyEnd and/or weeklyBegin/weeklyEnd values. The dailyBegin/dailyEnd values are in hours past midnight, and weeklyBegin/weeklyEnd values are day-of-week, in the range 1=Monday through 7=Sunday. For example, an overnight rate could be specified as dailyBegin=23, dailyEnd=6. Similarly, a weekend rate would have weeklyBegin=6, weeklyEnd=7.
- Two-part tariffs (fixed daily fee plus usage fee) are specified by including a non-zero periodicPayment, which specifies the daily fixed charge. The actual payment will be 1/24 of the periodicPayment every hour.
- Signup payments in either direction (negative for a fee, positive for a bonus) are paid when a Customer subscribes to a Tariff. A negative signup payment (paid by the customer to the broker) must be fully refunded to the customer if the Tariff is ever revoked.
- Early withdrawal penalties are specified by including a non-zero minDuration (in msec) and a non-zero earlyWithdrawalPayment (negative if the customer pays the broker).
- Variable rates must specify minValue, maxValue, and expectedMean values, along with a noticeInterval in hours. More detail on specifying and updating variable rates is provided below.

- Interruptible rates allow for some portion of the Customer’s load or production to be curtailed during a timeslot in order to reduce overall energy costs or to reduce the cost of balancing. An interruptible rate is specified with a value for `maxCurtailement` between 0 and 1, which is the maximum portion of the Customer’s capacity that can be switched off in a given timeslot. Most customers will respond to a curtailment by shifting the curtailed load to the following timeslot, or possibly spread it across several future timeslots. Not all customers will fully recover a curtailed load; for example, some hot water may be consumed at a temperature lower than the setpoint of the water heater.
- For energy storage devices, `RegulationRates` can be included that specify separate payments for use of the device for up-regulation (customer delivers energy to the grid, or reduces its energy consumption) or down-regulation (customer draws additional energy from the grid). Typically the price for up-regulation is positive (the customer is paid) and is higher than the price the customer must pay for consumption, while the price for down-regulation is negative, but less than what the customer must pay for consumption. Note that while regulation rates may be labeled as “fast” or “slow”, only “slow” regulation rates are currently used.

It is not currently possible to write tariffs that bundle multiple power-types, such as household consumption and electric-vehicle charging.

Figure 4 shows the evolution of a single tariff from the time it is published. Brokers can submit tariffs to the market at any time (*pending*). New tariffs are published periodically by the market to customers and to all brokers, at which point they become *offered*. Once a customer subscribes, the broker is notified of the new subscription, and the tariff becomes *active*. Brokers are notified of various events on active tariffs, including customer subscribe and unsubscribe actions, and customer meter readings. Tariffs can have an expiration date, after which they are *expired* and new subscriptions are not allowed.

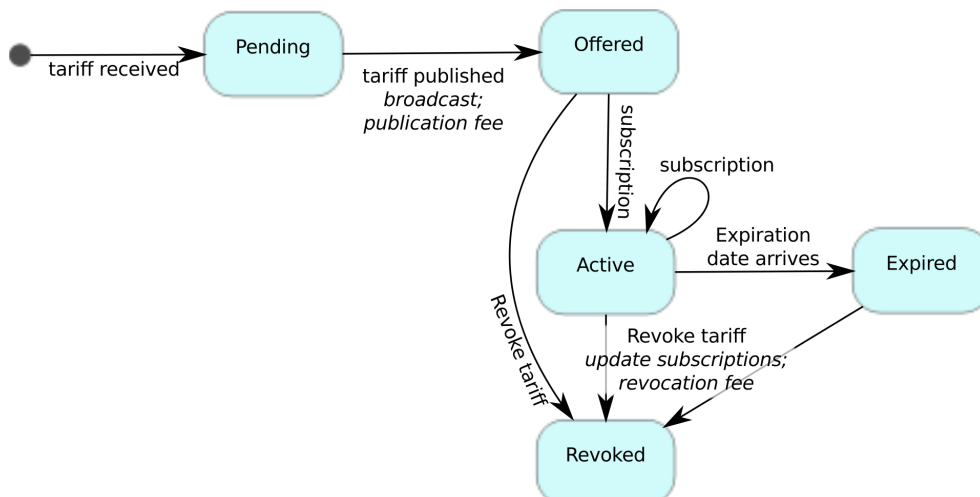


Figure 4: Tariff state transitions.

3.1.1 Dynamic pricing

In addition to time-of-use and tiered pricing, brokers can specify tariffs with variable or “dynamic” pricing. Dynamic prices must be communicated to subscribed customers some number of timeslots before the timeslot to which they apply. Brokers must therefore use some type of forecasting to determine the best price to set for each timeslot.

There are several environmental features that factor into the prices that the broker may want to charge. At a basic level, a broker typically already knows something about the price of energy to be delivered in the future from its interactions with the wholesale market. It may also want to forecast demand and supply of customers for the target timeslot. Major factors in the determination of this demand and supply include (i) the estimated or realized load and supply for timeslots preceding the target timeslot, (ii) the weather forecast conditions for the target timeslot, and (iii) customer load-shifting behaviors in response to exercise of curtailment or regulating capacity.

A variable-price tariff must specify a minimum p_{min} and maximum p_{max} price per kWh, an expected mean price $p_{em} \in [p_{min}, p_{max}]$, and a notification interval t_{notify} . Tariffs that specify minimum, maximum, and/or expected mean prices that do not satisfy these constraints will be rejected. The actual price $p_t \in [p_{min}, p_{max}]$ for a given timeslot t must be communicated to customers no later than $t - t_{notify}$. If a price is not communicated successfully for a given timeslot, then the customer will be charged p_{em} in that timeslot.

The tariff market keeps track of the actual price p_{actual} per kWh paid by customers subscribed to each variable-rate tariff. The current value of p_{actual} for each variable-rate tariff, as well as the total quantity of energy bought/sold through that tariff, is available to customers when they evaluate tariffs.

3.1.2 Capacity controls

Brokers may be motivated to offer tariffs for controllable (either *curtailable* or *storage*) capacity for two reasons:

- To reduce wholesale energy costs, a broker may directly exercise *economic controls* for a specific timeslot. For simple curtailable consumption devices, the Rate currently in effect specifies a maximum curtailment ratio r_{max} in the range $0 \leq r_{max} \leq 1$, while storage devices governed by a RegulationRate allow a wider range $-1 \leq r \leq 2$. An economic control specifies a control ratio r and a timeslot n , and must be received by the simulation server by the end of timeslot $n - 1$. For curtailable consumption devices, the ratio is constrained to $0 \leq r \leq r_{max}$, where a value $r = 1$ will result in power being shut off to the device for an entire timeslot. For battery-type storage devices, a ratio in the range $1 < r \leq 2$ will cause the battery to be discharged into the grid; $r = 2$ will discharge all the energy that is available for the given timeslot. For battery or thermal storage devices, a ratio in the range $-1 \leq r < 0$ will dump electrical energy into the device; $r = -1$ will add the maximum amount of energy the device can absorb in the timeslot. These controls are for specific timeslots, so a broker must re-issue them to extend such controls across multiple timeslots.
- To reduce balancing charges, a broker may authorize the balancing market to exercise controls against customers subscribed to its tariffs during the balancing phase, just in case doing so would be beneficial to the broker. Such controls are called *balancing controls*. A broker may issue balancing orders to the balancing market to authorize these controls for tariffs without

regulation rates, specifying the tariff, an allowable control ratio, and a price/kWh. The price is typically positive for up-regulation or consumption curtailment (the balancing market pays the broker), and negative for down-regulation or production curtailment. For tariffs with regulation rates, the prices are taken directly from those rates, and the allowable control quantities are supplied by the customers. Once received by the server, a balancing order remains in effect until it is superseded by issuing a new order.

Economic controls and balancing orders may be used concurrently for the same tariff in the same timeslot, but the economic control takes precedence, and so the actual control available to the balancing order is the difference between the allowable control and the control specified in the economic control.

In order to make such tariffs attractive to customers, brokers must factor in the future cost of customer inconvenience resulting from service interruptions. They must also deal with the load-shifting behavior of customers, because curtailment or regulation will generally result in compensating load showing up in future timeslots.

3.1.3 Revoke and supersede

In addition to changing hourly prices on variable-rate tariffs, it is possible to “modify” a tariff by revoking it and superseding it with a replacement tariff. The superseding tariff must be received (but not necessarily published) before revoking the original tariff. All subscriptions to the original tariff will be moved to the superseding tariff during the next tariff-publication cycle. However, for customers whose subscriptions are changed in this way, the withdrawal penalty for the superseding tariff is set to zero, and they will have an opportunity to re-evaluate their subscriptions before actually using or producing any energy against the superseding tariff.

3.2 Portfolio management

The primary goal of a broker is to earn a profit. To do this, it may offer tariffs for energy sources and loads that result in a portfolio that is profitable and balanced, at least in expectation, over some period. For example, a broker will benefit from having reasonably-priced energy sources that can be expected to produce energy when demand is expected to be highest within its load portfolio. Predictability is also important, and will generally improve both with volume (because noise as a proportion of demand or supply will be lower with larger numbers of randomly-behaving sources and load, even if they are correlated to some extent) and with a balanced portfolio of uncorrelated energy sources and customers.

A secondary goal is to manage financial and supply/demand imbalance risk. Such risk can be managed by acquiring producers and consumers that can be expected to balance each other in real time, by acquiring storage capacity, by acquiring controllable consumption and production capacity that can be used as needed (balancing capacity), and by trading futures contracts on the wholesale market.

Energy production includes energy acquired through the wholesale market, and local producers (household and small-business sources) acquired by offering tariffs. Energy sources can be more or less predictable, and may be controllable to some extent, as discussed in Section 2. Energy obtained from the wholesale market is predictable. Less predictable sources include most renewable sources such as wind and solar plants, which fluctuate with weather conditions and/or time of day.

Energy consumption includes energy sold in the wholesale market, and local loads (e.g., households and businesses) acquired by offering tariffs.

Energy storage is a special type of consumption that can be used to absorb excess energy or in some cases to source energy during times of shortage. Energy can be absorbed by storage capacity that is not fully charged (down-regulation), and (if discharging is supported) sourced by capacity that is above its contracted minimum charge level (up-regulation). Storage capacity that is below its minimum charge level is considered to be a load that is possibly responsive to real-time price signals.

3.3 Information available to brokers

Here we summarize the information available to brokers at various times during the game. All of this information arrives in the form of asynchronous messages at appropriate times during a simulation. Data structure details are available in the code documentation available on the project website.

At the beginning of a simulation, after brokers have logged in but before the clock begins to run, the following **public information** is sent to each broker:

Game parameters: The parameters used to configure or instantiate the specific game. See Section 8.1 for details.

Broker identities: The identities (usernames) of the participating brokers in the current game. A particular competition participant maintains the same identity over the different rounds of a competition.

Customer records: Names and characteristics of the various customer models running in the simulation. See Section 4 for details.

Default tariffs: At game initialization, the customer market offers only the tariffs published by the Default Broker. All customers start out subscribed to the appropriate default tariff. There will be one for each different “power-type” available in the configured set of customer models.

Bootstrap Customer data: Consumption and production data for each customer model for the 14 days preceding the start of the simulation, under the terms of the default tariffs.

Bootstrap Market data: Delivered prices and quantities for energy purchased by the default broker in the wholesale market over the 14 days preceding the start of the simulation. Quantities may differ from customer consumption, because the default broker starts the bootstrap period with no data on customer behavior or its response to weather, and therefore is not able to accurately balance supply and demand for at least the first week.

Bootstrap Weather data: Weather reports for the 14 days immediately before the start of the simulation.

Weather report, Weather forecast : The current weather and the forecast for the next 24 hours.

The following information is sent to brokers once every 6 simulation hours, when tariffs are published:

Tariff updates: New tariffs, revoked tariffs and superseding tariffs submitted by all brokers. This is **public information**, sent to all brokers.

Tariff transactions: When a Broker's tariffs are published, a Tariff publication fee is charged. When customers change subscriptions, brokers receive transactions that describe the changes, along with signup bonus and early-exit penalty amounts. This is **private information** for the tariff owner.

The following **public information** is sent to all brokers once per timeslot.

Wholesale market clearing data: Market clearing prices and total quantities traded for each of the 24 trading slots in the wholesale market. This may be missing if no trades were made in a given time slot.

Wholesale market orderbooks: Post-clearing orderbooks from the most recent clearing for each open time slot, containing prices and quantities of all unsatisfied bids and asks.

Total aggregate energy consumption Total energy production and consumption for the current timeslot.

Weather report and weather forecast Weather conditions for the current time slot, and forecast for the next 24 hours.

Every time peak-demand assessment is run (nominally once/week), each broker will receive a set of Capacity transactions, one for each assessed demand peak. These transactions specify the demand threshold, the broker's contribution to the peak, and the associated fee.

The following **private information** is sent to individual brokers once per timeslot.

Tariff transactions: Customer meter readings and associated credits/debits.

Balancing and distribution transactions: Charges (or credits) from DU for each individual broker to clear the balancing market and to distribute energy.

Portfolio supply and demand: Production and consumption transactions for the broker's current customer portfolio, broken down by customer subscription (customer-tariff pairs).

Wholesale market transactions: Cleared or partially-cleared bids and asks submitted by the broker.

Market positions: Broker's updated net import/export commitments, for each of the 24 open trading time slots on the wholesale market.

Cash position: Broker's updated cash position (bank balance) after all current accounting transactions have been applied.

4 Customers

Consumers and producers are simulated using a range of *customer models*. These customer models interact with brokers primarily through the tariff market mechanism – by subscribing to tariffs offered by brokers, and by consuming and producing energy. Each customer model is characterized by a core set of information that is communicated to brokers at the beginning of a simulation. This information includes:

- **Name:** The mnemonic handle for a customer model, separate from the internally generated unique ID for each customer.
- **Population:** An integer count of the number of indivisible entities (households, offices, electric vehicles) represented by the customer model. This typically corresponds to the number of metering endpoints deployed by the DU to service the customers represented by the model. For example, if a customer model represents a single household, it would have a population of 1 even though multiple persons might occupy the household. If a model represents an office building, it might represent each tenant or each floor of the building as a separate entity.
- **PowerType:** Indicates whether a customer *consumes* or *produces* energy. It also indicates whether consumption or production is *controllable* or incorporates *storage*; i.e., the consumption or production capacity can be remotely controlled in response to economic controls or due to balancing controls that the balancing market is authorized to exercise.
- **Controllable capacity:** Three numbers are given – the total capacity in kWh, the maximum up-regulation rate (increasing energy supply to the grid) in kW, and the maximum down-regulation (decreasing energy supply) rate in kW. These numbers are zero for customers with no controllable storage capacity and for those whose storage capacities cannot be controlled. Numbers are per-individual in population models, and represent an average across the population. For further details, see the discussion in Section 4.2 below.
- **MultiContracting:** Customers with non-singular populations may have the ability to allocate a partition of the population over multiple tariffs, which may be offered by multiple brokers. Note however that all entities of the population must be allocated to some tariff at every given point in the simulation.

The currently available customer models vary along the three key dimensions of population size, power type and ability to subdivide their populations across multiple tariffs. In implementation, the customer models are broadly one of two classes:

1. *Elemental Models:* This class of models attempts to simulate customer behavior at a fine level of granularity. For example, such customers are modeled using the number of persons per household, their work/vacation schedule, the usage patterns of the individual appliances that they use, and so on [9]. Two such models are currently available representing households and office buildings (respectively in the `household-customer` and `officecomplex-customer` software modules).
2. *Factored Models:* The fine granularity of the behavioral simulations employed by the elemental models severely constrains the size of populations that can be simulated by such models. As an alternate approach, factored models simulate the aggregate behavior of larger populations and other complex entities using a generalized set of *factors* that influence their behavior. Such factors control both the tariff selection process and the consumed/produced capacities exhibited by such customers. Thoughtfully configured combinations of all of these factors can be used to instantiate specific customer types such as relatively homogeneous collections of households, offices, campuses, hospitals, factories, wind farms, solar farms, etc.

In a research environment, one can choose which of these customer models are deployed in the simulation and how they are configured. The rest of this section describes the general behavior of

both classes of customer models. Implementation variances result in slight differences, which will be highlighted as necessary.

The observable behavior of the customer models can be categorized into three areas: (i) choosing tariffs, (ii) providing interruptible capacities for balancing by the DU, and (iii) generating meter readings. We will describe each of these aspects in the following sections.

4.1 Choose tariffs

Customers actively participate in the customer market by choosing new tariffs through periodic evaluation of offered tariffs. The key part of customer tariff evaluation is calculation of the expected cost or gain over the lifetime of a contract relationship. This quantity is composed of (i) per-kWh payments related to estimated consumption and/or production, (ii) fixed periodic payments, and (iii) one-time sign-up and early-withdrawal fees or bonuses.

Since early exit from contracts is allowed (possibly with a penalty), customer models may evaluate available tariffs at any time. In this case, a proper tariff-switching evaluation has to consider the early exit fees from leaving the current tariff.

To compute a utility value, this monetary evaluation is complemented by an additional assessment of several non-monetary tariff attributes, such as broker reputation, energy sources, interruptibility properties, and realized price of variable-rate tariffs. From the currently available tariff list customers need to select a suitable one (see Figure 5). This is a two-step problem:

1. Derive the utility value for the current tariff and the new tariffs to be considered. Details are in Section 4.1.1.
2. Compare evaluated tariffs and choose a suitable one. Details are given in Section 4.1.2.




Figure 5: Tariff selection problem.

Customers do not always evaluate tariffs when given the opportunity; in fact, mostly they ignore tariff publications, considering them to be junk mail. This behavior is modeled by an *inertia factor* $I \in [0, 1]$ giving the probability that the customer will *not* evaluate tariffs during a particular tariff-publication event. However, to model the market opening at the beginning of a simulation,

we expect customers to be paying attention, and so the actual inertia parameter I_a must start out with a value of 0 as

$$I_a = (1 - 2^{-n})I \quad (3)$$

where n is a count of the tariff publication cycles starting at 0. In other words, all customers evaluate tariff offerings in the first publication cycle, but their interest tails off quickly. For a population model, $(1 - I_a)$ is the portion of the population that will evaluate tariffs and possibly switch during a particular tariff publication cycle.

4.1.1 Tariff utility

The utility of a given tariff T_i is computed as a function of per-kWh payments $p_{v,i}$, periodic payments $p_{p,i}$, a one-time signup payment $p_{signup,i}$, a potential one-time withdrawal payment $p_{withdraw,i}$ in case the customer withdraws its subscription before the tariff’s minimum duration, and an inconvenience factor x_i to account for inconvenience of switching subscriptions, and of dealing with time-of-use or variable prices or capacity controls:

$$u_i = f(p_{v,i}, p_{p,i}, p_{signup,i}, p_{withdraw,i}, x_i) \quad (4)$$

The specifics of the function f could vary slightly across customer model implementations, but in general it is the normalized difference between the cost of using the default tariff and the cost of the proposed tariff, less the inconvenience factor. For consumption tariffs, cost is estimated using an energy usage profile $C_{t,i}$ over the expected duration $t = [0..d_e]$ of a potential new subscription to tariff T_i . Note that the expected usage profile C_t might vary across tariffs to account for potential load-shifting driven by time-of-use prices.

When a storage-capable customer evaluates a tariff containing a RegulationRate, evaluation must also consider up-regulation payments and savings on energy due to down-regulation payments. Note that up-regulation payments are normally positive from the customer’s viewpoint. Up-regulation actions in the form of simple curtailments result in little or no additional energy cost, unless regulation actions cause usage to shift from a lower-cost period to a higher-cost period. For battery-storage customers, up-regulation can also result in discharging some of the battery capacity, which must be replaced with purchased energy in a future timeslot. On the other hand, down-regulation payments are generally negative from the customer’s viewpoint, and so the customer’s resulting “revenue” comes from the fact that down-regulation actions displace normal energy purchases. Assuming the expected balance in each timeslot is zero, up-regulation actions should be as likely as down-regulation actions. However, only storage devices support down-regulation, so a given storage customer may see more down-regulation than up-regulation, depending on the makeup of the customer population. The situation is a bit more complex for thermal storage devices, since a down-regulation event will store additional energy, which then does not need to be supplied by purchased energy until it dissipates. In other words, exercising down-regulation in a thermal storage device temporarily reduces up-regulation capacity.

It is up to storage-capable customers to estimate three tariff evaluation parameters: expected up-regulation kWh/timeslot due to curtailment $E(\text{kWh}_{up})$, expected up-regulation kWh/timeslot due to discharge $E(\text{kWh}_{dis})$, and expected down-regulation kWh/timeslot $E(\text{kWh}_{down})$. These values may be configured as static values or learned during a simulation. They are used to modify the per-kWh payments $p_{v,i}$ for a tariff i as

$$p'_{v,i} = p_{v,i} - p_{up,i}E(\text{kWh}_{up}) - (p_{up,i} - p_{v,i})E(\text{kWh}_{dis}) - (p_{v,i} - p_{down,i})E(\text{kWh}_{down}) \quad (5)$$

The cost of using the default tariff is

$$cost_{default} = \sum_{t=0}^{d_e} (C_{t,default} p_{v,default} + p_{p,default}) \quad (6)$$

where $p_{v,default}$ is the per-kWh cost of the default tariff (assumed to be fixed), and $p_{p,default}$ is the periodic payment specified in the default tariff. The cost to switch to tariff i for the same usage profile is

$$cost_i = \sum_{t=0}^{d_e} (C_{t,i} p_{v,i,t} + p_{p,i}) + (p_{signup,i} + F_d p_{withdraw,i} + p_{withdraw,0}) \quad (7)$$

where we include both the cost of withdrawing from the current tariff $p_{withdraw,0}$ (which is zero if the minimum duration requirement for tariff 0 has already been met) and the expected cost of withdrawing from Tariff i , discounted by a factor $F_d = \min(1.0, d_i/d_e)$, which preferentially discounts shorter commitment intervals d_i . One of the options is staying with the current tariff T_0 , in which case we have no signup fee/bonus and no withdrawal cost:

$$cost_0 = \sum_{t=0}^{d_e} (C_{t,0} p_{v,0,t} + p_{p,0}) \quad (8)$$

The normalized cost difference η_i^C for consumption tariffs is then the difference between the cost of the default tariff and the proposed or current tariff, normalized by the cost of the default tariff

$$\eta_i^C = \frac{cost_{default} - cost_i}{cost_{default}} \quad (9)$$

Note that in general, energy consumption is represented by a positive value from the customer's standpoint, and payments from customer to broker are negative values. Therefore, the "cost" values in these formulas are negative (except in very unusual cases) for both consumption and production tariffs. However, in the case of a production tariff we will benefit if we choose a tariff with a larger payout, so the sign of the cost difference is reversed:

$$\eta_i^P = \frac{cost_i - cost_{default}}{cost_{default}} \quad (10)$$

For "normal" competitive consumption tariffs, we expect to see $0 < \eta_i < 1$. A tariff that is less attractive than the default tariff will have $\eta_i < 0$, while production tariffs and some very strange consumption tariffs could produce $\eta_i > 1$. An example would be a positive signup bonus that exceeds the cost of using energy over the evaluation period.

Finally, utility is the normalized cost difference less the inconvenience factor:

$$u_i = \eta_i - w_x x_i \quad (11)$$

where $w_x \in [0, 1]$ is an attribute of individual customers, and x_i is a linear combination of factors that penalize tariff features including variable pricing, time-of-use pricing, tiered rates, and capacity controls. These factors are scaled by $\log(max/min)$, which means that a 3:1 price range is penalized half as much as a 9:1 price range. To reduce gratuitous subscription "churn" among essentially equivalent tariffs, for $i \neq 0$, x_i also includes penalties for switching tariffs and for switching brokers.

Most customers have very little loyalty to their brokers, and will therefore set the broker-switch penalty close to zero except in the case of tariff revocation (see Section 4.1.3). In the future, this definition of inconvenience may be extended to cover customer preferences over sustainability of energy sources, and possibly other factors related to customer preferences. Storage-capable customers do not consider regulation rates to be inconvenient, since they are in full control of deciding how much capacity to make available for regulation in each timeslot.

Note that tariff utility u_i can be negative even if the corresponding normalized cost difference η_i is positive, due to the influence of the inconvenience factor x_i . However, values of $u_i > 1$ should occur for consumption tariffs only if a broker offers a tariff that pays the customer to take energy. On the other hand, a production tariff that offers more than twice the default rate for producing energy could easily have $u_i > 1$.

When a tariff contains one or more variable rates (dynamic pricing), customers compute a risk-adjusted estimate of the actual cost. Four values must be combined to generate an estimate for a variable-rate tariff:

$$p_v = \alpha(w_{em}p_{em} + w_{max}p_{max}) + (1 - \alpha)p_r \quad (12)$$

where p_{em} is the broker's claim of expected mean price, p_{max} is the brokers commitment to the maximum value for the rate, and p_r is the realized price for kWh_{total} , the total energy sold through the tariff so far. The weights are constrained such that $(w_{em} + w_{max}) = 1$. The parameter α is used to adjust the weight given to the realized price based on kWh_{total} , as

$$\alpha = 1 - w_r \left(1 - \frac{1}{1 + \frac{kWh_{total}}{kWh_0}}\right) \quad (13)$$

where $w_r \in [0, 1]$ and kWh_0 are parameters specific to each customer. The assumption is that the actual realized price is more predictive for a tariff with a more substantial price history (larger amount of energy sold).

4.1.2 Choice based on tariff utility

The set of tariffs considered is a subset of tariffs that are applicable to the given PowerType. Because tariff evaluation has some cost, and because we wish to discourage the practice of “flooding” by brokers who want their tariffs to have a better chance of being chosen, customers evaluate only the most recently published N tariffs from each broker, where N contains at most 5 of each applicable type. So for an electric vehicle, there could be EV tariffs, interruptible-consumption tariffs, general storage tariffs, and simple consumption tariffs that all apply. If a broker has published 5 of each type, then for that broker, $N = 20$.

An overall tariff choice does not necessarily follow a deterministic choice of the highest utility value, because customers are not entirely rational. This is especially important for population models that represent larger groups of customers. A smoother decision rule is therefore employed to allocate customers to tariffs, based on the multinomial logit choice selection model, which allocates the choice proportionally over multiple similar tariffs. The logit choice model assigns probabilities to each tariff, t_i , from the set of evaluated tariffs, \mathbb{T} , as follows:

$$\mathbb{P}_i = \frac{e^{\lambda u_i}}{\sum_{t \in \mathbb{T}} e^{\lambda u_t}} \quad (14)$$

The parameter λ is a measure for how rationally a customer chooses tariffs: $\lambda = 0$ represents random, irrational choice, while $\lambda = \infty$ represents perfectly rational customers always choosing the tariff with the highest utility². Depending on the customer model type this choice probability can be used in two ways – either to represent somewhat randomized, not perfectly rational tariff choice in case of single customer models or to assign population shares to different tariffs in case of a population customer model.

4.1.3 Revoked and superseded tariffs

If a customer is subscribed to a tariff that is superseded and canceled, then by definition $d_i = 0$ for the new (superseding) tariff and therefore there is no withdrawal penalty. In addition, the evaluation inertia I for the affected customers is reduced to $I_s = 0$, with the result that all subscribers to the superseded tariff re-evaluate their tariff options immediately, before they consume or produce energy against the superseding tariff. Customers will find it somewhat inconvenient to switch brokers at this point, because to accept the superseding tariff requires no action by them.

4.2 Provide balancing capacity

Customers can provide brokers with different forms of “demand management” capabilities that can be used to control costs or for balancing, as determined by the **PowerType**. These differ in availability and the amount of energy available in a timeslot. Some provide up-regulation (reducing demand or increasing supply), and some provide down-regulation.

- **Interruptible consumption:** Certain types of appliances (water heaters, heat pumps) can support remote interruption, thereby providing up-regulation. This capability is indicated by a non-zero value of the customer’s **controllableKW** attribute, which specifies the maximum power usage that can be curtailed. The actual amount of up-regulation that can be achieved at any given time depends on how much the customer would have been using at that time. If a broker has interruptible capacity under contract, its use can be offered to the balancing market to avoid balancing charges.
- **Consumption with storage:** Customers with energy storage (batteries, electric vehicles, or thermal storage devices such as water heaters or refrigerated warehouses) can provide both up-regulation and down-regulation, limited by the storage unit’s capabilities and state of charge. This capability is indicated by three attributes in the customer description, as shown in Figure 6: **upRegulationKW** gives the maximum discharge rate, **downRegulationKW** gives the maximum charge rate, and **storageCapacity** gives the maximum energy available for regulation. A thermal storage device would typically have **upRegulationKW** equal to zero, because their stored energy cannot be returned in electrical form. Most storage devices, on the other hand, would have non-zero **controllableKWh**.
- **Controllable production:** While intermittent producers like solar and wind typically cannot provide balancing capabilities, non-intermittent producers like CHPs or bio-gas units may offer some ability to control capacity for balancing purposes.

²In implementation, λ is less than ∞ to avoid numeric overflow issues.

- **Withdraw energy from storage:** Up-regulation can also occur by withdrawing stored energy from electric vehicle batteries or other electrical storage capacities that currently hold more energy than the customers need.




Figure 6: Visualization of storage device behavior.

Brokers can acquire controllable capacity by offering tariffs for power types that provide capacity controls. These include interruptible consumption, electric vehicle, and thermal storage types. To allow simple curtailment, the tariff must include Rates that specify a `maxCurtailment` value between 0 and 1. Actual up-regulation capacity available at any given time will nearly always be less than the specified `controllableKW` value – for example, a curtailable water-heater can provide up-regulation only if it is currently heating the water.

When a capacity is managed using an economic control exercised by the broker or a balancing control exercised by the balancing market, the customer may forfeit that capacity (for example, a customer may have multiple heat sources) or shift some or all of it to future time slots. The degree and nature of shiftability is a customer-specific attribute, tied to the physical nature of that customer’s capacity.

Brokers can acquire the ability to manipulate battery or thermal storage devices, including electric vehicles, by offering tariffs that include `RegulationRates` specifying the prices for up-regulation and down-regulation. Constraints on the availability of regulating capacity are shown visually in Figure 6. If the device is an electric vehicle, then when it is plugged in, its state-of-charge is “initial”, and there is some “target” state-of-charge at a specific time in the future. At any given time, if the device is not fully charged, it can absorb energy at some maximum charge rate; if not fully discharged, it can supply energy at some maximum discharge rate. For a thermal storage device, there is a current temperature and some “target” temperature in the future. The maximum discharge rate for a thermal storage device is zero, since it is not possible to extract energy as electricity.

The maximum charge rate is the most power it can absorb while heating up or cooling down. For a heating device, up-regulation shuts off power and allows temperature to drop, and down-regulation raises the temperature.

In general, a storage device responds to economic or balancing controls by altering its state within the constraints defined by its “feasible region” as shown in Figure 6. Therefore, a heat storage device that is already absorbing energy at its maximum charge rate cannot respond to a down-regulation request, but it can respond to an up-regulation request by reducing its energy use, as long as its state remains within its feasible region. Similarly, a heat storage device that is not currently using power cannot respond to an up-regulation request, but can respond to a down-regulation request by raising its temperature further, as long as its temperature remains within its feasible region.

4.3 Generate meter readings

The meter readings generated by customers may depend on different factors. Intuitively we can group these into three basic groups – static, broker-dependent and game-dependent factors. Static factors are model primitives (such as the number of household members, work shift hours, equipment) that characterize the customer’s fundamental load profile independent of developments in the game. Broker-dependent factors influencing the realization of customer load profiles are the tariff (time-of-use pricing induces customers to shift consumption) as well as balancing capacity actions (responding to current or previous curtailment). Lastly, game-dependent factors include load adjustment triggered at runtime by the game environment, e.g. randomization, simulated time-of-day, current weather conditions (e.g. turning on A/C, output from solar panels).

Currently implemented customer models consider the type of customer entity (e.g., household vs. factory) and the size of population to generate a base load. That base load is then adjusted for broker-dependent and other dynamic factors. The dynamic factors currently used include day-of-week, time-of-day, current weather (including temperature, cloud cover, wind speed, and wind direction), and a 48-hour weather forecast. The capacity is further adjusted to reflect attributes of the tariffs to which the customer is currently subscribed. Under adverse prices, consumption and production are both lowered to some degree (the degree depends on the specific customer). Customers with smart shifting capabilities also adapt by moving capacity to future time slots; such effects may benefit the customers when they are faced with tiered pricing (and therefore don’t want to currently consume beyond a particularly tier), time-of-use pricing (the customer knows that they can expect better rates in future time slots), or variable-rate pricing (the customer knows or estimates that it may get better rates in the future and is therefore willing to absorb the risk and potential disutility of postponing consumption or production).

5 Wholesale market

The wholesale market in Power TAC operates as a periodic double auction (PDA) and represents a traditional energy exchange like NordPool, FERC, or EEX³. The brokers can buy and sell energy contracts for future time slots. In the wholesale market brokers interact with each other directly as well as with generation companies (GenCos) and other wholesale market participants as described below in Section 5.3.

³See <http://www.nordpoolspot.com>, <http://www.ferc.gov>, or <http://www.eex.com/en>.

5.1 Trading and time slots available for trade

Brokers can submit orders to the wholesale market for delivery between one and 24 hours in the future. The time slots available for trading are marked as “enabled”; changes in time slot status are communicated to brokers at the beginning of each time slot. Orders submitted for non-enabled (disabled or not yet enabled) time slots are silently discarded. The market collects submitted orders continuously; the orders considered for clearing are exactly the set that have arrived since the start of the last clearing.

Each order is a 4-tuple (b, s, e, p) that specifies a broker b , a time slot s , an amount of energy e in megawatt-hours, and optionally a limit price per megawatt-hour p . Energy and price quantities are treated as proposed debits (negative values) and credits (positive values) to the broker’s energy and cash accounts. So an order $(b_1, s_{12}, 4.2, -21.0)$ represents a bid (a buy order) from broker b_1 to acquire 4.2 MWh of energy in time slot s_{12} for at most 21 €/MWh. Orders that specify a limit price p are called “limit orders”, while orders that do not specify a limit price are called “market orders.”

Order quantities must be larger than a minimum order quantity e_{min} (nominally 0.1 kWh) to prevent brokers from “spamming” the market with infinitesimal orders.

5.2 Market clearing

When the simulation clock is advanced to a new time slot, the wholesale market clears the orderbook for each of the enabled time slots. At the same time, an updated list of enabled time slots is sent to each broker. This minimizes the period of time in which the set of enabled time slots from the broker’s viewpoint differs from the set of enabled time slots from the market’s viewpoint.

In the clearing process, as shown in Figure 7, demand and supply curves are constructed from bids and asks to determine the clearing price for each enabled time slot. The clearing price is the intersection of the supply and demand curves. Note that bids propose a positive energy amount and a negative cash amount, and asks have negative energy and positive cash. Also note that market orders (orders that do not specify a price) are sorted first, as though they had the highest bid prices or the lowest ask prices.

If there is not a unique price where the supply and demand curves cross, as in this example, then the clearing price is set at the mean of the lowest usable bid and the highest usable ask price. All bids with prices higher than the last cleared bid, and all asks with prices below the last cleared ask, are fully executed. In most cases, either the last cleared bid or the last cleared ask is partially executed. If the last matched bid is a market order, then the clearing price is determined by the highest ask price, with an added margin (nominally 20%). Similarly, if the last matched ask is a market order, the clearing price is determined by the lowest bid price, less a margin. If all bids and asks are market orders, the clearing price is set to a (rather high) default value; this case is highly unlikely in practice, since the wholesale players never use market orders.

In the example of Figure 7 we see bids sorted by decreasing (negative) price, and asks sorted by increasing price. Both bid 1 and ask 1 do not specify a price; these are unconstrained “market orders” and are always considered first. Bids 1-8 are all matched by lower-priced asks, and asks 1-6 are all matched by higher-priced bids, although only the first 2 MWh of ask 6 is matched. Ask 7 and bids 9-10 cannot be matched. The cleared volume is 27 MWh, and the clearing price is 16, i.e. the mean of the prices in ask 6 and bid 8.

After the market is cleared the following steps are performed:




Figure 7: Market clearing example: bid 8 and part of ask 6 are the last to clear.

- Clearing price and volume are broadcast to all brokers. In the example of Figure 7, this would be (27, 16).
- Post-clearing orderbooks are published for each cleared time slot, giving the un-cleared bids and asks, without broker information. In the example, the orderbook would include two asks ((-3, 15), (-7, 16)), and two bids ((5, -14), (7, -12)).
- Brokers are informed about their own executed transactions.
- Updated cash and market positions are computed and communicated to individual brokers.
- All orders that arrived before the start of the clearing process (and were therefore included in the clearing) are discarded.

5.3 Wholesale suppliers and buyers

A typical wholesale market in Europe or North America serves a broad region, with many providers and millions of customers. The Power TAC scenario, on the other hand, models a single city in such a region. As a result, the retail brokers in an individual city have limited market power.

In addition to the competing brokers, three additional models operate in the wholesale market:

1. The Grid Genco is an abstract entity that simulates a population of generating facilities distributed over a large geographic area, serving a number of distribution areas including the

Power TAC simulation environment. It essentially generates a supply curve drawn from the pricing statistics of nodes in a market with congestion pricing, such as the MISO or PJM ISO organizations in North America. The supply curve is composed of a succession of ask orders with prices and quantities drawn from distributions that characterize the full price curve observed in the MISO and/or PJM LMP markets, scaled for the demand range among the customers in the simulation environment.

2. The Grid Buyer simulates regional demand in the north central region of the U.S. The demand profile is temperature-dependent (higher demand when temperature is far enough from the comfort zone to require heating or air-conditioning), and varies daily and weekly. Demand is generated by a time series trained on two years of MISO North actual demand, along with matching weather data. In each timeslot, the Buyer buys enough energy in the market to satisfy its expected demand for each open timeslot, using the current weather forecast for temperature adjustment. Weather forecast changes over the 24 hours of trading for a given timeslot will cause the Buyer to buy and sell energy to adjust its market position.
3. A wholesale buyer provides liquidity with stochastic behavior that simulates a population of buyers and speculators. A real-world example might be an industrial site that uses electric power when the price is low enough for process heat or electrolysis. Its behavior is very simple: Given two parameters, a quantity q_b and a mean price p_b , and a random value $\eta \in [0, 1]$, it computes a price $p_{b,s} = -p_b \ln(1 - \eta)$ for each time slot s and places a bid $(b_{b,s}, q_b/p_{b,s}, p_{b,s})$ in each open time slot. This exponential distribution produces large numbers of low-priced high-quantity bids, and a few higher-priced low-quantity bids.

Supply and demand profiles generated by these models are scaled to adjust the market power available to Power TAC brokers; higher scale factors result in lower market power. The scale factor ψ is the ratio of mean overall market demand including the Grid Buyer to mean demand originating within the Power TAC customer population. A typical tournament setting is $\psi = 10$. Note that the Power TAC demand is not the net demand after solar production is subtracted, since the regional market has less than 1% penetration of distributed solar production.

6 Market-based balancing

In electricity markets, supply and demand have to be balanced almost perfectly in real time. A major task of the Independent System Operator (ISO)⁴ is to monitor the grid and to maintain balance while keeping voltage, frequency, and power factor within very tight bounds. This task becomes more challenging as more small-scale “non-dispatchable” renewable energy sources, such as solar and wind, are connected to the grid [19]. Many of these sources (e.g. wind) are only partially predictable.

The grid balancing problem has been studied on various levels (wholesale vs. retail) and with different approaches [16]. Since Power TAC does not represent the transmission-level grid, the balancing function of the ISO is carried out by a “Balancing Market” (abbreviated BM, possibly operated by the DU or the ISO) that has access to the “regulating market” portion of a wholesale ancillary services market that trades in balancing capacity. Brokers accumulate credits and debits to their energy budgets for each time slot by selling (exporting) energy or buying (importing)

⁴In Europe the name Transmission Systems Operator (TSO) is used instead of ISO.

energy in the wholesale market, and by the energy consumption and production activities of their contracted customers. The total net energy budget for a time slot s and a broker b is denoted by $x_{b,s}$. The sign of x is negative if the broker has an energy deficit – its portfolio and market position uses and exports more energy than it produces and imports. To carry out its responsibility to balance supply and demand in each time slot, the BM may exercise capacity controls (see below) on behalf of brokers, and it may import or export energy through the regulating market at prices that are normally much less attractive than the prices faced by brokers in the forward wholesale market (see Figure 8).

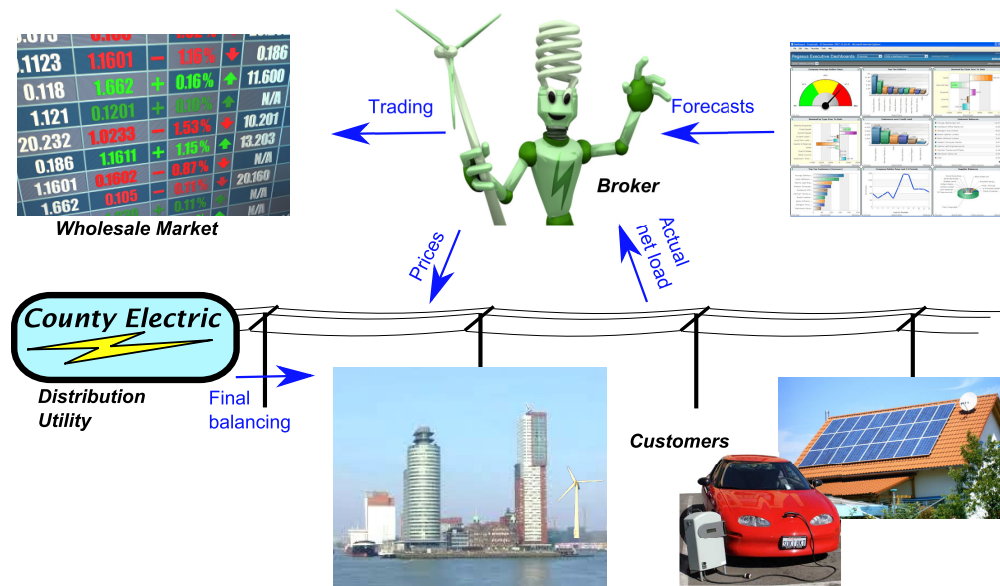


Figure 8: Entities and activities during balancing.

Detailed background and examples on market-based balancing can be found in [8]. For 2018, the BM implements a method called *static with controllable capacities*. The intent is to create a market that motivates brokers to balance themselves as closely as possible through portfolio development and wholesale market trading, and to offer controllable capacities to the BM in the form of *balancing orders* that allow the BM to exercise capacity controls among their contracted customers in order to achieve balance. For storage and interruptible customers subscribed to tariffs that include a regulation rate, the balancing order is constructed from the regulation rate. For storage customers, the available capacity in any given timeslot is provided by the customer; for interruptible consumption customers, if the capacity is not directly provided by the customer, it will be computed from the maximum curtailment given in the rate currently in effect. For these customers, the broker does not send balancing orders, and the market will ignore them if it receives them. For customers subscribed to tariffs that specify non-zero curtailment ratios and lack regulation rates, the balancing order is provided by the broker. Each broker-supplied balancing order specifies a tariff, a maximum curtailment ratio between 0 and 1, and a price. Note that curtailment ratios between 1 and 2 apply only to storage devices (batteries) that can be discharged into the grid. A balancing order allows the BM to directly manage subscriber capacities up to the specified ratio of their actual usage, for the price/kWh specified in the balancing order or the associated tariff, and remains in effect as long as the associated tariff is offered to customers or until it is

replaced by the broker. The prices brokers receive from exercise of balancing orders arise from the market clearing process and will always be at least the prices specified in the regulation rates or the broker-supplied balancing orders.

There are several constraints on the amount of energy available for capacity controls. For curtailment, the Rate currently in effect specifies a maximum curtailment ratio, and an economic control may have already been exercised against a particular tariff. Therefore, the available adjustment is the product of the unexercised ratio and the actual capacity of the customers subscribed to the tariff. For storage capacity, the customer computes its own constraints and notifies the BM of its available capacity for both up-regulation and down-regulation prior to the clearing process.

Brokers must submit their balancing orders before the customer models run (near the start of each timeslot), and the BM runs its balancing process after customer consumption and production quantities are known for the current time slot. At this point, the BM can determine the actual quantities available for adjustment against each balancing order. A balancing order remains in effect until it is superseded by submitting a new balancing order for the same tariff with different parameters.

The BM acts to resolve the net shortage (surplus) over all brokers at minimal cost (maximal profit). To achieve this in case of a shortage, given a set of balancing orders, the BM

1. discards the orders that cannot contribute to the solution; if overall balance is negative (up-regulation needed), then only consumption curtailment is used, and if overall balance is positive, then only production curtailment is used.
2. includes “dummy” orders with essentially infinite capacity that represent procurement or sale of energy in the regulating market at costs of $c_0(x_{RM})$, a linear function of the quantity x_{RM} being traded in the regulating market. For up-regulation, $c_0(x_{RM}) = P^+(s) + \phi^+ x_{RM}$, and for down-regulation $c_0(x_{RM}) = P^-(s) + \phi^- x_{RM}$, where ϕ^+ and ϕ^- are the slopes of the cost functions for up-regulation and down-regulation respectively. Note that in case there are balancing orders with prices above P^+ or P^- , the dummy orders will be split around such balancing orders. Competition values for ϕ^+ and ϕ^- are given in Table 2.
3. Sorts the remaining orders by price, with the lowest first.
4. In price order, the BM selects the lowest-price orders up to the required capacity. Note that in general, the highest-price order selected may only be partially exercised. Also, since prices are typically negative for down-regulation (the broker pays for the energy), the lowest-price down-regulation order is the one for which the broker pays the most.
5. The price for a broker b depends then on both its own imbalance, as well as on its balancing orders. This computation is the sum of a Vickrey-Clarke-Groves (VCG) payment p_{vcg} [20, 6], and an imbalance payment p_{imb} as defined in more detail below.

The payment for brokers consists of two parts: a payment for the use of its controllable capacity p_{vcg} , and a payment for its imbalance p_{imb} . Both payments typically are negative (the broker pays) in case of being short or when selling down-regulation capacity (e.g., curtail production, deposit energy in storage), and positive (the broker is paid) when it has a surplus that helps cover an overall shortage, or sells up-regulation by curtailing its consumption or withdrawing energy from storage.

The setting for choosing controllable capacity is very similar to a one-sided auction, and for this part the VCG payment is used. The VCG payment to a broker b for controllable capacity is

defined to be its marginal contribution to the overall cost of achieving balance. We compute this as the difference between the cost of the optimal combination of orders leaving out the orders of broker b , compared to the costs to the other brokers of the optimal combination using the orders of all brokers including b . Additionally, we resolve the following issues by the second part of the payment p_{imb} .

- We cover the costs of the BM for resolving the imbalance, including both the costs of “dummy orders” as well as the net payments of the brokers (note that in case of shortage at least some brokers offering up-regulation capacity will typically receive money).
- We make it uninteresting for brokers to create an imbalance to sell extra controllable capacity, and
- we provide an incentive to be as closely balanced as possible for brokers that are contributing to the imbalance.
- Additionally, the total payment by the brokers should be as low as possible (in other words, it is not a goal of the BM to earn a profit by performing this balancing task).

The idea of the imbalance payments is to let the brokers that contribute to the overall imbalance pay for both the costs of the BM as well as for the opposite imbalance other brokers may have (since they reduce balancing costs). For each broker, we remove the part the broker can influence (in this case the costs of its own controllable capacity) from the equation. Omitting the timeslot index s for brevity, we denote the set of orders for controllable capacity by C , and that of a broker b by C_b . The cost of the BM for a given net imbalance X (following from the VCG payments and possibly some dummy orders) is denoted by $\text{BM}_{\text{costs}}(C \setminus C_b, X)$. A broker b with a non-zero imbalance x_b that does not contribute to the overall imbalance (i.e., $x_b \cdot X \leq 0$) then “pays” $\frac{\text{BM}_{\text{costs}}(C \setminus C_b, X)}{X} x_b$. The payment for a broker b that contributes to the overall imbalance is defined the same in case there are no non-contributing brokers with controllable capacity. However, when such non-contributing brokers (denoted by B) do exist, we must make sure that the payment for contributing brokers such as b is sufficient to cover the payments for all non-contributing brokers. To guarantee this, we exclude all balancing orders of brokers B in computing the costs for contributing brokers, so broker b pays $\frac{\text{BM}_{\text{costs}}(C \setminus \{C_b \cup_{k \in B} C_k\}, X)}{X} x_b$.

In the (rare) case that the net imbalance of all brokers is exactly 0, brokers still need to pay something. In PowerTAC we use $-P^+ \cdot x_b$ in case $x_b < 0$ and $P^- \cdot x_b$ in case $x_b > 0$.

The VCG prices ensure, for a single time slot, that brokers cannot gain from pricing their orders higher or lower than their real costs, and that they can gain and never lose from placing orders for balancing capacity if they have any. In other words, brokers should bid their (estimated) actual costs for balancing capacity in the balancing market. With the second payment, if you expect other brokers to be (almost) balanced, it is better to be balanced as well; otherwise it is best to be imbalanced in the opposite direction from the overall imbalance.

The following example, graphically depicted in Figure 9, illustrates the balancing mechanism described above.

Example 1. Assume brokers $A_0, A_1, A_2,$ and A_3 have imbalances of 0, +40, -80, and -140 kWh, respectively, for a total imbalance $X = -180\text{kWh}$. We have six balancing orders bo_1 through bo_6 , and a dummy order RM . The total imbalance falls within the range of one of the orders, bo_5 .

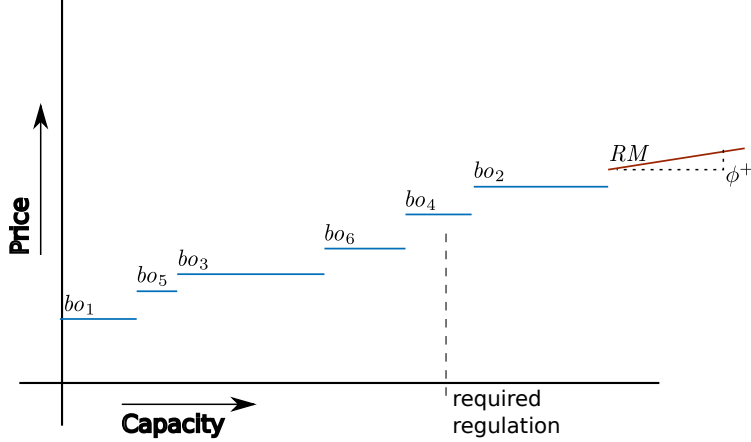


Figure 9: The balancing orders are ordered on price (increasing in case of a deficit). Only capacity up to the required regulation is used.

All orders with lower prices will be exercised, and bo_5 will be partially exercised. The signs in this example are from the standpoint of the brokers. This means that negative cash values represent payments from brokers to the BM, and negative energy values represent amounts the brokers have sold but not acquired, or amounts the brokers can consume by curtailing production.

The next step is to set prices for each broker's balancing orders, using the VCG mechanism. For each broker that has orders to be exercised, we must discover the price that would have to be paid for its capacity if its orders were not in the mix. To see how this works, assume the orders are as follows: bo_1 is (A_0 , 35 kWh, 0.003/kWh); bo_2 is (A_0 , 62 kWh, .0091/kWh); bo_3 is (A_1 , 67kWh, .0051/kWh), bo_4 is (A_1 , 30kWh, .008/kWh), bo_5 is (A_2 , 20kWh, .0042/kWh); bo_6 is (A_2 , 39kWh, .0062/kWh); and RM is (BM, xx kWh, .01/kWh, $\phi^+ = 0.001/kWh$). Sorted on the cost, we thus have the following balancing orders: bo_1 , bo_5 , bo_3 , bo_6 , bo_4 , bo_2 (See Figure 9). To balance we need all of bo_1 , bo_5 , bo_3 , bo_6 and only 19kWh of bo_4 .

The (VCG) pricing of the orders for broker A_1 can be found by removing bo_3 and bo_4 (i.e., 67 + 19 (out of 30) kWh), which requires the addition of all 62kWh of bo_2 and 24 kWh from order RM. The marginal cost of leaving broker A_1 's orders out is therefore $62 \cdot 0.0091 + 24 \cdot (0.01 + 0.024) = 1.3802$.

To compute the imbalance portion of the payment, we treat the contributing and non-contributing brokers separately. Since the overall balance is negative, the set of non-contributing brokers thus is just $\{A_1\}$. Contributing brokers A_2 and A_3 pay their shares of the cost to the BM to resolve the total imbalance, assuming that there are no balancing orders from the non-contributors. For each broker, the hypothetical total cost paid by the BM is the sum of the VCG payments to the other contributors, plus the residual amount the BM would have to purchase from the regulating market; the broker's share is the product of the total cost and the ratio of its imbalance to the total imbalance. For e.g. broker A_2 we first compute the costs for the BM in case bo_3 and bo_4 from A_1 and bo_5 and bo_6 from A_2 are removed. For this we compute the VCG payments for all other brokers (A_0 and A_4) and sum these (26.481 for A_0 since otherwise everything is to be resolved by the RM; A_4 does not have controllable loads), add the cost of the remaining 83 kWh at the marginal rate of 0.093, i.e., 7.719, divide by the overall imbalance $X = -180$, and multiply by the imbalance of A_2 of -80 ,

| broker | Imbalance | VCG payment | 2nd payment | total |
|--------|-----------|-------------|-------------|----------|
| 0 | 0 | 0.904 | 0 | 0.904 |
| 1 | +40 | 1.3802 | 5.0564 | 6.4366 |
| 2 | -80 | 0.5248 | -15.2 | -14.6752 |
| 3 | -140 | 0 | -17.6976 | -17.6976 |

Table 1: Broker payments for the example

giving a payment of 15.2. Broker A_1 is a non-contributing broker, so we remove bo_3 and bo_4 and compute the VCG payments for A_0 (15.035) and A_2 (6.903), and the BM cost for the extra 24 kWh from the regulating market (0.816) to get the hypothetical total BM cost (22.754) and A_1 's share of -5.056444.

All payments are summarised in Table 1.

The total budget for the BM, including all VCG payments and all secondary payments, in this case amounts to 25.0322.

7 Distribution utility

The Distribution Utility (DU) operates the distribution grid that connects the wholesale transmission grid to brokers and customers. It also acts as the “default broker,” offering basic and rather unattractive tariffs for all power types. The costs of the distribution utility must be covered by fees paid by brokers through distribution fees and capacity fees.

7.1 Distribution fees

The DU must recover the capital and O&M costs for the local distribution network. In the real world, billing for these costs are often called “meter charges” because they make up most of the fixed monthly charge paid by customers in regulated markets. The amount of the fee is often a function of the capacity of the customer’s connection to the grid, so a residential customer with 100-amp service will pay less than a commercial customer with 1000-amp service.

In Power TAC, the distribution fee will be a fixed charge per timeslot for each customer in a broker’s customer portfolio. There are two categories of customers: small and large. Small customers include households, EVs, offices, and distributed solar (Solar Leasing). Large customers include industrial and warehouse models, hospitals, and larger retail producers such as Windmill Coop and the Medical Center solar production.

The hourly meter charge m per customer will be configurable to allow for experimentation, but the value will be fixed for any given tournament. For small customers the hourly charge $.01 \leq m_{small} \leq .02$; for large customers $.03 \leq m_{large} \leq .10$.

7.2 Transmission capacity fees

In the real world, if we ignore taxes, the cost of delivered energy has two main components: the cost of generating it, and the cost of delivering it. Generation cost is commonly dominated by fuel prices, while delivery cost consists of capital and O&M (operations and maintenance) cost for the overall capacity to deliver it, including generation capacity and grid capacity. Of course,

the capacity to deliver energy must be sufficient to support the maximum demand for power that might be anticipated over some future horizon. In an environment where most energy is generated by fossil fuels, the fixed (capacity) and variable (fuel) costs of delivered energy are within an order of magnitude.

Much of the cost of capacity in modern grids is in the transmission infrastructure, which in the Power TAC scenario is behind the wholesale market. Since the need for capacity is driven by peak demand rather than mean demand, it makes sense that brokers should pay for their (customers') contribution to demand peaks.

It is common in the real world to assess costs for peak demand retrospectively; in year n , a distribution utility might be charged for their peak demand in year $n - 1$. The Power TAC scenario has two features that make this approach infeasible. First, the 2-month scenario is much too short for an annual assessment cycle, and second, customer portfolios are continually changing. Therefore, costs for transmission capacity will be assessed once every ν timeslots; the nominal assessment interval is $\nu = 168$ (one week). Each assessment will consist of three steps:

1. In each timeslot t chosen for assessment, compute the mean \bar{d}_t and standard deviation $\sigma_{d,t}$ of the net demand d over all timeslots back to the start of the bootstrap session. Compute a peak threshold $z_t = \bar{d}_t + \gamma\sigma_{d,t}$ where $1.0 \leq \gamma \leq 2.0$ is a parameter that will be fixed for any given tournament.
2. Find the n_{peak} highest demand events over the previous ν timeslots, where demand is defined as the difference between total customer consumption and total customer production.
3. For each identified peak $p > z_t$, a capacity charge ϕ_b will be assessed to each broker b in proportion to their customers' contributions to the peak. The total capacity charge Φ across all brokers at each assessment will be weighted by the amount by which the peak exceeds the threshold as $\Phi = \lambda(p - z_t)$, where $100 \leq \lambda \leq 1000$ is a parameter that will be fixed for any given tournament.

8 Competition format and interaction

As opposed to previous TAC competitions where the number of brokers were fixed in each game, in Power TAC the number of broker agents varies. This is expected to stimulate more dynamic agent design and a better abstraction of real-world conditions, simulating a range of competitive environments from duopolies to more competitive multi-broker oligopoly markets. Given the set of competing brokers in a given tournament, the Game Master will pick a few game-size values, and games will be scheduled such that each combination of registered brokers of each game-size are scheduled for the same number of games.

A typical Power TAC tournament is made up of one or more *rounds*, such as qualifying, seeding, semi-finals, finals. Depending on the number of participants N , each round consists of one or more *blocks*, each defined by the number of brokers B per *game* (or simulation) within the block. Each block in turn is composed of n sets of games, each set including every combination $\binom{B}{N}$ of N brokers taken B at a time.

8.1 Initialization and Default Broker

To create a fair start of each game, the simulation begins with all customers subscribed to the tariffs of the default broker, the marketing arm (such as it is) of the DU. These initial tariffs are intended to be fairly unattractive, so that customers will switch to more attractive tariffs very quickly once they are offered by the competing brokers.

A standard competition simulation begins after 15 days of simulation have already run with the default broker’s tariffs as the only available tariffs. Customer, market, and weather data from the last 14 days of this pre-game “bootstrap” period are collected and sent to brokers at the beginning of a game. More specifically, the bootstrap information includes:

Customer information: for each customer model, and for each power type supported by that model (such as solar production, consumption, interruptible consumption), the hourly energy consumption or production is given for each 1-hour time slot during the 14-day bootstrap data-collection period. Values are negative if the default broker is supplying the energy, positive if the customer is supplying energy.

Market information: for each time slot in the data-collection period, the total energy quantity purchased by the default broker in the wholesale market in MWh, along with the aggregated price/MWh.

Weather information: the weather reports for each time slot in the bootstrap data-collection period.

This data is intended to allow brokers to generate a reasonable initial model of the market in time to compose an initial set of tariff offerings as early in the simulation as possible.

In order to interpret the market prices in the bootstrap dataset, it is necessary to understand the bidding behavior of the default broker. The default broker estimates the net energy it needs to deliver to its customers by populating a vector for each of its customer subscriptions (each combination of customer and tariff) of size $7 \cdot 24$, or one cell for each time slot in a week. During the second through n th week, these cells contain the exponentially-smoothed ($\alpha = 0.3$) net consumption value for the customer in that time slot, counting from the start of a week. During the first week, it uses the actual consumption observed in the given hour h during the previous 24 hours, and during the first day it uses the usage observed in the previous time slot.

Given the default broker’s estimated net energy requirement (summed over all its models) for each of the following 24 time slots, it attempts to build a market position equal to its estimated need for that time slot. This is done by submitting an order for a quantity equal to the difference between its current position and its estimated need, with a limit price $l_{s,t}$ for an order placed at time t for energy in time slot s , except that if $s = t + 1$ (the last chance to purchase or sell energy for time slot s) then no limit price is given; the broker is willing to pay the market price. The limit price is bounded by minimum and maximum prices l_{min} and l_{max} , and computed as follows: First, a previous price is computed as

$$l_{prev} = \begin{cases} l_{s,t-1} & : \text{ if order in previous time slot } t - 1 \text{ did not clear} \\ l_{max} & : \text{ otherwise} \end{cases} \quad (15)$$

Then, given a random value ν in $[0, 1]$, the limit price is computed as

$$l_{s,t} = \max \left(l_{min}, 2 \frac{l_{min} - l_{prev}}{s - t - 1} \right) \quad (16)$$

The standard competition parameters can be found in Table 2. Values for these parameters are sent to a broker at the start of every game. For details see the software documentation.

Table 2: Parameters used in Power TAC tournament games.

| Parameter | Symbol | Standard Game Setting |
|---|--------------------------------|----------------------------|
| Length of pre-game bootstrap period | | 14 days |
| Nominal length of game | E | 60 days |
| Probability of game end for each time slot after time slot 1320 (start of day 55) | p | $\frac{1}{121}$ |
| Minimum game length | Min(TS) | 1320 |
| Expected game length | E(TS) | 1440 |
| Timeslot length | τ | 60 minutes |
| Time compression ratio | ρ | 720 (5 seconds/time slot) |
| Open time slots on wholesale market | | 24 |
| Market closing time | | 1 time slot ahead |
| Minimum order quantity | e_{min} | 0.1 kWh |
| Wholesale scale factor | ψ | 10 |
| Meter charge, small customer | m_{small} | [0.01 - 0.02]€/timeslot |
| Meter charge, large customer | m_{large} | [0.03 - 0.10]€/timeslot |
| Demand assessment interval | ν | 168 hours |
| Peak demand threshold coefficient | γ | [1.0 - 2.0] |
| Peak demand charge | λ | [100 - 1000]€/point |
| Balancing cost | c_0 | [0.02 - 0.06]€/kWh |
| Slope of regulating market price | ϕ^+, ϕ^- | $10^{-6}, 10^{-6}$ €/kWh |
| Default broker's min and max bid order prices | $l_{min}(bid), l_{max}(bid)$ | -100, -5 |
| Default broker's min and max ask order prices | $l_{min}(ask), l_{max}(ask)$ | 0.1, 30 |
| Tariff publication fee | | [1000 - 5000] € |
| Tariff revocation fee | | [100 - 500] € |
| Tariff publication interval | | 6 time slots |
| Daily bank debt interest rate | $[\beta_{min}, \beta_{max}]$ | 4.0%/365 \dots 12.0%/365 |
| Daily bank deposit interest rate | $[\beta'_{min}, \beta'_{max}]$ | 0.5β |
| Weather report interval | | 1 hour |
| Weather forecast interval | | 1 hour |
| Weather forecast horizon | | 24 hours |

8.2 Simulation duration

The game ends after a random number K time slots after day 55 (time slot 1320), $K = 0, 1, \dots$. For each time slot, starting at the start of day 56, there is a fixed probability p that the game ends after that particular time slot. As a consequence, the number of time slots in excess of day 55, K , follows a geometric distribution. The expected length of a standard tournament game is 1440

timeslots.

Given the random end of game and that each Power TAC day lasts 120 seconds in real time, the expected duration of a Power TAC tournament game is around 2 hours overall.

8.3 External metrics and game logs

In order to allow games to be followed in real time, and also analyzed in depth at a later date, the simulator generates a *state log* while it runs, from which the entire simulation can be reproduced. From this data, additional metrics, including the following, can be monitored during or after a game. These metrics are used by the game viewer to provide a visual representation of the game as it proceeds.

- Bank balance for each broker
- Balancing performance for each broker
- All tariff offers and orders exchanged by brokers and customers
- Portfolio of each broker

8.4 Winner determination

At the conclusion of a competition, the relative performance of the participants must be evaluated and compared. This is accomplished by rank ordering all participants according to one or more defined performance criteria and declaring the best performer in this rank order to be the winner of the competition. This principle also applies to Power TAC; albeit with significant differences compared to previous TAC competitions. Here we describe the performance criteria used to rank order the Power TAC participants. Note that a wide range of performance criteria, such as minimizing carbon emissions, maximizing the share of renewable energy, and other factors can be converted to monetary units by introducing taxes and incentives as part of the market structure.

8.4.1 Performance criteria

For each broker b participating in game g with a block c of a given tournament round, a profit $\pi_{b,c,g}$ is calculated as the (monetary) payments $pay_{b,c,g}$ minus costs $cost_{b,c,g}$ minus fees $fee_{b,c,g}$:

$$\pi_{b,c,g} = pay_{b,c,g} - cost_{b,c,g} - fee_{b,c,g} \quad (17)$$

- **Payments** are monetary transfers from customers (consumer) to brokers and are based on agreed contract conditions (including payments for exercising storage capacity) and the actual (ex-post) measured energy consumptions of the respective customer (consumer) after curtailments are exercised. Other payments for instance include sales in the wholesale market, and possible payments from external balancing.
- **Costs** are monetary transfers from brokers to customers (producers) and are based on the agreed contract conditions between the respective customer (producer) and broker and the actual (ex-post measured) energy produced after curtailments are exercised. Other costs for instance include procurement in the wholesale market.

- **Fees** are (i) the cost for external balancing energy (see Section 6) used, and (ii) energy distribution fees (in €/kWh) or capacity fees (in €/kW) levied by the DU for energy delivered to customers. Other fees include costs for publishing or revoking tariffs.

8.4.2 Final ranking algorithm

After all blocks within a round have run, e.g. at the end of a finals round, z -scores of the accumulated profits for each broker are calculated to facilitate comparisons among games with different numbers of brokers. If we denote the accumulated profits of a broker in a game as π_{bc} , the average accumulated profits of all brokers in the game as $\bar{\pi}_c$ and the standard deviation of all brokers in the game as S_c , then the standardized accumulated profits of broker b in block c , z_{bc} , is obtained as:

$$z_{b,c} = \frac{\pi_{b,c} - \bar{\pi}_c}{S_c}, \quad (18)$$

where

$$\pi_{b,c} = \sum_{g=1}^{N_{b,c}} \pi_{b,c,g}, \quad (19)$$

where $N_{b,c}$ is the number of games broker b played during block c .

At the end of a round, after all blocks C have ended, an overall measure of relative broker performance will be obtained by summing over the standardized broker performance per block:

$$z_b = \sum_{c=1}^C z_{bc} \quad (20)$$

where C is the number of blocks in the round.

8.5 Tournament structure

A typical Power TAC tournament consists of several rounds. Note that seeding, quarter-final, and semi-final rounds will be included in a tournament schedule only if the number of participants is clearly too large for a reasonable final-round schedule. It will be the responsibility of the Game Master to make this determination.

Qualification Round A chance for each team to test their broker against brokers from other teams in a real competition environment. This is mainly done to check overall functionality of a broker and its communication with the competition server. A broker may be eliminated at the end of a qualification round for several reasons:

- it has not demonstrated an ability to function consistently and correctly by the end of the round;
- its behavior results in severe market distortions, such as underpricing its tariffs and consistently losing money in such a way that other brokers are shut out, although exploiting flaws in the design or implementation of the game is not grounds for disqualification;
- it fails to participate in any meaningful way, even though it stays connected.

A typical qualification round will run continuously for several days, around the clock.

Seeding Round This round will result in a ranking that is used to determine broker pools for quarter final rounds, in case the number of entrants needs to be trimmed before the final round. As with a qualification round, a seeding round might result in elimination of a broker that does not perform according to the game specification or appears to be purposely disruptive to other brokers. As with a qualification round, seeding rounds will typically run 24 hours/day for several days.

Quarter Finals This is the first real elimination round, since only half of the teams will proceed to the semi finals.

Semi Finals Elimination round; only half of the teams will proceed to the finals.

Final The winner of this round wins the overall specific competition.

8.6 Competition rules

In the following list we highlight the competition rules that each participant team has to follow; failure to do so will lead to disqualification from the overall tournament. The decision rests with the current game master.

- Much of the information in the game logs is private to individual brokers during a game, and is not provided to other brokers. Brokers must not attempt to access it through external means (i.e. through the game viewer or the server logs). The use of such external information, either manually or automatically, is regarded as external ‘tuning’ of the broker. As such, according to the existing competition rules, it is forbidden within a game, and within quarter-final, semi-final, and final rounds of a tournament. Tuning with any available data, including game logs, is allowed between tournament rounds.
- Data that brokers discover on their own during a game can be used to fine-tune their behaviors in games within a tournament round.
- Collusion is not allowed between the different brokers.
- To discourage anti-competitive collusion, no team is allowed to enter the competition with two different broker identities. Multiple teams from the same institution may be allowed, but only if they are clearly separate and committed to winning without collusion.
- For efficient tournament scheduling, each team must have at least two copies of its broker either running or waiting for its next game at all times during a tournament round, since brokers are required to participate in different pools at the same time.

9 System architecture

9.1 Tournament deployment

Power TAC is designed to run as an annual competition, a model that has been very effective in stimulating research. Each year, research groups build or update their brokers and enter them in the competition. The competition systems architecture is shown in Figure 10.

The tournament configuration is intended to support multi-round tournaments, with large numbers of spectators. The administration portion of the web application supports tournament scheduling, broker registration, and access to records of past games. The web application also manages a set of visualizers giving access to running games on potentially several simulation servers.




Figure 10: Competition systems architecture.

A single web app can control multiple servers on multiple hosts, along with associated visualizer modules that support scalable browser-based observation of games. Weather data will be served by a remote service, hosted on its own database. The tournament database holds summary information for completed games, including access information for retrieving game logs.

Brokers register with the web app, and join a game by requesting credentials and a URL for an active simulation. With this information, it then logs into the simulation server and runs its game interactions.

9.2 Research deployment

After the competition, teams are encouraged to release their agent code (in either binary or source form), so all teams can design and run their own experiments using a range of broker behaviors and market design details. Teams publish results, incorporate new insights into their broker designs for the following year.

The Power TAC simulation, with or without a copy of the game visualizer, can also be deployed in a research or software development situation, as shown in Figure 11. The goal of the research configuration is to support development of brokers and server models (customers, markets, etc.) and to support empirical research. In this configuration, the server must be easily deployable on a desktop workstation, without requiring special privileges, and with minimal dependencies on other installed software, such as a database. In addition, this configuration meets the following requirements:

- Single-simulation setup from a simple web interface.
- Optionally allow broker login without credentials.

- Visualizer support for at least one browser.

In the research configuration, the simulation server is identical to the tournament version, and a portion of the web app is installed in the server. Through the web interface, a user can configure and start a game, and use the visualizer to watch the game. Weather and price data may be contained in flat files, or a research server could potentially access weather data from a tournament installation. Game data is dumped to flat files at the conclusion of each game, and a log-analysis tool is available for parsing it and extracting whatever data may be of interest to answer specific research questions.

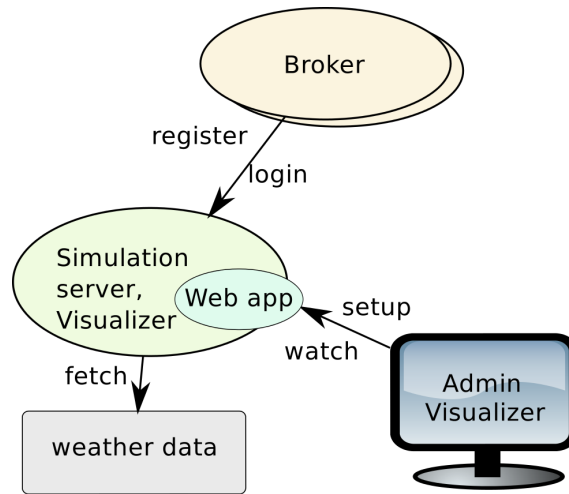


Figure 11: Research systems architecture.

Each year, the simulation may be updated to add new challenges, and if necessary to tune the market designs and level of realism to enhance the relevance of the shared enterprise for both research value and policy guidance.

References

- [1] U. S. E. I. Administration. *Annual Energy Review 2010*. Oct. 2011.
- [2] M. Amin and B. Wollenberg. Toward a smart grid: Power delivery for the 21st century. *IEEE Power & Energy Magazine*, 3(5):34–41, 2005.
- [3] M. Bichler, A. Gupta, and W. Ketter. Designing Smart Markets. *Information Systems Research*, 21(4):688–699, 2010.
- [4] S. Borenstein, J. B. Bushnell, and F. A. Wolak. Measuring market inefficiencies in California’s restructured wholesale electricity market. *The American Economic Review*, 92(5):1376–1405, 2002.
- [5] J. Casazza and F. Delea. *Understanding electric power systems: an overview of the technology and the marketplace*, volume 13. Wiley-IEEE Press, 2 edition, 2003.
- [6] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

- [7] J. Collins, W. Ketter, and N. Sadeh. Pushing the limits of rational agents: the Trading Agent Competition for Supply Chain Management. *AI Magazine*, 31(2):63–80, 2010.
- [8] M. de Weerd, W. Ketter, and J. Collins. A theoretical analysis of pricing mechanisms and broker’s decisions for real-time balancing in sustainable regional electricity markets. In *Conference on Information Systems and Technology*, pages 1–17, Charlotte, Nov. 2011.
- [9] S. Gottwalt, W. Ketter, C. Block, J. Collins, and C. Weinhardt. Demand Side Management - A simulation of household behavior under variable prices. *Energy Policy*, 39:8163–8174, 2011.
- [10] P. Jordan, B. Cassell, L. Callender, and M. Wellman. The ad auctions game for the 2009 trading agent competition. Technical, University of Michigan, 2009.
- [11] P. Joskow and J. Tirole. Retail electricity competition. *The Rand Journal of Economics*, 37(4):799–815, 2006.
- [12] W. Ketter, J. Collins, and P. Reddy. Power TAC: A Competitive Economic Simulation of the Smart Grid. *Energy Economics*, 39:262–270, 2013.
- [13] W. Ketter, M. Peters, J. Collins, and A. Gupta. Competitive Benchmarking: An IS Research Approach to Address Wicked Problems with Big Data and Analytics. *MIS Quarterly*, 40(4):1057–1080, 2016.
- [14] W. Ketter, M. Peters, J. Collins, and A. Gupta. A Multiagent Competitive Gaming Platform to Address Societal Challenges. *MIS Quarterly*, 40(2):447–460, 2016.
- [15] D. L. Mills. *Computer network time synchronization : the network time protocol*. CRC/Taylor & Francis, Boca Raton, 2006.
- [16] C. Möller, S. Rachev, and F. Fabozzi. Balancing energy strategies in electricity portfolio management. *Energy Economics*, 33(1):2–11, 2011.
- [17] P. Siano. Demand response and smart grids: A survey. *Renewable and Sustainable Energy Reviews*, 30:461–478, Feb. 2014.
- [18] A. Somani and L. Tesfatsion. An Agent-Based Test Bed Study of Wholesale Power Market Performance Measures. *IEEE Computational Intelligence Magazine*, 3(4):56–72, Nov. 2008.
- [19] L. Vandezande, L. Meeus, R. Belmans, M. Saguan, and J.-M. Glachant. Well-functioning balancing markets: A prerequisite for wind power integration. *Energy Policy*, 38(7):3146–3154, 2010.
- [20] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- [21] P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24(1):17–27, 1998.

A Assumptions

In particular we make the following *assumptions*:

1. Within the simulated region, grid constraints (line capacity limitations) are assumed to be non-existent, i.e. energy flows within the region are unconstrained. Local distribution grids are typically over-dimensioned with respect to their line capacities, thus this assumption is not a strong restriction but may have to be rethought in future once much more distributed generators and storage facilities are installed.
2. Power factor effects, i.e. phase shifts between voltage and current, are not taken into account. Modeling these effects would possibly influence the brokers' decision making on which consumers and producers to add to their portfolios but is out of scope at this time.
3. Energy distribution and transformation losses are ignored. In Germany these losses are estimated at 3%; for North America they are estimated at 5.5% [1]. These losses can be considered as being more or less constant within a distribution grid and identical for all grid participants. Thus the validity of the simulation results is not affected.
4. Two kinds of producers (energy production facilities) are distinguished. One kind (photovoltaic arrays, wind turbines) produce energy when active, and are under control of their respective owners. The second kind (PEV batteries, some CHP units) is called "controllable" and may be switched on or off, or have its output adjusted remotely within its capacity range.
5. Technical load balancing (i.e. the real-time operations of the local distribution grid) is accomplished outside the action domain of the competition participants using a combination of controllable generators and spinning reserves.
6. The simulation will model time as a series of discrete "time slots" rather than as continuous time. This models the trading intervals in the regional wholesale market, and enables the simulation to model a period of days rather than minutes or hours.
7. The temporal distribution of energy consumption and generation *within* a time slot is not taken into account. This means for example that balancing energy demand for a time slot is calculated as the difference of the sum of generation and the sum of consumption for that time slot and not as the instantaneous difference between the two timeseries.
8. Some portion of the load, including the charging and discharging of plug-in Electric Vehicles (PEVs), could be controlled by voluntary or automated means, using prospective or real-time price signals.

B Acknowledgements

The authors wish to acknowledge the valuable contributions of many people who have contributed to the design of the Power TAC scenario and the development of the software infrastructure that supports it.

Carsten Block was one of the initial designers of the competition scenario (besides Wolfgang Ketter and John Collins), and of the software infrastructure.

We would like to thank further the members of the TAC, CS, and IS community for their valuable feedback on our design, and for the great help with the development of the Power TAC platform.

Special thanks go to Jurica Babic, Tomislav Brisevac, Govert Buijs, Antonios Chrysopoulos, Jan van Dalen, Travis Daudelin, David Dauer, Josh Edeen, Niels Egberts, Christoph Flath, Jeroen Ruigrok, Ryan Finneman, Sebastian Gottwalt, Puthyrak Kang, Yixin Lu, Adis Mustedanagic, Nguyen Nguyen, Erik Onarheim, Shashank Pande, Martin Paschmann, Markus Peters, Vedran Podobnik, Prashant Reddy, Daniel Schnurr, Andreas Symeonidis, Konstantina Valogianni, and Manuela Veloso.