

The Power Trading Agent Competition

Wolfgang Ketter, John Collins, Prashant Reddy,
Christoph Flath, and Mathijs de Weerd

ERIM REPORT SERIES <i>RESEARCH IN MANAGEMENT</i>	
ERIM Report Series reference number	ERS-2011-027-LIS
Publication	December 2011
Number of pages	46
Persistent paper URL	http://hdl.handle.net/1765/30683
Email address corresponding author	wketter@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES *RESEARCH IN MANAGEMENT*

ABSTRACT AND KEYWORDS	
Abstract	<p>This is the specification for the Power Trading Agent Competition for 2012 (Power TAC 2012). 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. 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 timeslot.</p> <p>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 modelling a single region, and therefore we do not model location-marginal pricing. Customer models include households and a variety of commercial and industrial entities, many of which have production capacity (such as solar panels or wind turbines) as well as electric vehicles. 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 and for real-time balancing of supply and demand. The balancing process is 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. The broker with the highest bank balance at the end of the simulation wins.</p>
Free Keywords	autonomous agents, electronic commerce, energy, preferences, portfolio management, power, policy guidance, sustainability, trading agent competition
Availability	<p>The ERIM Report Series is distributed through the following platforms:</p> <p>Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal</p> <p>Social Science Research Network (SSRN), SSRN ERIM Series Webpage</p> <p>Research Papers in Economics (REPEC), REPEC ERIM Series Webpage</p>
Classifications	<p>The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems:</p> <p>Library of Congress Classification, (LCC) LCC Webpage</p> <p>Journal of Economic Literature, (JEL), JEL Webpage</p> <p>ACM Computing Classification System CCS Webpage</p> <p>Inspec Classification scheme (ICS), ICS Webpage</p>



The Power Trading Agent Competition

**Wolfgang Ketter, John Collins^a,
Prashant Reddy^b, Christoph Flath^c, Mathijs de Weerd^d**

December 2011
ERS-2011-027-LIS

^aComputer Science Department, University of Minnesota, USA

^bMachine Learning Department, Carnegie Mellon University, USA

^cKarlsruhe Institute of Technology, Germany

^dDelft University of Technology, Netherlands

Department of Decision and Information Sciences
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 2012 (Power TAC 2012). 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. 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 timeslot.

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 do not model location-marginal pricing. Customer models include households and a variety of commercial and industrial entities, many of which have production capacity (such as solar panels or wind turbines) as well as electric vehicles. 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 and for real-time balancing of supply and demand. The balancing process is 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. The broker with the highest bank balance at the end of the simulation wins.

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	5
2.5	Accounting	5
2.6	Weather reports	6
3	Brokers	6
3.1	Actions available to brokers	6
3.1.1	Design, offer and modify tariffs	6
3.1.2	Dynamic pricing decisions	8
3.1.3	Wholesale market trading	9
3.1.4	Portfolio management	10
3.2	Information available to brokers	11
4	Customer market	12
4.1	Customer types	13
4.2	Tariff market interaction	13
4.2.1	Derive tariff utility	14
4.2.2	Choose from a list of tariffs	15
4.3	Provide balancing capacity	15
4.4	Consume and produce power	16
4.5	Available customer models	17
5	Wholesale market	18
5.1	Trading and timeslots available for trade	18
5.2	Market clearing	18
5.3	Wholesale suppliers and buyers	20
6	Balancing market	20
6.1	Adjusting energy demand and supply	21
6.2	Market-based balancing mechanisms	24
6.2.1	Desirable balancing mechanism properties	25
6.2.2	Scenario I: no controllable capacities	26
6.2.3	Scenario II: static with controllable capacities	27
6.2.4	Scenario III: dynamic with controllable capacities	29
7	Competition format and interaction	29
7.1	Competition initialization and Default Broker	29
7.2	Competition ending	30
7.3	External metrics and game logs	32
7.4	Winner determination	32

7.4.1	Performance criteria	32
7.4.2	Final ranking algorithm	33
7.4.3	Tournament structure	33
7.5	Competition rules	34
8	System architecture	34
8.1	Tournament deployment	34
8.2	Research deployment	35
A	Assumptions	39
B	Acknowledgements	40

1 Background and motivation

We know how to build “smart grid” [1] components that can record energy usage in real time and help consumers better manage their energy usage. However, this is only the technical foundation. Variable energy prices that truly reflect energy scarcity can motivate consumers to shift their loads to minimize cost, and for producers to better dispatch their capacities [14]. This 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 power. Unfortunately, serious market breakdowns such as the California energy crisis in 2000 [3] 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 [21], 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 [12], 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 regions¹. Therefore, we are presenting 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. 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 [2]. The *Power Trading Agent Competition* is an example of a Trading Agent Competition (TAC)² applied to energy markets. Earlier successful examples of TAC include the Trading Agent Competition for Supply-Chain Management (TAC SCM) [7] and the Trading Agent Competition for Ad Auctions (TAC AA) [13].

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. Brokers compete with each other trying to attract customers by offering *tariff* contracts to a population of anonymous small customers (households, small businesses), and by negotiating individual contracts with larger customers (such as major manufacturing facilities, or greenhouse complexes with many Combined Heat and Power (CHP) units). Contract terms may include fixed or varying prices for both consumption and production of energy, along with other

¹See [15] for a complete overview of related work.

²See <http://www.tradingagents.org>

incentives such as rebates for energy conservation, or even sign-up bonuses or 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.

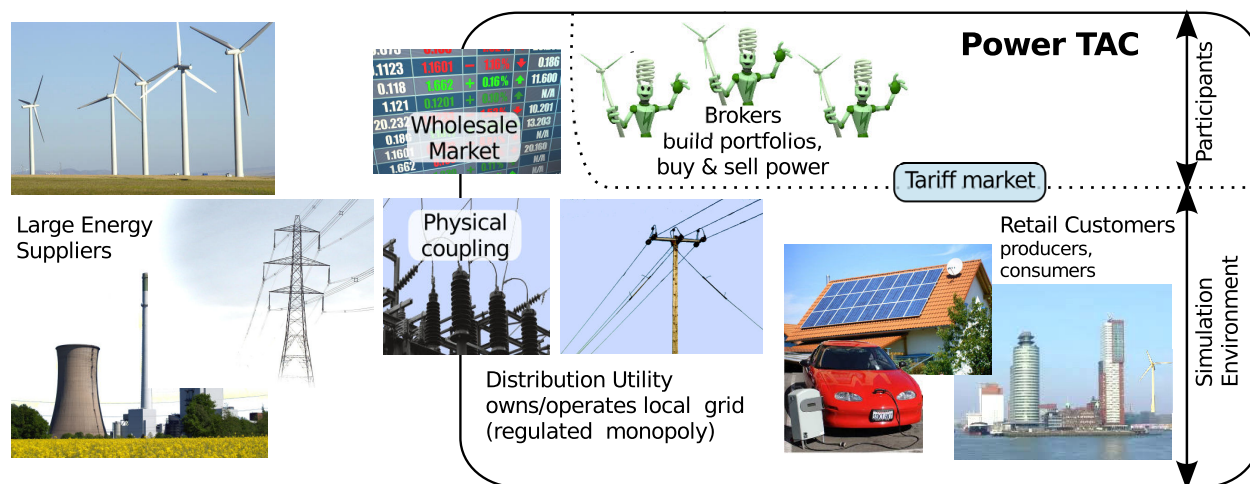


Figure 1: Major elements of the Power TAC scenario.

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). In the future, we anticipate integrating the 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 two markets, a tariff market and a wholesale market. Over a planning horizon from weeks to months, brokers build portfolios of consumer, producer and electric vehicle customers by offering tariff contracts and negotiating individual contracts³. At the operational level, over a time horizon of 24 hours, brokers must balance the fluctuating energy demands of their contracted power 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, and by purchasing or selling energy in the wholesale energy market. Retail market dynamics thus influence the wholesale market and vice versa.

A broker’s primary goal in portfolio development (see Figure 2) 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 agent will benefit from having reasonably-priced energy sources that can be expected to produce power when demand is expected to be highest within its load portfolio.

³Individual contract negotiation will be implemented for the 2012 competition.

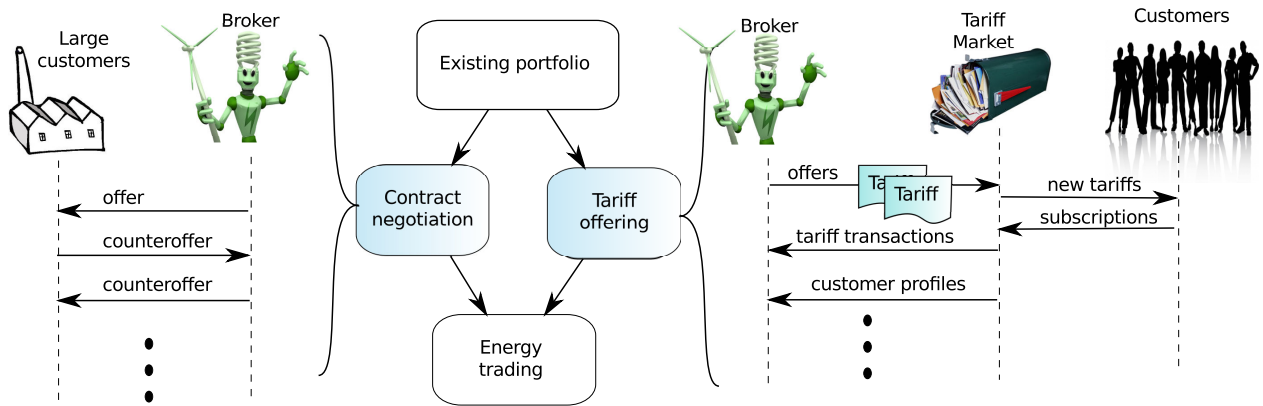


Figure 2: Portfolio management process. Tariff offerings proceed in parallel with individual contract negotiation.

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.

2.1 Simulation time

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

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 can buy and sell energy through two different mechanisms, *tariffs* and *individual contracts* (although individual contracts will likely not be implemented for the 2012 competition). For most customers, such as households, small businesses, and small energy producers, brokers may offer tariffs that specify pricing and other terms, and customers must choose among the tariffs on offer. For larger producers or consumers that do not interact directly with the wholesale markets (for example, a large industrial facility, a university campus, or a greenhouse complex with many CHP units), brokers may negotiate individual contracts. Tariff offerings and contract negotiations may be conducted at any time, without regard to the daily and hourly cycle of the simulation, as depicted in Figure 2. However, tariffs will be published to retail customers in batches, nominally 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.

Contract and tariff terms and conditions must be described in a language that has clear semantics along with the necessary features to describe a variety of possible business agreements between brokers and their customers. The development of a common semantic model and a common pricing model to describe various kind of energy tariffs are considered top priorities on the EPRI / NIST Smart Grid roadmap for the development of a smart grid [26]. With no common standard in place to build on for Power TAC, we use with the work of Tamma et al. [23], an ontology that describes a negotiation process including (i) the involved parties, (ii) the object to negotiate on, and (ii) the negotiation process, i.e. the economic mechanism itself.

Within the Power TAC domain, negotiations and the contracts (including tariffs) that are the subject and result of negotiations must be able to specify

Time: including points in time, time intervals, periodicity (days, weeks, months, etc.), and temporal relationships (before, after, during, etc.). 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). Some contracts or tariffs will also need to specify amounts of energy that can be remotely controlled (interrupted), for example by shutting off a domestic water heater for 15 minutes every hour during peak demand periods. Such remotely-controllable sources or loads are called “balancing capacity.”

Money: Agreements must 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, etc.

A broker must use tariff offerings and contract negotiations to develop a portfolio of contracted consumers and producers. To do this, brokers will need to estimate and reason about consumer and producer preferences in order to design appropriate tariffs and to appropriately respond to counteroffers from potential contract customers. Brokers will also need to estimate future consumer and producer behavior to build a portfolio that has well-balanced demand and supply over time and that provides sufficient balancing capacity to achieve an acceptably low risk of imbalance.

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, clearing once every simulated hour. Participants include the brokers and a set of wholesale participants that provide bulk power and liquidity to the market.

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 power from the transmission grid to the 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. Brokers must pay distribution fees for the use of the distribution grid in proportion to the quantities of energy their customers transport over the grid.
2. It is responsible for the real-time balance of supply and demand on the distribution grid. In this role it operates a “balancing market” (see Section 6 that creates incentive for brokers to balance their own portfolios of energy supply and demand in each timeslot.
3. It offers “default” tariffs for energy consumption and production. In this role it simulates the electric utility in a non-competitive regulated tariff 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 power, 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 Accounting

Cash accounting aggregates customer transactions for tariff subscription and withdrawal, and power consumption and production. Other transactions include tariff publication fees, market settlements, interest on debt, and credits and debits related to taxes and incentives. Market position accounting tracks the current commitments in the wholesale market for each broker in each future timeslot. This information is needed by the Distribution Utility to run the balancing process in the current timeslot.

Each agent has an 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 timeslot. Agents are allowed to carry a negative balance during the course of the game.

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

$$b_{d+1} = (1 + \beta/365)b_d + \text{credits}_d - \text{debits}_d \quad (1)$$

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

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

$$b_{d+1} = (1 + \beta'/365)b_d + \text{credits}_d - \text{debits}_d \quad (2)$$

Typical annual savings interest is $\beta' = 5\%$.

Values for β and β' are provided to the agent at the beginning of the game (see Table 1 on page 31 for standard tournament values). An updated cash position report is the last message sent by the simulation server to the broker in each timeslot.

2.6 Weather reports

Weather forecasts and current-hour weather conditions are sent to brokers in each timeslot. Some customer models will use this information to influence energy consumption (temperature, for example), and production (wind speed, cloud cover). Brokers who have subscribed customers that are weather-sensitive will also need this data to predict production and consumption. In most cases, this component will be a proxy for an external data source containing real-world weather and forecast history data for some real-world location. The location and date range for the weather dataset is not revealed to brokers.

3 Brokers

3.1 Actions available to brokers

Figure 3 provides an overview of the timeline and information exchange between a broker and the simulation environment in each timeslot. Note that the specific order of events is more flexible than what is shown. Specifically, the sequence of major processes in the simulation environment is fixed (additional detail is given in Figure 6), but brokers can send messages at any time, as long as they arrive before the server needs them.

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

Create new tariffs (Tariff Market): Design and offer new tariffs to customers.

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

Price adjustments (Customers): Adjust prices in a current tariff, if tariff terms allow it.

Contract negotiation (large Customers): Participate in bilateral negotiation to define individual contracts (not implemented in the current version).

Balancing offer (Distribution Utility): Offer controllable capacities for real-time balancing, to the extent allowed by tariff terms.

Create asks and bids (Wholesale Market): Create asks and bids to sell or procure energy for future timeslots.

We now describe each of these activities in more detail.

3.1.1 Design, offer and modify tariffs

To manage their portfolios, brokers design and offer tariffs. They may also modify an existing tariff by superseding it with a new one, then revoking the original tariff. The detailed structure of a tariff offering is shown in Figure 4. This structure supports a number of features within a simple, compact object graph. Many concepts are represented in the `TariffSpecification` itself (payments, energy-type), but the rate structure is broken out. This allows for a range of rate structures without requiring space (memory and bandwidth) for unused features. It also allows a simple convention of empty references for unused features. Here are some common tariff features that can be represented with this structure:




Figure 3: Overview of Power TAC activities within one timeslot. A broker interacts with the wholesale and tariff markets, and receives information from the weather service, customers, the balancing market, and the accounting service.

- tiered rates, in which customers pay/receive one rate for a portion of usage (up to 20 kWh/day, for example), and a different rate for the remainder;
- time-of-use rates;
- weekday/weekend rates;
- two-part tariffs (fixed daily fee plus usage fee);
- signup payments in either direction (fee or bonus);
- early withdrawal penalties;

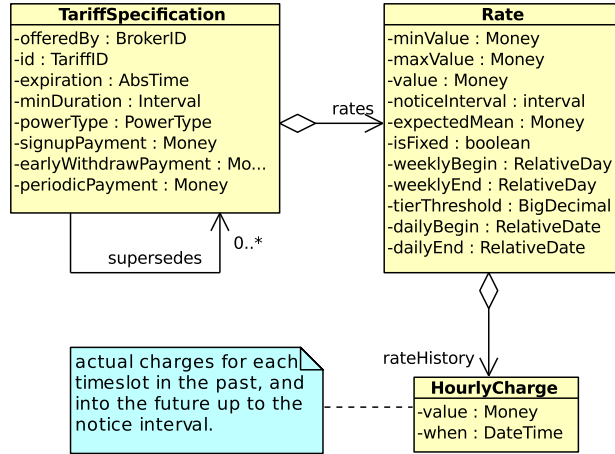


Figure 4: Tariff structure.

- variable rates with minimum and maximum values, estimated mean values, and notice intervals.

It is not currently possible to write tariffs that bundle multiple power-types, such as household consumption and electric-vehicle charging. Such bundling is certainly practiced in the real world, but for the time being, the complexity of evaluating bundled tariffs is avoided. On the other hand, bundling of tariff instances within the scope of a negotiated agreement seems reasonable and easily represented with minor modifications.

Figure 5 shows the evolution of a single tariff from the time it is published. Brokers can submit tariffs to the market at any time (*pending*). Periodically new tariffs are published by the market to customers and to all brokers, at which point they are *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. If a broker wishes to modify an existing tariff, the process is to first offer a new tariff that *supersedes* the existing tariff, and then force customers to unsubscribe from the existing tariff by *revoking* it. As long as some other tariff has already been submitted that *supersedes* the revoked tariff, then all subscriptions are automatically transferred to the superseding tariff, but with a minimum contract duration of 0. If there is no superseding tariff, then customers are forced back to the default tariff.

3.1.2 Dynamic pricing decisions

An important tool in a broker’s ability to balance consumption and production from its portfolio of customers and wholesale market commitments is the ability to change prices for customers dynamically using variable-rate tariffs. Since such dynamic prices are typically communicated to the customers some number of timeslots before the timeslot to which they apply, the broker must use some type of forecasting to determine the optimal price to set for the target timeslot, i.e., the future timeslot for which it is now required to communicate prices.

There are several environmental features that factor into the prices that the broker may want to

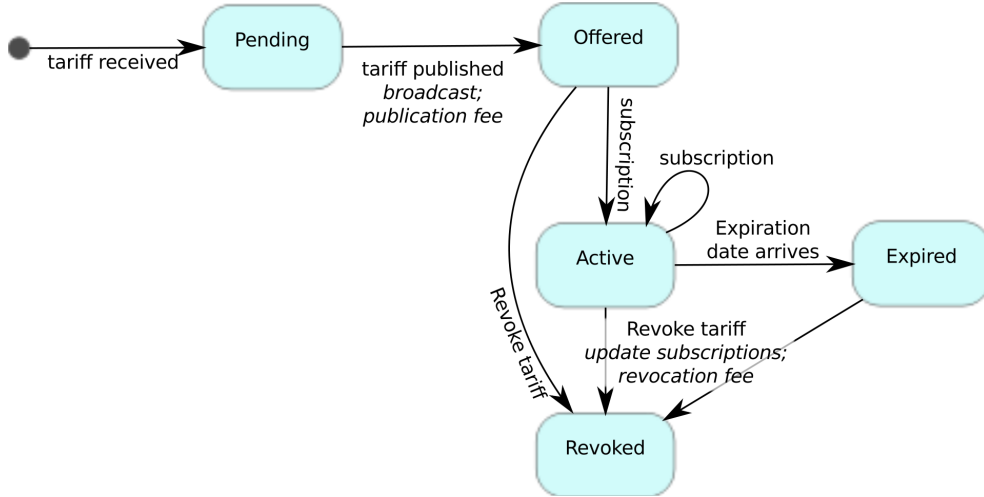


Figure 5: Tariff state transitions.

charge. At a basic level, a broker typically already knows something about the price of power 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. Two major factors in the determination of this demand and supply are (i) the estimated or realized load and supply for timeslots preceding the target timeslot, and (ii) the weather forecast conditions for the target timeslot.

At a more advanced level, a broker can also try to forecast the prices in the wholesale market as well as the DU’s balancing market and use those forecasts in setting its tariff prices for the target timeslot. For example, if the broker believes that it will likely be cheaper to buy energy in the wholesale market than to increase production from its portfolio, it may choose to not increase its dynamic tariff prices for producers, which would normally incentivize them to increase production, even when it needs to respond to a potential short-supply condition in the target timeslot.

3.1.3 Wholesale market trading

Dynamic adjustment of prices for consumers and producers who are on variable-price tariffs and the advance reservation of interruptible capacity as balancing power are two possibilities to balance a broker’s portfolio over time. The third is to buy missing, or to sell excess, capacity on the wholesale market. Details of the wholesale market clearing process are given in Section 5. In Figure 6 we see in more detail the timing of interactions between the broker and the wholesale market, along with the information needed by brokers to make trading decisions.

The wholesale market is cleared at the beginning a timeslot n . The process starts with an announcement of the timeslots open for trading in the following timeslot, typically timeslots $[n + 1 \dots n + 24]$. Next, all outstanding orders that have been submitted since the beginning of the previous timeslot $n - 1$ are cleared, and the results announced in the form of cleared trades (amounts and prices) and orderbooks (uncleared bids and asks) for each cleared timeslot. From the broker’s perspective, the information it needs to make trading decisions for future timeslots starts at the beginning of a timeslot. This information includes weather reports, customer usage and production reports, balancing transactions, tariff subscription changes, transactions, and updates to its current market and cash positions. Assuming reasonable network performance, all this information will

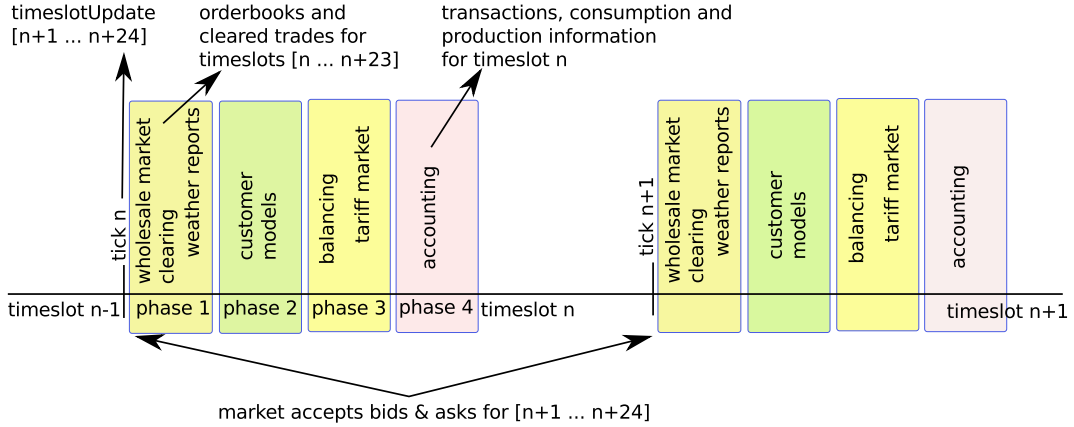


Figure 6: Simulation process phases and associated market information.

arrive in time to make final trading decisions for the following market clearing.

3.1.4 Portfolio management

The primary goal of a broker is to publish tariffs and negotiate contracts for power sources and loads that result in a portfolio that is profitable and balanced, at least in expectation, over some period of upcoming execution activities and timeslots. For example, an agent will benefit from having reasonably-priced energy sources that can be expected to produce power 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) and with a balanced portfolio of uncorrelated power 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 interruptible or controllable consumption and production capacity that can be used as needed (balancing capacity), and by trading futures contracts on the wholesale market.

Power sources include cleared bids in the wholesale market, small local producers (household and small-business sources) acquired by offering tariffs, and large local producers (e.g., small wind farms or CHP plants) acquired through individually negotiated contracts.

Power sources can be more or less predictable, and may have a non-zero controllable component as discussed in Section 2. Predictable sources include power obtained from the wholesale market as well as the continuous portion of the output from many CHP and hydro plants. Less predictable sources include most renewable sources such as wind and solar plants, which fluctuate with weather conditions and/or time of day.

Loads include cleared asks in the wholesale market, small local loads (e.g., households and small businesses) acquired by offering tariffs, and large local loads (e.g., industrial facilities and large office parks) acquired through individual contracts.

Storage capacity can be used to absorb excess power or to source power during times of shortage. Power can be absorbed by capacity that is not fully charged, and sourced by capacity that is above its contracted minimum charge level. Storage capacity that is below its minimum charge level is considered to be a load that is possibly responsive to real-time price signals.

Storage capacity can be contracted through the tariff market or the contracting process. For example, individual owners of plug-in electric vehicles (PEVs) could subscribe to tariffs that provide for both charging of the batteries as well as limited discharging as needed for load balancing by the contracted broker. On the other hand, a battery-exchange service for electric vehicles might negotiate a contract for the use of a portion of its current battery inventory for balancing purposes.

3.2 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 7.1 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.

Default tariffs: At game initialization, the tariff 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 power 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 if the default broker’s balance is not accurately balancing supply and demand.

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 per **Tariff Period**, which is typically once every 6 simulation hours.

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

Portfolio changes: New and dropped customer subscriptions, consisting of the customer model ID, the tariff ID, and the number of individual customers within the customer model. This is **private information**, sent to the tariff owner.

Tariff transactions: Tariff publication fees, signup bonus and early-exit penalty transactions corresponding to the subscription changes. This is **private information** for the tariff owner.

The following **public information** is sent to all brokers once per **Timeslot**, which is typically once every 1 simulation hour.

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 timeslot.

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

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

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

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

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 timeslots on the wholesale market.

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

4 Customer market

The simulation can include a range of customer models, including electric vehicles, CHPs, solar panels and wind turbines, and multiple models of private households. An important feature of these models is their responsiveness to price changes. A special focus lies on modeling substitution effects between timeslots as longer-term price elasticity effects would be very limited in 60 days of simulation time. In the literature such effects have been analyzed by means of synthetic aggregate models [18] or micro-founded bottom-up models [10]. Power TAC's dynamic customer models can extend both approaches to describe a rich customer population. Moreover, customers can not only be parameterized to reflect varying behavior but can even be swapped for other implementations. This adaptability is a key aspect of Power TAC's research proposition to analyze and guide the development of local energy markets.

In the game context customers perform three major tasks; choosing tariffs, recording meter readings and providing balancing capabilities. From a technical perspective customers are realized in the form of plugins. A customer model plugin instantiates a population of a customer type. Such population models can represent large groups of relatively homogeneous customers, which helps to reduce computational complexity. The plugin approach allows researchers to investigate questions related to specific consumer types or behavioral assumptions by using only relevant customer models.

4.1 Customer types

At least the following customers types will be implemented:

- households – typical residential consumption behavior, including limited production from solar and possibly small-scale CHP plants.
- offices – typical flat consumption throughout working hours, limited consumption at other times.
- factories – similar to office consumption but with greater magnitudes and more variations.
- electric vehicles – large loads (positive when charging and negative when feeding back to the grid) only when connected to grid, otherwise zero.
- institutions – Universities, municipalities, hospitals.

A customer's load profile is further specified by the power types it supports. A customer includes at least one of these types:

- consumption — power flow from grid to customer.
- interruptible consumption — power flow from grid to customer that can be interrupted by the DU within certain bounds, typically characterized by heat-storage capacity.
- production — power flow from customer to grid; this power type is further split into sub types that allow differentiation of power sources.
- storage — power flow to and from the grid; continuous operation in one direction is limited by storage capacity.

4.2 Tariff market interaction

The tariff market facilitates the matching of consumers and brokers. Customer models actively participate in the tariff market by choosing new tariffs through periodic evaluation of the tariffs offered by the brokers.

The key part of customer tariff evaluation is calculation of the expected cost (gain) over the lifetime of a contract relationship. This quantity is composed of the expected variable payments from estimated consumption (production), periodical payments as well as sign-up fees or bonuses. Especially the derivation of expected variable payments is crucial: It needs to properly reflect a customer's consumption (production) choice under the tariff to be evaluated. Therefore, tariff choice needs to be fundamentally driven by consumption choice under a tariff as described in the next

subsection. Especially for complex tariffs this is a key design challenge for creating customer models. Since early exit is possible, customer models may evaluate available tariffs at any time. Clearly in this case, a proper switching evaluation has to additionally factor in the exit fees from leaving the current tariff. This monetary evaluation is complemented by an additional assessment of other tariff aspects, e.g. broker reputation, energy sources, interruptibility properties or early exit fees. The tariff comparison is therefore described by a utility value for each available tariff. This value moderates costs and other factors. The tariff utility function and the corresponding tariff choice logic are the key characteristics of customer model actions in the tariff market. Elicitation of these tariff preferences is thus a major aspect of a successful broker strategy.

From the currently available tariff list customers need to select a suitable one (see Figure 7). This is a two-step problem:

1. Derive the utility value for the current tariff and the new tariffs to be considered — this could be either all tariffs or just a (random) subset.
2. Compare all evaluated tariffs and choose (most) suitable one




Figure 7: Tariff selection problem.

The implementation of the tariff selection problem is described in the remainder of this section.

4.2.1 Derive tariff utility

To derive the utility of any given tariff, customers need to jointly evaluate costs, energy sources, broker reputation and tariff risk to determine a tariffs suitability. For customer tariff utility we assume generalized additive independence between the attributes. Tariff utility can then be represented as

$$u_i = -(c_v + c_f)\alpha_{cost} - r_i\alpha_{risk} - I_i\alpha_{inertia}. \quad (3)$$

The alphas are customer-specific weighting parameters for the different tariff-specific realizations of the sub-disutility types. The sub-disutility values for tariff costs ($c_v + c_f$), tariff risk (r_i) and inertia (I_i) are evaluated using functions common to all customers:

Variable tariff costs c_v Consumption payments are determined by sampling k random days, deriving each day's optimal consumption under the tariff to be evaluated and finally averaging the realized cost: $c_v = \frac{\sum_k c_v^*(k)}{k}$. For variable tariffs this calculation is performed using the average realized values.

Fixed tariff payments c_f Fixed tariff payments consist of sign-up fees/ bonuses of the new tariff, $c_{\text{sign-up}}$, daily periodic payments c_{daily} as well as exit fees of current tariff c_{exit} . These costs need to be normalized to a one day time span. While this is trivial in case of the periodic payment, it requires the expected tariff life \tilde{t} for the other payments.⁴ The normalized values of the fixed payments are summed to obtain the fixed other payments value, $c_f = c_{\text{daily}} + \frac{c_{\text{sign-up}} + c_{\text{exit}}}{\tilde{t}}$.

Tariff risk r_i Under a dynamic contract customers face the risk of unfavorable rate developments. Hence, they evaluate a dynamic tariff's rate risk using the variance of the realized prices.

Customer inertia I_i Customers have behavioral cost of changing a tariff. These are reflected by the inertia term. Given the current tariff j , I_i is defined as

$$I_i = \begin{cases} 1 & \text{if } i \neq j \\ 0 & \text{if } i = j. \end{cases} \quad (4)$$

With this procedure customers can assess the utility of any tariff offered. This utility is the foundation of the customer tariff selection as described in the next section.

4.2.2 Choose from a list of tariffs

An overall tariff choice does not need to strictly follow a deterministic choice of the highest utility value. This is especially important for population models that wrap a larger group of customers.

A smoother decision rule which allocates the selection choice proportionally over multiple similar tariffs is therefore needed. A *logit choice model* facilitates this type of tariff choice randomization. Instead of providing a discrete tariff decision, a choice probability \mathbb{P}_i is obtained for each tariff i from the set of tariffs considered \mathbb{T} :

$$\mathbb{P}_i = \frac{e^{\lambda u_i}}{\sum_{t \in \mathbb{T}} e^{\lambda u_t}} \quad (5)$$

The parameter $\lambda \geq 0$ 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 utility. 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.

4.3 Provide balancing capacity

Customers can provide brokers with different forms of balancing capacities, determined by the **PowerType**. These differ in availability and the amount of balancing energy available.

⁴The derivation of \tilde{t} may be customer-specific.

Interruptible consumption: Certain types of appliances (water heaters, heat pumps) can support remote interruption by the DU. If a broker has such interruption under contract, its use can be offered to the DU to avoid balancing charges.

Pledged energy from storage: By pledging stored energy customers with energy storage can provide balancing capacity — limited by the storage unit’s discharge power and level of charge.

Controllable micro generation: While intermittent producers typically cannot provide balancing capabilities, non-intermittent producers like CHPs or bio-gas units can pledge extra generation capacity for balancing purposes.

4.4 Consume and produce power

Customer models ultimately represent the entities connected to the grid. As such the game implications of their actions can be represented as timeslot meter readings for both consumption (positive reading) and generation (negative reading). 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 capability actions (respond to current or previous load interruption). Lastly, game-dependent factors include all load adjustment triggered at runtime by the game environment, e.g. randomization, the current season and weather conditions (e.g. turning on A/C, output from solar panels).

Of central interest in the Power TAC research setting is the effect of customer tariff choices on realized load patterns. This relationship between a customer’s tariff and the meter reading is described by an economic consumption or generation logic. In the following sections typical implementations for these consumption/ generation logics are described. Clearly, other load influence factors (weather, balancing actions) do of course affect this tariff-dependent consumption logic by changing the base load level or inducing ex-post distortions.

Fully static: These are customer models that do not adjust their consumption to the rates of their current tariff, i.e. the meter readings of these customers are independent of their selected tariff. This could be due to lack of shifting capabilities or relative insignificance of electricity costs (rich customers, certain industrial customers). This is also the appropriate model for non-controllable generation facilities (e.g., solar or wind).

Static amount, flexible timing: Customer models who can change the timing of their loads (e.g. through automatic appliance scheduling) will not change their consumption amount under a given tariff but will try to minimize their cost by scheduling the activities appropriately. Such household models are typically bottom- up models where consumption originates from the activity/ appliance level [10].

Flexible amount, static timing: This type of customer model implements a simple demand behavior: for each timeslot the optimal consumption amount is decreasing in the timeslot electricity price. Such customer models reflect synthetic consumption profiles determined in a top-down approach considering aggregate electricity consumption as a continuous good with positive and decreasing marginal utility. Controllable generation with well-defined cost functions (e.g. micro-CHP) is also captured by this modeling approach.

These models are especially helpful for economic analyses as their behavior can be described in compact mathematical form.

Fully dynamic: Fully dynamic consumption features both flexible consumption amounts as well as flexible timing. Such models can be both top-down as well as bottom-up. While bottom-up models in this group formulate appliance-level usage decisions taking into account prices and available income, top-down models specify cross-price elasticities between timeslots [18]. Fully dynamic bottom-up models endogenize price for activity occurrence and scheduling.

4.5 Available customer models

In the following we list and describe the customer models which are implemented in the current Power TAC release. This list will be continuously updated

Household model

This model represents a neighborhood of residential customers (houses) as described by [10]. The houses are aggregated in a population (village) which handles the tariff market interactions such as tariff subscriptions or allocation of aggregate consumption to the tariffs.

The houses themselves are characterized by a randomly initialized number of household members (e.g., mostly present or working persons) and a set of appliances (e.g., stove, heater, fridge). The household electricity consumption is driven by a combination of the household member occupancy profile and the appliance runtime characteristics. See Figure 8 for an illustration.

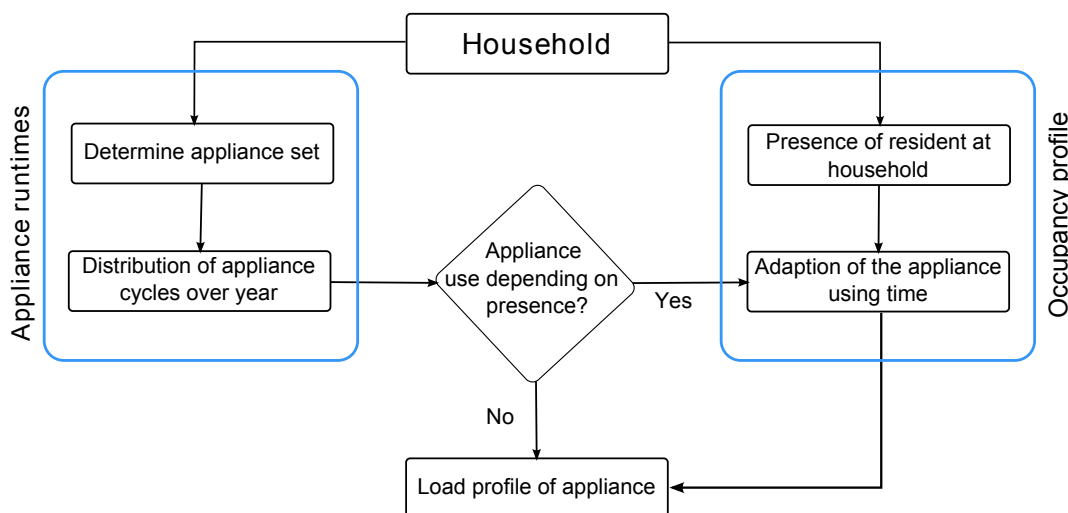


Figure 8: Interaction of presence and appliance model [10]

A key element of this model is the ability to automatically shift certain loads with the goal to minimize electricity cost given the daily price vector of the current tariff. The household’s basic load profile for the whole game is derived at initialization using the occupancy and appliance model. At runtime the model periodically evaluates and executes shifting opportunities.

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 EEX⁵. The brokers can buy and sell power contracts for future timeslots to optimize their portfolio. 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.

5.1 Trading and timeslots available for trade

Brokers can submit orders to the wholesale market for delivery between one and 24 hours in the future. The timeslots available for trading are marked as “enabled”; changes in timeslot status are communicated to brokers at the beginning of each timeslot. Orders submitted for non-enabled (disabled or not yet enabled) timeslots are silently discarded. Depending on the market configuration brokers may also be able to delete submitted orders from order books. 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 timeslot 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 timeslot 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.”

5.2 Market clearing

When the simulation clock is advanced to a new timeslot, the wholesale market clears the orderbook for each of the enabled timeslots. Note that at the beginning of the clearing process an updated list of enabled timeslots is sent to each broker, but the set that is considered in clearing is the set that was enabled immediately before the clearing process started. This is done to minimize the period of time in which the set of enabled timeslots from the broker’s viewpoint differs from the set of enabled timeslots from the market’s viewpoint.

In the clearing process, as shown in Figure 9, demand and supply curves are constructed from bids and asks to determine the clearing price of each orderbook (one for each enabled timeslot) at the intersection of the two, which is the price that maximizes turnover. 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 are sorted first, as though they had the highest bid prices or the lowest ask prices.

⁵See <http://www.nordpoolspot.com>, <http://www.ferc.gov>, or <http://www.eex.com/en>.

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 bid and the highest ask price supporting this maximum turnover. 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 9 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 the mean of the prices in ask 6 and bid 8, or 16.

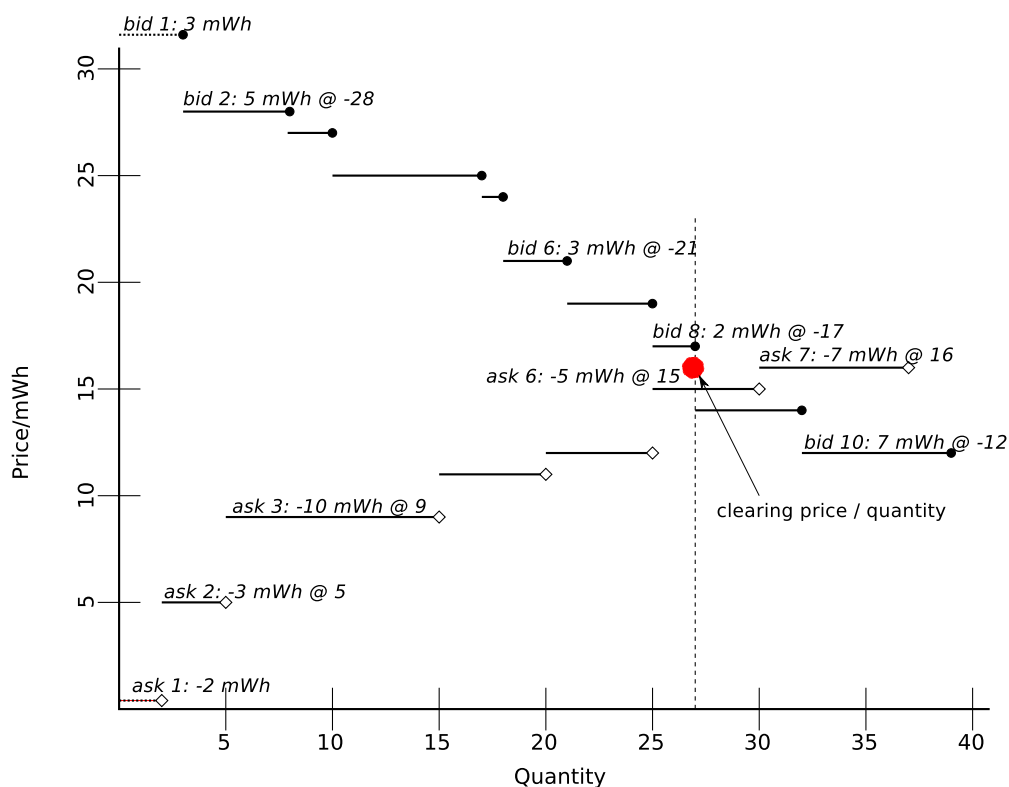


Figure 9: Market clearing example: bid 8 and part of ask 6 are the last to clear.

After the market is cleared the following steps are performed:

- Clearing price and volume are publicly broadcast (public information). In the example of Figure 9, this would be (27, 16).
- Post-clearing orderbooks are published for each cleared timeslot, 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 (private information).
- Updated cash and market positions are computed and communicated to brokers (private information).
- All orders are discarded.

5.3 Wholesale suppliers and buyers

To ensure liquidity to the wholesale market, the simulation includes both wholesale energy providers as well as wholesale buyers. The wholesale suppliers are called Generation Companies, or Gencos for short. Each Genco g has a nominal capacity \hat{C}_g , a fixed cost/MWh c_g , a commitment leadtime τ_g , and a reliability value r_g . Actual capacity $C_{g,s}$ in timeslot s varies around the nominal value by either a mean-reverting random walk, or by current weather conditions in the case of wind turbines. Given a variability parameter v , a mean-reversion rate m , and a uniformly distributed random value ν on $[0..1]$, the random walk is defined as

$$C_{g,s} = C_{g,s-1} + v(2\nu - 1)\hat{C}_g + vm(\hat{C}_g - C_{g,s-1}) \quad (6)$$

At any given time, each Genco is “in operation” with a probability r_g . If a Genco is in operation, it will submit an ask to the market for its uncommitted capacity at its fixed cost in each future timeslot that is farther in the future than its commitment leadtime τ_g . Once it has sold at least some power for a given timeslot, it is committed, and will attempt to sell the remainder by continuing to submit asks in each enabled timeslot, including those closer to the current time than its commitment leadtime. If it fails to sell at least some power in a given timeslot by its commitment time, then it will withdraw its capacity from the market for that timeslot.

Once a Genco has sold power for a given timeslot, it will deliver the power, regardless of its capacity or operational status. We assume it has the ability to purchase power from others, if necessary, to meet its commitments.

The exact set of Genco entities in the simulation and their parameters are not specified, but will be revealed to brokers at the beginning of a simulation. The available set of Gencos will be sufficient to cover the demand in the simulation. This can be assured by providing one high-priced, high-capacity Genco with a minimal leadtime.

In addition to the Gencos, there is a wholesale buyer b_b with stochastic behavior that simulates a population of buyers and speculators. Its behavior is very simple: Given two parameters, a quantity q_b and a mean price p_b , and a random value ν , it computes a price $p_{b,s} = -p_b \ln(1 - \nu)$ for each timeslot s and places a bid $(b_b, s, q_b/p_{b,s}, p_{b,s})$ in each open timeslot. This exponential distribution produces large numbers of low-priced high-quantity bids, and a few higher-priced low-quantity bids.

6 Balancing market

In electricity markets, supply and demand have to be balanced almost perfectly in real time. A major task of the Independent Systems Operator (ISO)⁶ on the wholesale (transmission) level and

⁶In Europe the name Transmission Systems Operator (TSO) is used instead of ISO.

of the Distribution Utility (DU) on the regional (distribution) level 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 [24]. 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 [17].

In Power TAC, brokers accumulate credits and debits to their energy budgets for each timeslot by selling (exporting) power or buying (importing) power in the wholesale market, and by the power consumption and production activities of their contracted customers. To carry out its responsibility to balance supply and demand in each timeslot, the DU may exercise capacity controls (see below) on behalf of brokers, and it may import or export power through an “ancillary services” or “regulating” market at prices that are normally much less attractive than the prices faced by brokers in the wholesale market (see Figure 10).

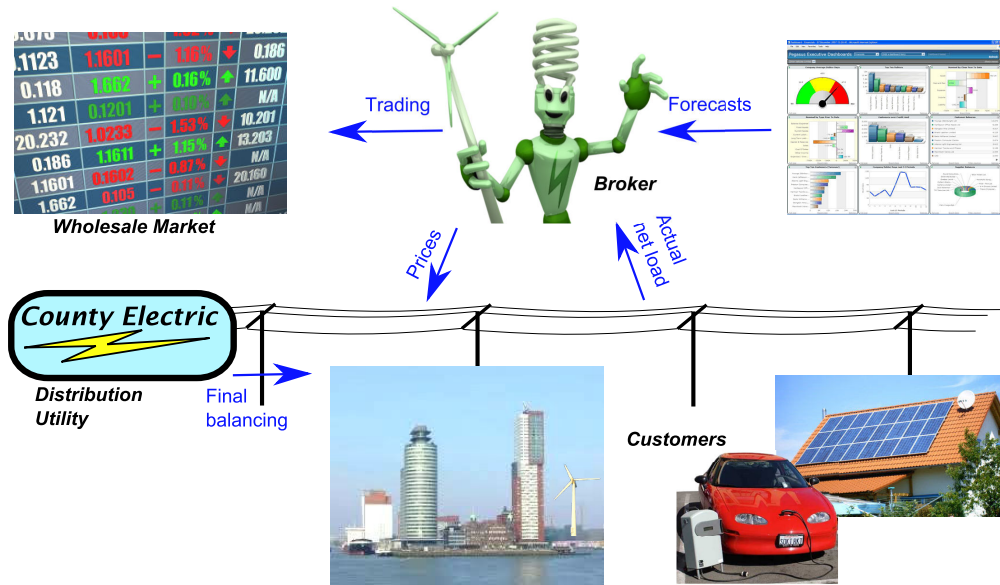


Figure 10: Entities and activities during balancing.

Brokers acquire balancing capacity by offering price concessions in exchange for the ability to remotely interrupt loads or sources for limited periods of time. Balancing capacity consists of “interruptible” or “controllable” load or source devices. These are connected to controllers installed at a customer site that allow the DU to interrupt or modulate power flow for a certain time periods, dependent on the type of contract the broker has with its customer. Most examples of balancing capacity are associated with thermal or battery storage devices, such as CHPs (Combined Heat and Power) systems that produce power when heat is needed, and domestic water heaters that can be interrupted for periods of time without significantly impacting customer convenience.

6.1 Adjusting energy demand and supply

Here we explain more formally how the simulation computes balance, how brokers can act to avoid imbalances, and the actions taken by the DU to achieve balance. The total energy consumption

$e_c(b, s)$ for broker b in timeslot s is

$$e_c(b, s) = e_{\text{ex}}(b, s) + \sum_{i=1}^{|\mathcal{C}_b|} e_i(s) \quad (7)$$

or the sum of the loads during timeslot s of each energy consumer in the set \mathcal{C}_b , the consumers in the portfolio of broker b , plus the energy exported e_{ex} from the grid by broker b during timeslot s through sales commitments in the wholesale energy market (see Section 3.1.3). Similarly, the total energy production for broker b in timeslot s is

$$e_g(b, s) = e_{\text{im}}(b, s) + \sum_{j=1}^{|\mathcal{G}_b|} e_j(s) \quad (8)$$

or the sum of outputs during timeslot s of each energy producer in the set \mathcal{G}_b of producers in the portfolio of broker b , plus the energy imported e_{im} by b through purchase commitments in the wholesale market.

In this context, balance between supply and demand means that supply equals demand for each broker in each timeslot,

$$\forall s \in \mathcal{S}, e_g(b, s) - e_c(b, s) = 0 \quad (9)$$

Note that $e_g(b, s)$ can include an arbitrary portion of contracted controllable production capacity, and $e_c(b, s)$ may include, as described in the following, an arbitrary portion of contracted controllable load. Broker actions to buy or sell energy in the wholesale market, and to contract for balancing capacity, can affect only future timeslots, not the current timeslot. Ultimately, it is the job of the DU to ensure exact balance between supply and demand in real time. Any imbalance remaining after summing supply and demand across all brokers will be balanced by the DU, by invoking brokers controllable sources and loads, and by increasing or decreasing power draw from the transmission system through the wholesale regulating market. Costs for regulating power, along with DU fees, are charged to the brokers who are responsible for the residual imbalance as we shall see in the following section.

For each timeslot s of length τ , each broker b should ideally balance expected supply and demand closely enough that the DU can achieve exact balance without requiring regulating services. Expected demand is the total expected load, or the sum of committed power exports and the expected loads $E(e_c(b, s))$ of each consumer i in the broker's consumer portfolio \mathcal{C}_b during timeslot s (see Equation 7):

$$E(e_c(b, s)) = e_{\text{ex}}(b, s) + \sum_{i=1}^{|\mathcal{C}_b|} E(e_i(s)) \quad (10)$$

Expected supply is committed power imports plus total expected production capacity of all generators g within the broker's portfolio \mathcal{G}_b during timeslot s (see Equation 8):

$$E(e_g(b, s)) = e_{\text{im}}(b, s) + \sum_{j=1}^{|\mathcal{G}_b|} E(e_j(s)) \quad (11)$$

These values are maximum values in case some customers in the broker's portfolios have agreed to external control, presumably in exchange for better prices. For example, a combined heat and

power generator with a nominal output of 50kW can be adjusted by an external control so that its real production is within certain boundaries, e.g., [40kW–50kW]. Similarly, a domestic water heater may be configured to permit remote shutoff for up to 15 minutes every hour. The total controllable load for a broker b during timeslot s is $\epsilon_c(b, s)$, and the total controllable production capacity is $\epsilon_g(b, s)$, where it is understood that a control that increases demand is equivalent to a control that reduces supply, and vice-versa. As long as $e_g(b, s) - \epsilon_g(b, s) \leq \epsilon_c(b, s)$ and $\epsilon_c(b, s) - \epsilon_c(b, s) \leq \epsilon_g(b, s)$, then supply and demand during timeslot s is expected to be in balance. Within this range, the DU will either reduce load or reduce output as needed to achieve exact balance.

The activation of balancing power (or load) by the DU is done only during the current simulation timeslot s_n . In Figure 11(a), we can see in the current slot s_n that both the actual observed supply and demand have deviated from the forecasted overall supply and demand for broker b . But as the difference between $\epsilon_c(b, s_n)$ and $\epsilon_g(b, s_n)$ was smaller than $\epsilon_g(b, s_n)$, the controllable production capacity of broker b in this slot, the DU was able to automatically reduce supply such that overall demand and supply for timeslot s_n was rebalanced.

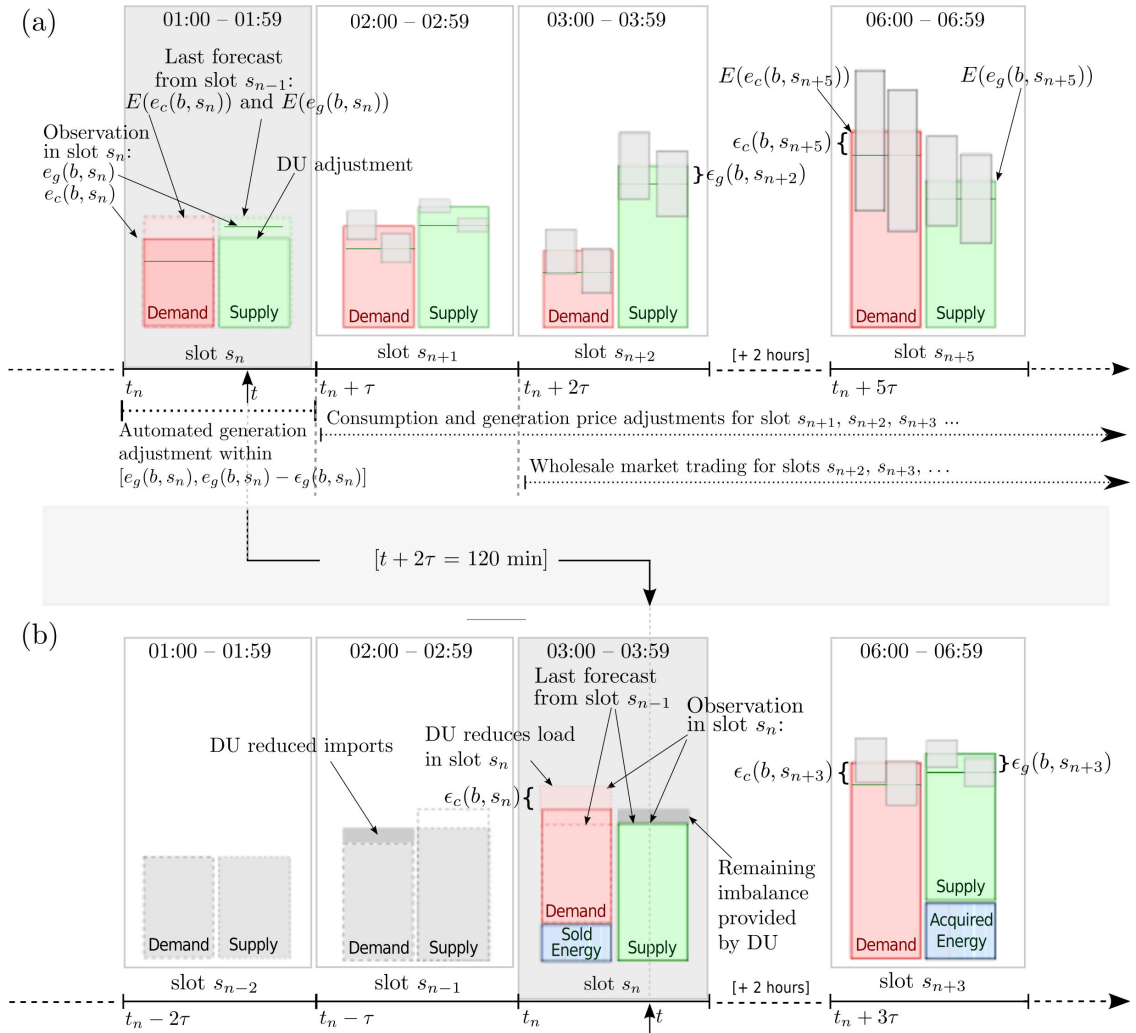


Figure 11: Broker's expected and actual energy supply and demand at two points in time.

For timeslot s_{n+1} in Figure 11(a), expected overall demand is forecasted to be within range of the available production capacity, but the uncertainty envelope (grey boxes) shows that this is not certain. In other words $E(e_g(b, s_{n+1})) - \epsilon_g(b, s_{n+1}) \leq E(e_c(b, s_{n+1}))$. After 2τ simulation time has elapsed (Figure 11(b)), this slot is now designated s_{n-1} , and we can see that the real consumption $e_c(b, s_{n-1})$ in this timeslot turned out to be lower than $e_g(b, s_{n-1}) - \epsilon_g(b, s_{n-1})$. This means that even after the DU reduced the broker's production capacity to its minimum level, the overall production still exceeded the overall consumption. In this case the DU either reduced imports through the regulating market, or matched the surplus with a shortage of power from some other broker, to absorb the excess generated energy.

In slot s_{n+2} in Figure 11(a), a significant difference between overall production and overall consumption is forecast. Internal balancing capacity is likely to be insufficient for leveling the expected difference. In order to avoid the (expensive) utilization of external balancing power, broker b can either sell some of its surplus energy on the wholesale market, or use its contracted pricing power to try to encourage (i) some or all of its consumers to increase their demand, or (ii) some or all of its producers to reduce their production.

Technical adjustments by brokers (e.g. a remote activation of capacities at consumer premises) is not allowed within the competition; only the DU acts in the current timeslot. But a consumer's energy consumption is subject to the energy consumption price for consumer i in a timeslot s , which is defined as $p_c(i, s)$. We define

$$\hat{e}_c(i, s_{n+2}) = E(e_c(i, s_{n+2}, p_c(i, s_{n+2}))) \quad (12)$$

as the predicted load for consumer i in timeslot s_{n+2} , given price $p_c(i, s_{n+2})$. If the broker changes the underlying consumption price to $p'_c(i, s_{n+2})$ the forecasted consumption of this consumer is expected to increase as

$$\hat{e}'_c(i, s_{n+2}) = e'_c(i, s_{n+2}, p'_c(i, s_{n+2})) \quad (13)$$

The ratio of demand change to price change

$$PE_i = \frac{\hat{e}_c(i, s, p) - \hat{e}_c(i, s, p')}{p - p'} \quad (14)$$

is called the "price elasticity" for consumer i . Price elasticities will to be modeled within the different consumer agents provided by the competition environment following empirical findings on price elasticity as described for example in [22, 19].

Some customers in the broker's portfolio (such as electric vehicle batteries that can be discharged into the grid) might have agreed to flexible pricing as well, and therefore their output will be sensitive to price in a similar way. In other words, the power generation capacity of broker b in timeslot s , $e_g(b, s)$, is likely to change if the generation price $p_g(j, s)$ is changed to $p'_g(j, s)$, decreasing if $p'_g(j, s) < p_g(j, s)$. Next we discuss how the DU sets prices for balancing services.

6.2 Market-based balancing mechanisms

We present three different scenarios and the related mechanisms to balance the market and when they will be used:

Scenario I: no controllable capacities This was implemented for the 2011 pilot release.

Scenario II: static with controllable capacities This will be implemented for the 2012 competition.

Scenario III: dynamic with controllable capacities This may be implemented as an option in the 2012 competition.

In the following we discuss the desirable properties and the different scenarios. More detailed background and examples on the balancing market can be found in [8].

6.2.1 Desirable balancing mechanism properties

The main goal of a real-time balancing mechanism is to have a balanced system, using the services of the wholesale regulating market, local storage or spinning reserves, or controllable loads and sources made available by brokers, such that demand and supply is matched exactly. To arrive at this goal, the balancing market prices imbalanced portfolios in a way that is intended to motivate brokers to achieve balance on their own. We first discuss desirable properties of such a pricing mechanism, and we analyze what information is private to the brokers. These properties and the relevant private information differ slightly depending on whether brokers have access to controllable loads and sources. We start with the properties that hold for all scenarios.

1. A desired property is to have an *efficient* system, i.e., which optimizes social welfare.
2. To arrive at this, we do not just want efficient solutions regarding how imbalances are resolved just in time, we, in fact, would like to have as little imbalance as possible between broker commitments in the day-ahead market and the actual net load experienced in real-time. The idea is that generally more efficient allocations are found when imbalances are resolved in the day-ahead market (or even earlier), simply because there are more options then to produce (or consume) additional power. For example, some generators have a start-up time of several hours. Consequently, the strategy of brokers to have a portfolio with (almost) *no net imbalance* should be *incentive compatible*.
3. Since the DU is responsible for the real-time balancing of the portfolio across all brokers, we can argue that the DU should be compensated fairly for its services. An additional desired property then is to ensure that the payments offered to the DU are always sufficient to cover its costs. A pricing mechanism meeting this criterion is called *weakly budget balanced*.

In scenario I (without controllable capacities) restoring the balance is done solely by the DU interacting with regulating capacity available in the wholesale market. However, to optimally restore the balance when brokers can have controllable capacities (scenario II), we need to extract additional information regarding costs and capacities of their controllable loads and sources.

4. Since manipulating the costs of potential controllable capacities can lead to sub-optimal solutions, an additional goal in this setting is to make the strategy of declaring *the true capacities and costs of controllable capacities* to be *incentive compatible*.
5. A second criterion in the case of controllable capacities is a so-called participation constraint, i.e., the mechanism should benefit participating brokers, or otherwise brokers just will not declare any controllable load at all. In other words, the mechanism should be *individually rational*.

In the following two sections we discuss mechanisms that can be used for the above two scenarios (with/without controllable capacities). These mechanisms meet all the desirable properties described above.

6.2.2 Scenario I: no controllable capacities

The relevant players in all three scenarios are the N brokers denoted by $\{1, 2, \dots, n\}$, and the system operator or distribution utility (DU), denoted by 0. In our analysis, we assume that energy production and consumption are more or less stable during a single timeslot. In any given timeslot, each broker $i \in N$ has an expected net local energy (potentially negative) surplus of $x_i \in \mathbb{R}$. Furthermore we use P^+ to denote the maximum market price of energy for this timeslot in the day-ahead market (over all day-ahead trade periods), P^- to denote the minimum market price over all day-ahead trade periods for this time slot, and $P^* \geq P$ the all-time highest price possible in any time slot.

The DU has the potential to import or export energy on the wholesale market (or apply its own ancillary services such as spinning reserves) to arrive at a perfectly balanced energy production and consumption. This comes at a cost of $c_0 : \mathbb{R} \rightarrow \mathbb{R}$ per unit. See Figure 12 for an example, illustrating that the cost of buying additional energy is higher than the benefit of selling additional energy at this last instance.

Given actual imbalances, the DU can compute the net imbalance $\hat{x} = \sum_{i \in N} \hat{x}_i$ over all brokers, and then apply its cost function c_0 to $-\hat{x}$ to determine the (expected) total costs for balancing, i.e., $-\hat{x} \cdot c_0(-\hat{x})$. Since in this first scenario there is no other way to recover from imbalances, this meets our requirement (1) of an efficient solution.

Payments need to be set such that the incentive compatibility and budget balance requirements are met. We denote these payments by p_2 . Since these payments are computed after the timeslot supply and demand are known, we can base them upon the real imbalances x_i . The second requirement in fact implies that the payment for an imbalance should always be higher than resolving the imbalance against the maximum market price P^+ in the day-ahead market, i.e., $p_{2,i} \geq -x_i \cdot P^+$ if $x_i < 0$, or otherwise $p_{2,i} \geq -x_i \cdot P^-$. Finally, the third requirement just says the payments from the brokers (who consume more than they produce) should be more than the payments to the brokers (who produce more than they consume) and the costs for recovering from the imbalance together, i.e., $\sum_i p_{2,i} + x \cdot c_0(-x) \geq 0$.

Given these constraints, there are infinitely many possible choices for these payments, since they are only bounded from below. However, we are convinced that a DU should not profit significantly from any imbalances, and the payments should be fair in the sense that brokers that produce too much in an over-consuming market, or brokers that consume too much in an over-producing market should not pay as much as the others. We therefore propose to minimize the difference between the payments and the costs (or profits) attached to resolving the imbalance in the day-ahead market. In the following mathematical programming model, let $p_{2,i}$ denote the payment of broker i ; this is the only variable, since x_i , P^+ , and P^- are given.

$$\begin{aligned}
& \text{minimize} && \sum_{i \text{ if } x_i < 0} (p_{2,i} + x_i \cdot P^+)^2 + \sum_{i \text{ if } x_i \geq 0} (p_{2,i} + x_i \cdot P^-)^2 \\
& \text{subject to} && p_{2,i} \geq -x_i \cdot P^+ && \text{if } x_i < 0 \\
& && p_{2,i} \geq -x_i \cdot P^- && \text{if } x_i \geq 0 \\
& && \sum_i p_{2,i} + x \cdot c_0(-x) \geq 0
\end{aligned} \tag{15}$$

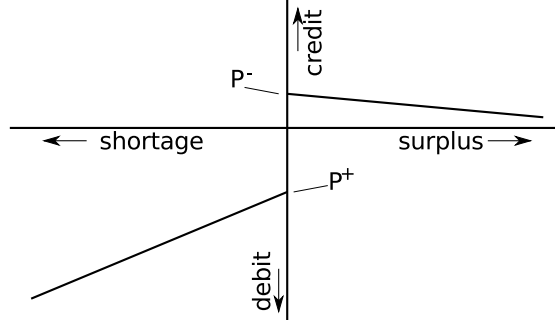


Figure 12: The price paid for an energy surplus is always lower than the lowest price observed in the day-ahead market, while the price charged for an energy deficit is always higher than the highest price observed in the day-ahead market.

This program is a (quadratic) convex program if $c_0(\cdot)$ can be modeled by a (set of) linear function(s); it then can efficiently be solved, e.g., using interior point methods [4].

According to this definition over-consuming ($x_i < 0$) brokers always have to pay a positive amount. In the program in Equation 15 the distribution of the costs of balancing is defined by the minimization criterion, which expresses that each broker should pay an equal portion above the minimum amount defined by the constraints. However, this minimization criterion can be chosen differently (e.g., never let over-producing brokers pay a positive amount).

Scenario I will be used for the pilot competition in summer of 2011, using data from the econometric analysis of [20] for parameter settings, scaled by the size of the customer population.

6.2.3 Scenario II: static with controllable capacities

The controllable capacities for each broker i are represented by a capacity range for its controllable production (or consumption) $[c_i^-, c_i^+]$, and a function describing the price (absolute costs per unit) of diverting from its production x_i , i.e., $c_i : [c_i^-, c_i^+] \rightarrow \mathbb{R}$, similar to the up- and downward regulation of the DU. We assume this to be a monotonically increasing (often step-wise) function, since it represents all contracts that include a controllable part, usually at different prices and capacities and first (for rational agents) the cheapest options are used. Examples of such upward regulation contracts are the possibility to turn-off lights or heat pumps, or turn on CHPs, and examples of downward regulation contracts are pre-loaded washing machines, the charging of batteries (also of electrical vehicles), and the possibility to temporarily tune down production capacity.

The distribution utility needs to make sure that for every $i \in N \cup \{0\}$ some extra production (or consumption) δ_i within the possibilities is chosen at minimal total costs such that all energy consumption and production is balanced, i.e.,

$$\begin{aligned} & \text{minimize} && \sum_{i \in N \cup \{0\}} \delta_i \cdot c_i(\delta_i) \\ & \text{subject to} && \delta_i \in [c_i^-, c_i^+] \\ & && \sum_{i \in N \cup \{0\}} (x_i + \delta_i) = 0 \end{aligned} \quad (16)$$

This may or may not include an increase/decrease of production regulated by the DU itself, dependent upon the costs. For now we assume this possibility to be unlimited, i.e., $[c_0^-, c_0^+] = [-\infty, \infty]$.

In case the correct information is used, Equation 16 meets the first requirement (of efficiency). Furthermore, if all the functions c_i are monotonically increasing, this problem is convex and can therefore again efficiently be solved [4].

The payments to incentivize the brokers to provide the correct information consists of two parts:

1. The same mechanism as in scenario I (with payments p_2) is used to make having no imbalance a dominant strategy.
2. An additional payment p_3 is introduced in this section to make declaring the true cost function for balancing capacity a dominant strategy.

The utility of a broker i for a solution δ is not just defined by these payments, but also by the costs of load control, i.e., $u_i(\delta) = -c_i(\delta_i) - p_{2,i} - p_{3,i}(\delta)$. Regarding the incentives, note that the first payment is completely independent of the others, and that the analysis of the incentives from the previous section thus automatically transfers. The focus of this section will be on the third payment.

No real-time matching among brokers First observe that the strategic opportunities for brokers relate to the cost function of their controllable capacities. However, even then the social costs for balancing can be reduced by real-time matching upward and downward regulating services among brokers. If this is possible, these brokers could also have realized this exchange in the day-ahead market. It turns out that forbidding such exchanges in the real-time balancing phase sufficiently restricts the setting to meet all given requirements. The additional conditions are that for all $i \in N \cup \{0\}$ it holds that

$$\begin{aligned} \delta_i &\geq 0 && \text{if } \sum_{i \in N} x_i < 0 && \text{(under-production)} \\ \delta_i &\leq 0 && \text{if } \sum_{i \in N} x_i > 0 && \text{(over-production)} \end{aligned} \tag{17}$$

With these conditions and step-wise cost functions, the mechanism is similar to a multi-unit auction (in the case of over-production) or a reverse multi-unit auction (in the case of under-production) [16]. With any type of cost functions, [11] mechanisms are the only mechanisms to achieve both an efficient allocation and a truthful declaration, in our case of the cost function of the controllable capacities. Within this class, the VCG mechanism [25, 6] ensures that brokers always have a nonnegative utility for participating (i.e., individual rationality) under two conditions that hold in this domain: (i) choice-set monotonicity, which says that removing an agent never increases the set of alternative solutions, and (ii) no negative externalities, which says that every agent has zero (or more) utility for any choice that is made without its participation. When VCG is applied in the above setting, given the optimal production vector δ , the payment for each broker i is defined as follows (note that the sign is flipped because VCG is defined on the maximum social welfare, not on the minimal costs).

$$p_{3,i}(\delta) = - \sum_{j \neq i} \delta_j^{-i} \cdot c_j(\delta_j^{-i}) + \sum_{j \neq i} \delta_j \cdot c_j(\delta_j), \tag{18}$$

where δ^{-i} denotes the optimal solution to Equation 16 in a situation where the controllable capacities of broker i cannot be used, i.e., $\delta_i = 0$. In case of over-production, all payments are positive, and thus the VCG mechanism meets all our requirements.

6.2.4 Scenario III: dynamic with controllable capacities

In practice, the process of Scenario II is repeated every time slot, and incentive compatibility is not automatically transferred to such a dynamic setting. In a dynamic setting, the static problem needs to be solved for each time slot t (of e.g. 60 minutes). The main difference is that all variables become time dependent (i.e., functions of t). The cost function p_i may not only be different for different time slots, because of specific contracts made previously, but also a reduction in consumption now may require an increase in consumption in a subsequent time slot. Consequently, these functions $p_i(t)$ become dependent upon decisions made for previous time slots. When VCG is repeatedly applied in such a setting, its property of truthfulness is not guaranteed anymore. However, the fact that VCG can be used in the static scenario promises good news for the application of a so-called dynamic-VCG mechanism [5].

7 Competition format and interaction

Number of broker agents As opposed to previous TAC competitions where the number of agents 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. We will pick a few game-size values and group them into different sized broker pools to simulate oligopolies as well as highly competitive markets.

7.1 Competition 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 period are collected and sent to brokers at the beginning of a game. More specifically, this "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 power consumption is given for each 1-hour timeslot during the 14-day bootstrap data-collection period. Values are negative if the default broker is supplying the power, positive if the customer is supplying power.

Market information: for each timeslot 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 timeslot 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 power 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 timeslot 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 timeslot, 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 timeslot.

Given the default broker's estimated net energy requirement (summed over all its models) for each of the following 24 timeslots, it attempts to build a market position equal to its estimated need for that timeslot. 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 timeslot s , except that if $s = t + 1$ (the last chance to purchase or sell power for timeslot 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

$$l_{prev} = \begin{cases} l_{s,t-1} & : \quad \text{if order in previous timeslot } t - 1 \text{ did not clear} \\ l_{max} & : \quad \text{otherwise} \end{cases} \quad (19)$$

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

$$l_{s,t} = \max \left(l_{min}, 2 \frac{l_{min} - l_{prev}}{s - t - 1} \right) \quad (20)$$

The standard competition parameters can be found in Table 1. Values for these parameters are sent to a broker at the start of every game. For details see the software documentation.

7.2 Competition ending

The game ends at a random number of K timeslots after day 55 (timeslot 1320), $K = 0, 1, \dots$. For each timeslot, starting day 55, there is a fixed probability p that the game ends by the end of that particular timeslot. As a consequence, the number of timeslots in excess of day 55, K , follows a geometric distribution. The expected number of timeslots in excess of day 55 is equal to $E(K) = (1 - p)/p$. The cumulative probability distribution that the game ends after at most k extra timeslots is equal to:

$$P(K \leq k) = 1 - (1 - p)^{k+1}, \quad \text{for } k = 0, 1, \dots \quad (21)$$

The probability ω that the game does not end before day 60 (timeslot 1440) is derived from the inverse cumulative distribution. More generally, we want the probability that the game takes more than k' timeslots to be at most equal to some ω :

$$P(K > k') \leq \omega \Leftrightarrow (1 - p)^{k'+1} \leq \omega \quad (22)$$

$$\Rightarrow k' \leq \frac{\ln \omega}{\ln(1 - p)} - 1 \quad (23)$$

The end-of-timeslot ending probability p will be based on:

Table 1: Parameters used in Power TAC tournament games.

Parameter	Symbol	Standard Game Setting
Number of brokers in a game	B	2, 4, and 8
Number of games in a round with 2 brokers	G_2	12
Number of games in a round with 4 brokers	G_4	6
Number of games in a round with 8 brokers	G_8	6
Length of pre-game bootstrap period		14 days
Nominal length of game	E	60 days
Probability that there are k timeslots after timeslot 1320 (start of day 55) before end of game	$[p_{min}, p_{max}]$	$[p_\omega, 1]$
Probability of game end for each timeslot after timeslot 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/timeslot)
Open timeslots on wholesale market		24
Market closing time		1 timeslot ahead
Distribution fee		[0.01 - 0.3]€/kWh
Balancing price basis	P	most recent clearing price
Balancing cost	c_0	[0.02 - 0.06]€/kWh
Default broker's min and max bid order prices	$l_{min}(\text{bid}), l_{max}(\text{bid})$	-100, -5
Default broker's min and max ask order prices	$l_{min}(\text{ask}), l_{max}(\text{ask})$	0.1, 30
Tariff publication fee		[100 - 500] €
Tariff revocation fee		[100 - 500] €
Tariff publication interval		6 timeslots
Annual bank debt interest rate	$[\beta_{min}, \beta_{max}]$	4.0 – 12.0%
Annual bank deposit interest rate	$[\beta'_{min}, \beta'_{max}]$	0.5 β
Weather report interval		1 hour
Weather forecast interval		1 hour
Weather forecast horizon		24 hours

$$P(K > k') \leq \omega \Rightarrow p \geq 1 - \sqrt[k'+1]{\omega} \quad (24)$$

If the probability that the game ends after 60 days (timeslot 1440 - timeslot 1320), $k' = 120$, is to be no more than 1%, $\omega = 0.01$, then the timeslot ending probability should be set at $p \geq 1 - \sqrt[121]{0.01} = 0.037$. The choice of p will be operationalized as a random drawing from a uniform distribution defined on the domain $[p_\omega, 1]$, where p_ω refers to the probabilities calculated before; for example, $p_{0.01}$ would be 0.037. Given the random end of game and that each Power TAC day

lasts 120 seconds in real time, an average Power TAC game will last around 2 hours overall.

7.3 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, an additional set of metrics (including the following) will be monitored throughout the game. These metrics are used by the game viewer to provide a visual representation of the game as it proceeds, and are stored within the game logs for post-mortem analysis.

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

7.4 Winner determination

Within a competition the performance of its participants has to be evaluated and compared at a certain point in time. This is usually accomplished by rank ordering all participants according to one or more defined performance criteria and to declare the best performer in this rank order winner of the competition. This principle also applies to Power TAC; albeit with quite some differences compared to previous TAC competitions. Consequently this section describes 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.

7.4.1 Performance criteria

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

$$\pi_{b,c,g} = pay_{b,c,g} - cost_{b,c,g} - fee_{b,c,g} \quad (25)$$

- **Payments** are monetary transfers from customers (consumer) to brokers and are based on the agreed contract conditions and the actual (ex-post) measured energy consumptions of the respective customer (consumer) as described in Section 6.1. 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 as described in Section 6.1. Other costs for instance include procurement in the wholesale market.
- **Fees** are (i) the cost for external balancing power (see Section 6) used, (ii) power distribution fees (in €/KWh) levied by the DU for power delivered to customers, and (iii) a carbon tax. The carbon tax is a fixed fee (in €/MWh) for each MWh of energy produced from non

renewable energy sources. The carbon tax remains constant throughout a competition and is publicly announced ahead of the start of the first round. Other fees for instance include publishing or revoking tariff.

7.4.2 Final ranking algorithm

After each competition round ends, e.g. at the end of the finals, z -scores of the accumulated profits for each broker are calculated to facilitate comparisons between one competition and another, i.e. between the 2-player, 4-player, and 8-player competition. If we denote the accumulated profits of a broker in a competition as $\pi_{b,c}$, the average accumulated profits of all brokers in the competition as $\bar{\pi}_c$ and the standard deviation of all brokers in the competition as S_c , then the standardized accumulated profits of broker b in competition c , $z_{b,c}$, is obtained as:

$$z_{b,c} = \frac{\pi_{b,c} - \bar{\pi}_c}{S_c}, \quad (26)$$

where

$$\bar{\pi}_{b,c} = \sum_{g=1}^{N_{b,c}} \pi_{b,c,g}, \quad (27)$$

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

After all competitions C have ended, an overall measure of relative broker performance will be obtained by summing over the standardized broker performance per competition:

$$z_b = \sum_{c=1}^C z_{b,c} \quad (28)$$

where C is the number of competitions.

7.4.3 Tournament structure

A typical Power TAC tournament consists of several rounds. Each competition, i.e. 2, 4, and 8-player games, has the following setup:

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.

Seeding Round This round will result in a ranking that is used to determine the broker pools for the quarter final. It might result in an elimination of brokers that don't perform according to the game specification or are purposely disruptive to other agents.

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.

Note: As opposed to previous TAC tournaments where the winner ranking was straightforward, i.e. after each round, agents in the top half of the performance ranking will proceed to the next round. In Power TAC we have three individual competitions (2, 4, and 8-player games) and the overall winner is the one agent with the highest overall accumulated z-score of all competitions (see Equation 28). For instance, an agent could reach only the quarterfinals in the 2-player competition, but takes second place in the 4-player competition, and first place in the 8-player competition, and still wins the overall tournament, since it has the highest accumulated z-score.

7.5 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.

- Information about external metrics and game logs are not provided to a broker directly, and agents should 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 agent. As such, according to the existing competition rules, it is forbidden within any specific round during the competition. Tuning with any available data on the other hand is allowed between the different tournament rounds.
- Data that agents discover on their own during a game can be used to fine-tune their agent in games within a round.
- Collusion is not allowed between the different agents.
- To discourage anti-competitive collusion, no team is allowed to enter the competition with two different agent identities.
- For efficient tournament scheduling, each team must be able to run two copies of their agent at any time in the tournament, since agents are required to participate in different pools at the same time.

8 System architecture

8.1 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 agents and enter them in the competition. The competition systems architecture is shown in Figure 13.

The tournament configuration is intended to support multi-round tournaments, with large numbers of visualizers. The administration portion of the web application supports tournament scheduling and access to records of past games. The web-app also serves as a proxy to allow visualizers access to running games on potentially several simulation servers.

A single web app can control multiple servers on multiple hosts, by storing game configuration in a shared database and then starting a server on a remote host, or notifying a running server of a game configuration that is ready to run. Weather and market price data will be served by remote

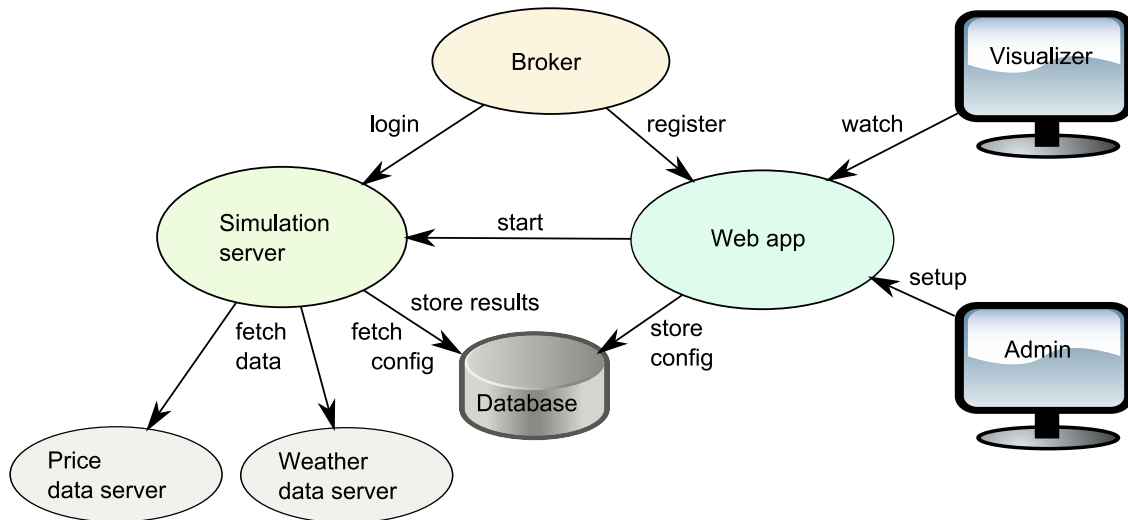


Figure 13: Competition systems architecture.

services, hosted on their own databases. The shared database will hold 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.

8.2 Research deployment

After the competition, teams are encouraged to release their agent code, so all teams can design and run their own experiments using a range of broker behaviors and market design details. The research systems architecture is shown in Figure 14. The results are published, and teams incorporate new insights into their agent designs for the following year.

The goal of the research configuration is to support development of agents 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 must meet the following requirements:

- Single-simulation setup from a simple web interface.
- Optionally allow agent login without credentials.
- Visualizer support for at least one browser.

Figure 14 shows the components of this 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 the weather and price services from a tournament installation. The game data is dumped to a flat file at the conclusion of each game.

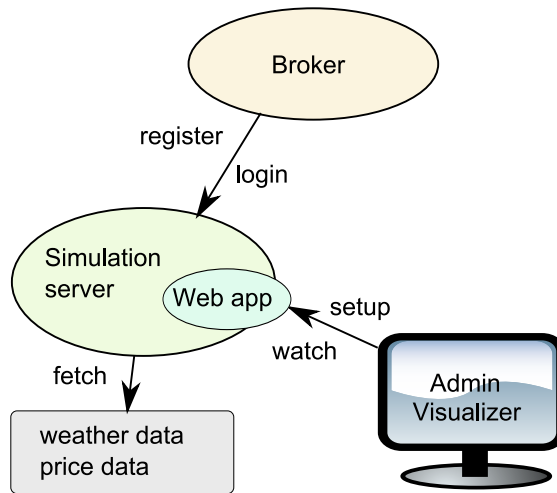


Figure 14: Research systems architecture.

Brokers may optionally log into the simulation server directly, without authentication. Otherwise, the web app will perform the authentication as in the tournament setup, and pass back credentials for access to the simulation server. 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] M. Amin and B. Wollenberg. Toward a smart grid: Power delivery for the 21st century. *IEEE Power & Energy Magazine*, 3(5):34–41, 2005.
- [2] M. Bichler, A. Gupta, and W. Ketter. Designing smart markets. *Information Systems Research*, 21(4):688–699, 2010.
- [3] 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.
- [4] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [5] R. Cavallo. *Social Welfare Maximization in Dynamic Strategic Decision Problems*. PhD thesis, Harvard University, 2008.
- [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 Weerdt, 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 –, Charlotte, November 2011.

- [9] Technology options for the near and long term. Technical report, U.S. Department of Energy Climate Change Technology Program, September 2006.
- [10] 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.
- [11] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [12] C. Hirsch, L. Hillemacher, C. Block, A. Schuller, and D. Möst. Simulations in the smart grid field study MeRegio. *it - Information Technology*, 52(2):100–106, 2010.
- [13] P. R. Jordan, B. Cassell, L. F. Callender, and M. P. Wellman. The ad auctions game for the 2009 trading agent competition. Technical report, University of Michigan, Department of Computer Science and Engineering, 2009.
- [14] P. Joskow and J. Tirole. Retail electricity competition. *The Rand Journal of Economics*, 37(4):799–815, 2006.
- [15] W. Ketter, J. Collins, and C. Block. Smart grid economics: Policy guidance through competitive simulation. Technical Report ERS-2010-043-LIS, RSM Erasmus University, Rotterdam, The Netherlands, 2010.
- [16] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [17] C. Möller, S. Rachev, and F. Fabozzi. Balancing energy strategies in electricity portfolio management. *Energy Economics*, 33(1):2–11, 2011.
- [18] S. S. Oren, S. A. Smith, and R. B. Wilson. Multi-product pricing for electric power. *Energy Economics*, 9(2):104–114, 1987.
- [19] A. Siddiqui, E. Bartholomew, and C. Marnay. Empirical analysis of the spot market implications of price-elastic demand. Berkeley Lab Publications LBNL-56141, Lawrence Berkeley National Laboratory, July 2004.
- [20] K. Skytte. The regulating power market on the Nordic power exchange Nord Pool: An econometric analysis. *Energy Economics*, 21(4):295–308, 1999.
- [21] 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, November 2008.
- [22] K. Spees and L. Lave. Impacts of responsive load in PJM: Load shifting and real time pricing. *The Energy Journal*, 29(2):101–122, 2008.
- [23] V. Tamma, S. Phelps, I. Dickinson, and M. Wooldridge. Ontologies for supporting negotiation in e-commerce. *Engineering Applications of Artificial Intelligence*, 18(2):223–236, March 2005.
- [24] 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.

- [25] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- [26] D. von Dollen. Report to NIST on the smart grid interoperability standards roadmap. Technical Report SB1341-09-CN-0031, Electric Power Research Institute (EPRI), June 2009.

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. power flows within the region are unconstrained. Local distribution grids are typically overdimensioned 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. Power distribution and transformation losses are ignored. In Germany these losses are estimated at 3%; for North America they are estimated at 5,5% [9]. 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 power 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 "timeslots" 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 timeslot is not taken into account. This means for example that balancing power demand for a timeslot is calculated as the difference of the sum of generation and the sum of consumption for that timeslot 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, Antonios Chrysopoulos, Jan van Dalen, Travis Daudelin, David Dauer, Joseph D'Costa, Josh Edeen, Ryan Finneman, Sebastian Gottwalt, Puthyrak Kang, Yixin Lu, Adis Mustedanagic, Nguyen Nguyen, Erik Onarheim, Markus Peters, Vedran Podobnik, Philippe Souza Moraes Ribeiro, Daniel Schnurr, Alexander Schuller, Andreas Symeonidis, Konstantina Valogianni, and Manuela Veloso.

Publications in the Report Series Research* in Management

ERIM Research Program: “Business Processes, Logistics and Information Systems”

2011

Sequencing Heuristics for Storing and Retrieving Unit Loads in 3D Compact Automated Warehousing Systems

Yugang Yu and René B.M. De Koster

ERS-2011-003-LIS

<http://hdl.handle.net/1765/22722>

A Local Search Algorithm for Clustering in Software as a Service Networks

Jelmer P. van der Gaast, Cornelius A. Rietveld, Adriana F. Gabor, and Yingqian Zhang

ERS-2011-004-LIS

<http://hdl.handle.net/1765/22723>

Implementing Standardization Education at the National Level

Henk J. de Vries

ERS-2011-007-LIS

<http://hdl.handle.net/1765/22812>

The Power Trading Agent Competition

Wolfgang Ketter, John Collins, Prashant Reddy, and Christoph Flath

ERS-2011-011-LIS

<http://hdl.handle.net/1765/23263>

Real-time Tactical and Strategic Sales Management for Intelligent Agents Guided By Economic Regimes

Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta, and Paul Schrater

ERS-2011-012-LIS

<http://hdl.handle.net/1765/23339>

Knowledge Sharing in Non-Knowledge Intensive Organizations: When Social Networks do not Matter?

Joey van der Capellen, Otto Koppius, and Koen Dittich

ERS-2011-013-LIS

<http://hdl.handle.net/1765/23489>

Towards a Value-based Method for Risk Assessment in Supply Chain Operations

Lingzhe Liu, and Hennie Daniels

ERS-2011-014-LIS

<http://hdl.handle.net/1765/23492>

Beyond Waste Reduction: Creating Value with Information Systems in Closed-Loop Supply Chains

Otto Koppius, Ozgur Ozdemir, and Erwin van der Laan

ERS-2011-024-LIS

<http://hdl.handle.net/1765/26892>

The Power Trading Agent Competition

Wolfgang Ketter, John Collins, Prashant Reddy, Christoph Flath, and Mathijs de Weerd

ERS-2011-027-LIS

<http://hdl.handle.net/1765/30683>

* A complete overview of the ERIM Report Series Research in Management:

<https://ep.eur.nl/handle/1765/1>

ERIM Research Programs:

LIS Business Processes, Logistics and Information Systems

ORG Organizing for Performance

MKT Marketing

F&A Finance and Accounting

STR Strategy and Entrepreneurship