The 2012 Power Trading Agent Competition

Wolfgang Ketter, John Collins, Prashant Reddy,
and Mathijs de Weerdt

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2012-010-LIS
Publication	September 2012
Number of pages	42
Persistent paper URL	http://hdl.handle.net/1765/37192
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

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS

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 time slot.
Free Keywords	autonomous agents, electronic commerce, energy, preferences, portfolio management, power, policy guidance, sustainability, trading agent competition
Availability	The ERIM Report Series is distributed through the following platforms: Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal Social Science Research Network (SSRN), SSRN ERIM Series Webpage Research Papers in Economics (REPEC), REPEC ERIM Series Webpage
Classifications	The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems: Library of Congress Classification, (LCC) LCC Webpage Journal of Economic Literature, (JEL), JEL Webpage ACM Computing Classification System CCS Webpage Inspec Classification scheme (ICS), ICS Webpage

The 2012 Power Trading Agent Competition

Wolfgang Ketter, John Collins[a],
Prashant Reddy[b], Mathijs de Weerdt[c]

September 2012
ERS-2012-010-LIS

[a]Computer Science Department, University of Minnesota, USA
[b]Machine Learning Department, Carnegie Mellon University, USA
[c]Delft 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 time slot.

The simulation environment models a wholesale market, a regulated distribution utility, and a population of energy customers, situated in a real location on Earth during a specific period for which weather data is available. The wholesale market is a relatively simple call market, similar to many existing wholesale electric power markets, such as Nord Pool in Scandinavia or FERC markets in North America, but unlike the FERC markets we are modeling a single region, and therefore we 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 . . . 9
  - 3.1.3 Capacity controls . . . 9
  - 3.1.4 Wholesale market trading . . . 10
  - 3.1.5 Portfolio management . . . 11
- 3.2 Information available to brokers . . . 12

4 Customer market **13**
- 4.1 Choosing tariffs . . . 15
  - 4.1.1 Deriving tariff utility . . . 15
  - 4.1.2 Choosing based on tariff utility . . . 16
- 4.2 Providing interruptible capacity . . . 17
- 4.3 Generating meter readings . . . 17

5 Wholesale market **18**
- 5.1 Trading and time slots available for trade . . . 18
- 5.2 Market clearing . . . 18
- 5.3 Wholesale suppliers and buyers . . . 20

6 Balancing mechanisms **20**
- 6.1 Scenario I: no controllable capacities . . . 22
- 6.2 Scenario II: static with controllable capacities . . . 22
- 6.3 Scenario III: dynamic with controllable capacities . . . 26

7 Competition format and interaction **26**
- 7.1 Competition initialization and Default Broker . . . 26
- 7.2 Competition ending . . . 28
- 7.3 External metrics and game logs . . . 29
- 7.4 Winner determination . . . 29
  - 7.4.1 Performance criteria . . . 30
  - 7.4.2 Final ranking algorithm . . . 30
  - 7.4.3 Tournament structure . . . 31
- 7.5 Competition rules . . . 31

8 System architecture **32**
8.1 Tournament deployment 32
8.2 Research deployment 32

A Assumptions **35**

B Acknowledgements **36**

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 [11]. 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 [14], but these studies are not able to explore the full range of unanticipated self-interested or destructive behaviors of the participants. Smart grid pilot projects, on the other hand, are limited in their ability to test system dynamics for extreme situations. They also lack the competitiveness of open markets, because a single project consortium typically controls and optimizes the interaction of all parts of the pilot regions. Therefore, we 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*[1] 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) [6] and the Trading Agent Competition for Ad Auctions (TAC AA) [10].

2 Competition overview

The major elements of the Power TAC scenario are shown in Figure 1. Competing teams will construct trading agents to act as self-interested "brokers" that aggregate energy supply and demand with the intent of earning a profit. In the real world, brokers could be energy retailers, commercial or municipal utilities, or cooperatives. Brokers will buy and sell energy through contracts with retail customers (households, small and medium enterprises, owners of electric vehicles), and by trading in a wholesale market that models a real-world market such as the European or North American wholesale energy markets. 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

---

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

[2] See 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.

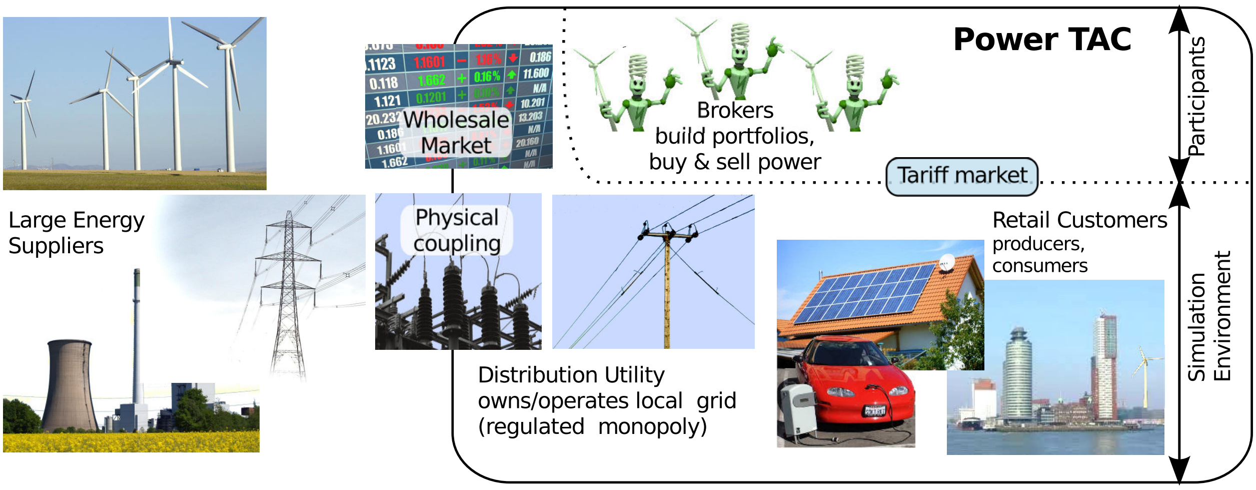

Figure 1: Major elements of the Power TAC scenario.

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 contracts[3]. 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.

---

[3] Individual contract negotiation will be implemented for the 2013 competition.

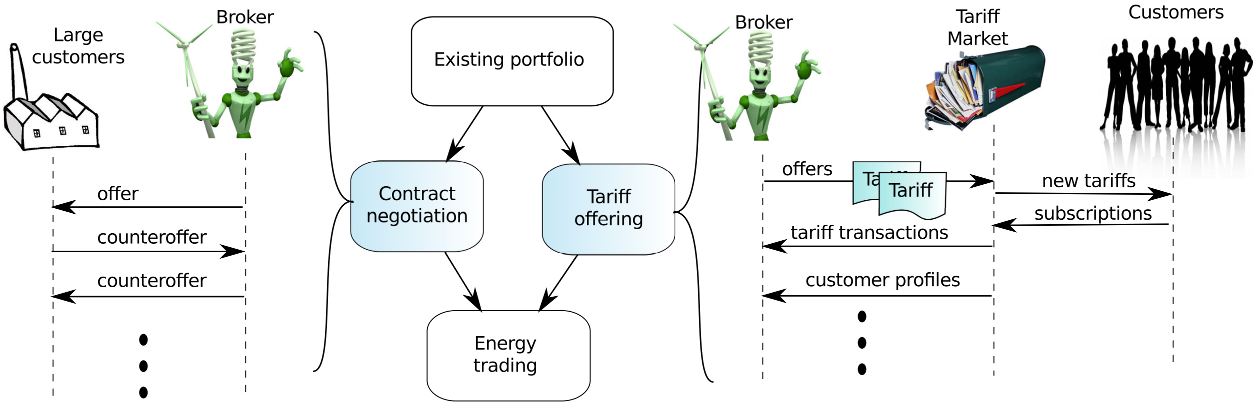

Figure 2: Portfolio management process. Tariff offerings proceed in parallel with individual contract negotiation.

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

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

2.2 Customer market

In the customer market, broker agents try to acquire energy generation capacity from local producers, and load capacity from local energy consumers. Brokers can buy and sell energy through two different mechanisms, *tariffs* and *individual contracts* (although individual contracts will likely not be implemented before the 2014 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 [17].

Within the Power TAC domain, negotiations and the contracts (including tariffs) that are the subject and result of negotiations are 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). Contracts or tariffs may also specify amounts of energy that can be remotely controlled or curtailed, 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 may specify payments to or from the customer based on time (one-time sign-up fee or bonus, fixed monthly distribution fees), or time and energy (fixed or variable prices for a kilowatt-hour).

Communication: contract award and termination, notification of price changes, 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 through 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 an incentive for brokers to balance their own portfolios of energy supply and demand in each time slot.
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, wholesale market settlements, balancing 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 time slot. This information is needed by the Distribution Utility to run the balancing process in the current time slot.

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 time slot. 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

$$b_{d+1} = (1 + \beta/365)b_d + credits_d - debits_d \tag{1}$$

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

$$b_{d+1} = (1 + \beta'/365)b_d + credits_d - debits_d \tag{2}$$

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 2 on page 28 for standard tournament values).

2.6 Weather reports

Weather forecasts and current-hour weather conditions are sent to brokers in each time slot. Some customer models will use this information to influence energy consumption (temperature, for example), and production (wind speed, cloud cover). Brokers 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 privileged information, 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 time slot. Note that the specific order of events is more flexible than what is shown. Specifically, the sequence of major processes in the simlation 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 time slot, 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 for an existing 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 time slots.

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,

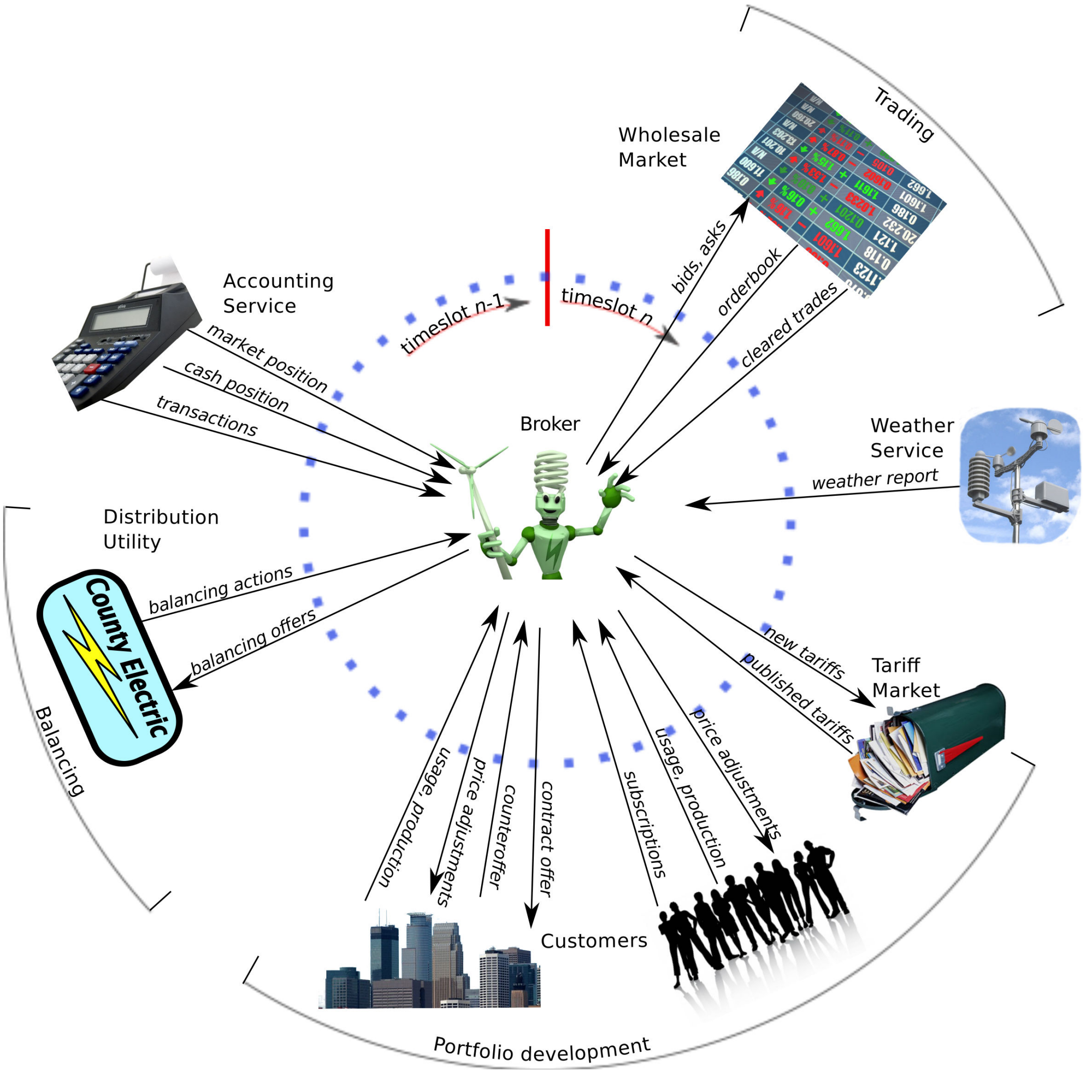

Figure 3: Overview of Power TAC activities within one time slot. A broker interacts with the wholesale and tariff markets, and receives information from the weather service, customers, the balancing market, and the accounting service.

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:

- 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;

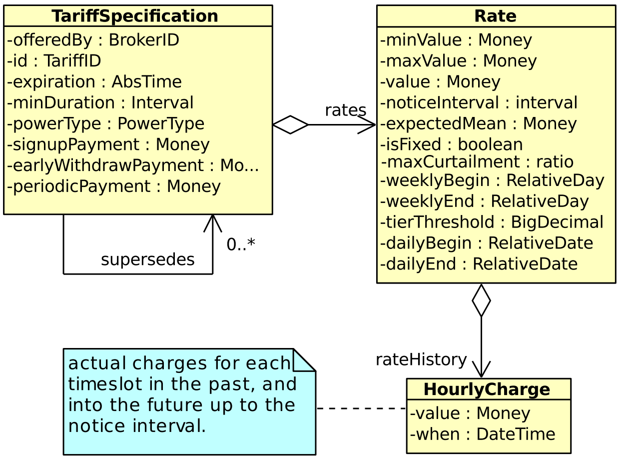

Figure 4: Tariff structure.

- two-part tariffs (fixed daily fee plus usage fee);
- signup payments in either direction (fee or bonus);
- early withdrawal penalties;
- variable rates with minimum and maximum values, estimated mean values, and notice intervals.
- interruptible rates in which some portion of the load may be curtailed during a timeslot in order to reduce overall power costs or to reduce the cost of balancing.

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.

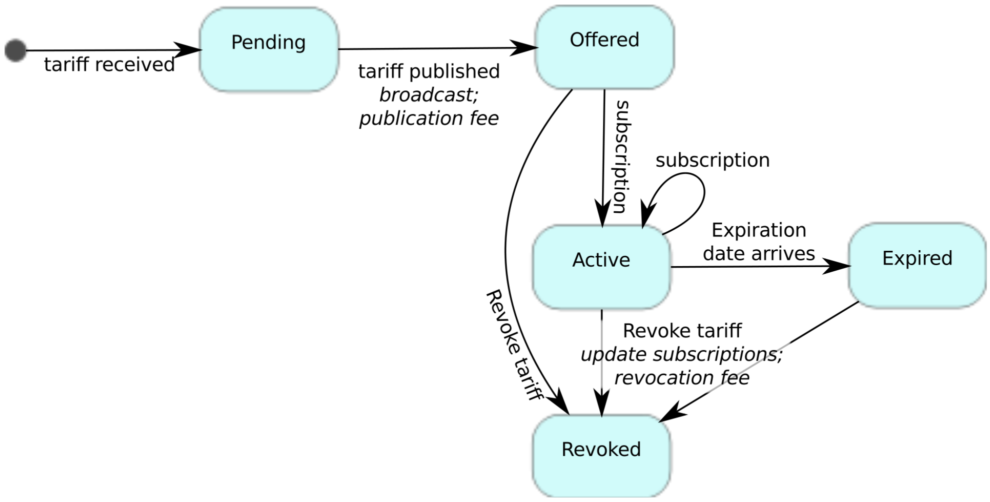

Figure 5: Tariff state transitions.

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 time slots before the time slot to which they apply, the broker must use some type of forecasting to determine the optimal price to set for the target time slot, i.e., the future time slot 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 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 time slot. Two major factors in the determination of this demand and supply are (i) the estimated or realized load and supply for time slots preceding the target time slot, and (ii) the weather forecast conditions for the target time slot.

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 time slot. 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 time slot.

3.1.3 Capacity controls

Brokers may be motivated to offer tariffs for controllable (also known as *interruptible* or *curtailable*) capacity for two reasons:

- to reduce wholesale power costs, a broker may directly exercise *economic controls* for a specific timeslot, up to the limit of the `maxCurtailmentRatio` specified in the rate that is in effect for a given tariff. An economic control is communicated by an `EconomicControlEvent` that specifies a curtailment ratio r and a timeslot n, and must be received by the simulation server before the customer models run in timeslot n. Since the customer models run close to the start of a timeslot, the broker should communicate economic controls before the end of timeslot $n-1$. These controls are for specific timeslots, so a broker must re-issue them to extend such controls across multiple timeslots.

- to reduce balancing charges, a broker may authorize the DU to exercise controls against its tariffs during the balancing phase, just in case doing so would be beneficial to the broker. Such controls are called *balancing controls*. Brokers may issue `BalancingOrder` messages to the DU in order to authorize these controls, specifying the tariff, an allowable curtailment ratio, and a price/kWh. The price is typically positive for consumption curtailment (the DU pays the broker), and negative for production curtailment. Balancing orders remain in effect until canceled by issuing a new order specifying a different curtailment ratio.

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

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

3.1.4 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.

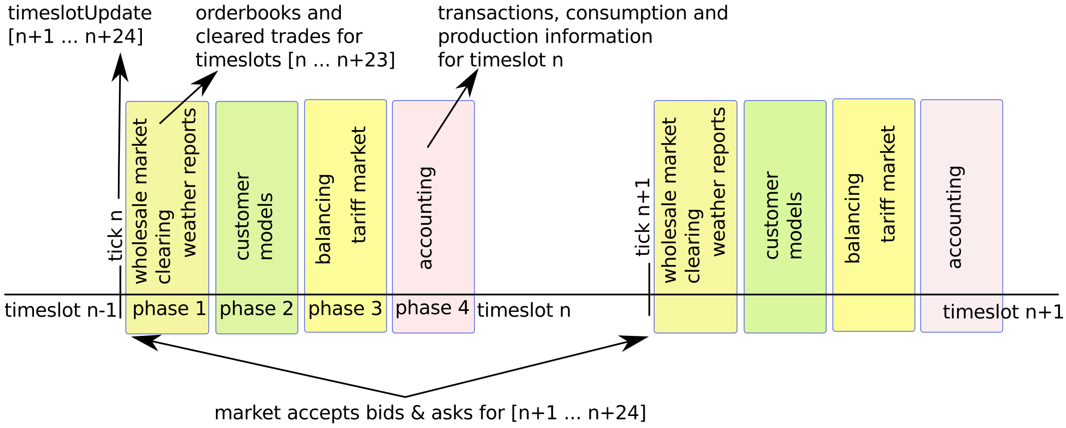

Figure 6: Simulation process phases and associated market information.

The wholesale market is cleared at the beginning a time slot n. The process starts with an announcement of the time slots open for trading in the following time slot, typically time slots

$[n+1 \ldots n+24]$. Next, all outstanding orders that have been submitted since the beginning of the previous time slot $n-1$ are cleared, and the results announced in the form of cleared trades (amounts and prices) and orderbooks (uncleared bids and asks) for each cleared time slot. From the broker's perspective, the information it needs to make trading decisions for future time slots starts at the beginning of a time slot. 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 arrive in time to make final trading decisions for the following market clearing.

3.1.5 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 time slots. 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.

Note that storage capacity is not implemented for the 2012 competition.

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 time slot.

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

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

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

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

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

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

4 Customer market

Consumers and producers in the customer market are simulated using a range of *customer models*. These customers interact with brokers primarily through the tariff market mechanism – by subscribing to tariffs offered by brokers. In future, larger customers will also be able to negotiate individual contracts. Each customer is fundamentally characterized by a core set of information captured in a corresponding `CustomerInfo` object, which includes:

- **Name**: The mnemonic handle for a customer, separate from the internally generated unique ID for each customer.
- **Population**: An integer count of the number of indivisible entities represented by this "aggregate" customer. This typically corresponds to the number of metering endpoints deployed by the DU to service this customer. For example, if a customer model represents a single household, it would likely have a population of 1 even though multiple persons occupy the household. If a model represents an office building, it may have each tenant or each floor of the building as a separate population entity.

- **PowerType**: This is an important characterization that primarily indicates where a customer *consumes* or *produces* power. It also indicates whether that consumption or production is *interruptible*; i.e., the consumption or production capacity can be taken offline directly by the DU in response to economic controls exercised by brokers or due to balancing controls that the DU is authorized to exercise. This field is realized in code as a sophisticated enumeration-like object that can be interrogated for the above information and it also captures subtype information such as whether a production customer is a solar producer or a wind producer.

- **MultiContracting**: Customers with non-singular populations may have the ability to allocate a partition of the population over multiple tariffs, which may be offered by multiple brokers. Note however that all entities of the population must be allocated to some tariff at any given point in the simulation.

- **CanNegotiate**: This field is a placeholder for future enhancement; it indicates whether a customer is allowed to negotiate individual contracts.

Thus, the currently available customer models vary along the three key dimensions of population size, power type and ability to multi-contract. In implementation, the customer models are broadly one of two classes:

1. *Elemental Models*: This class of models attempts to simulate customer behavior at a fine level of granularity. For example, such customers are modeled using the number of persons per household, their work/vacation schedule, the usage patterns of the individual appliances that they use, and so on [9]. Two such models are currently available representing households and office buildings (respectively in the `household-customer` and `officecomplex-customer` software modules).

2. *Factored Models*: The fine granularity of the behavioral simulations employed by the elemental models constrains the size of populations that can be simulated by such models. As an alternate approach, factored models simulate the aggregate behavior or larger population models and other complex entities using a generalized set of *factors* that influence their behavior. Such factors control both the tariff selection process and the consumed/produced capacities exhibited by such customers. Thoughtfully configured combinations of all of these factors can be used to instantiate specific customer types such as relatively homogeneous collections of households, offices, campuses, hospitals, factories, wind farms, solar farms, etc. Information about the specific instantiations of factored models currently available can be found in the online documentation and in the `factored-customer` module[4].

In a research environment, one can choose which of these customer models are deployed in the simulation and how they are configured, but in the competition setting this information is opaque to the brokers, i.e., the competition participants. The rest of this section describes the general behavior of both classes of customer models. Implementation variances result in slight differences, which will be highlighted as necessary.

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

---

[4] https://github.com/powertac/powertac-server/wiki/Factored-Customers

4.1 Choosing tariffs

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 or gain over the lifetime of a contract relationship. This quantity is composed of (i) the expected variable payments from estimated consumption or production, (ii) fixed periodic payments, and (iii) sign-up fees or bonuses. Since early exit from contracts is allowed (possibly with a penalty), customer models may evaluate available tariffs at any time. In this case, a proper switching evaluation has to consider the early 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, and realized price of variable-rate tariffs. Therefore, tariffs are compared using a utility value computed from the monetary implications and these other aspects. 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

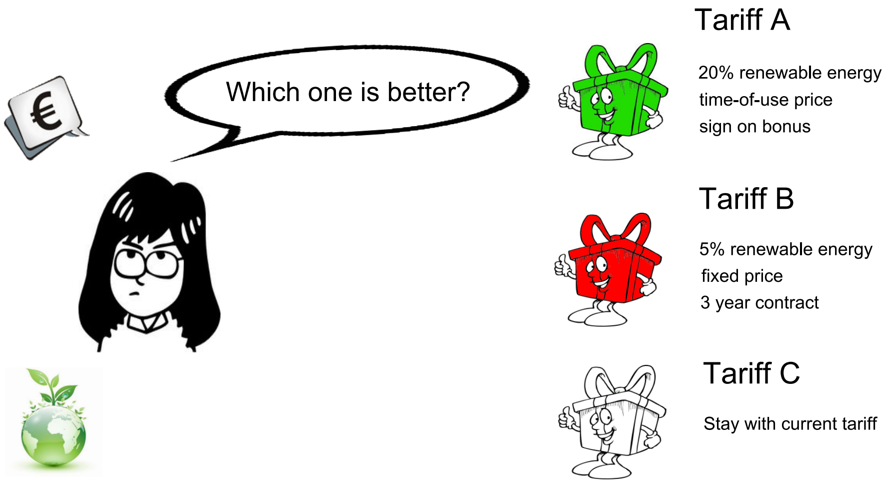

Figure 7: Tariff selection problem.

Note that customers also exhibit inertia which prevents them from immediately switching to the tariff(s) with the best utility. The implementation of the tariff selection problem is described in the remainder of this section.

4.1.1 Deriving tariff utility

The utility of a given tariff, T_i, is computed as a function of variable payments p_v, periodic payments p_p, lifecycle payments p_l (sign-up fees and bonuses), realized variable payments p_r, and the interruptibility characteristics of the tariff x_i:

$$u_i = f(p_v, p_p, p_l, p_r, x_i)$$

The specifics of the function f vary across customer model implementations, but some general characteristics of this utility evaluation mechanism are described here. Broker reputation is currently measured in terms of how the realized price on tariffs compares to the expected costs communicated by the tariff (the distinction is mostly relevant for tariffs with variable rates). Whenever realized price information is available (once the tariff has been active for some period), the expected variable payments for a tariff are computed as a linear combination of the communicated variable payments p_v and the historical realized payments p_r: i.e., $\alpha_v p_v + (1 - \alpha_v) p_r$.

The computation of the various payment components is normalized to nominal daily values. Thus, variable payments are computed using 24-hour forecast samples of consumption or production capacities for each customer. Thus, it is possible for different tariffs to obtain different utility value rankings purely from the variance in these forecast samples. Periodic payments are denoted as daily payments already. Sign-up fees and bonuses are normalized to daily values using expected length of typical contract periods. Finally, a customer-specific discount may be applied for a tariff that has the option to exercise interruptibility (for customers with interruptible power types).

4.1.2 Choosing based on tariff utility

The set of tariffs considered for evaluation at any given time may be constrained. For example, tariffs with expected costs that deviate "too far" from the default tariff may be disregarded. Similarly, a "flood of tariffs" from a single broker may be filtered down such that only the best of them is considered. At various points customers may evaluate new tariffs only in comparison to their currently subscribed tariffs whereas at other times they may consider all unexpired tariffs.

Customers may also decide to postpone evaluation of newly published tariffs at some times. So, each customer model is configured with an *inertia threshold*, $\iota \in \mathbb{R}[0, 1]$. At each tariff evaluation timeslot, τ, if a random draw from the uniform distribution, $x_\tau \sim \mathit{Unif}(0, 1)$ is at or below the inertia threshold, *i.e.*, $x_\tau \leq \iota$, then the newly published tariffs are queued for deferred processing at a subsequent tariff evaluation timeslot. On the other hand, if the random draw exceeds the inertia threshold, *i.e.*, $x_\tau > \iota$, then any newly published tariffs and any previously deferred tariffs are evaluated immediately.

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

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

The parameter λ is a measure for how rationally a customer chooses tariffs: $\lambda = 0$ represents random, irrational choice, while $\lambda = \infty$ represents perfectly rational customers always choosing the tariff with the highest utility[5]. 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.

---

[5] In implementation, λ is less than ∞ to avoid numeric overflow issues.

4.2 Providing interruptible 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.

- **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.

When capacity is curtailed by the DU due to an economic control exercised by the broker or due to a balancing control exercised by the DU, the customer may forfeit that capacity or shift some or all of it to future time slots. The degree and nature of shiftability is a customer-specific attribute that is deeply tied to the physical nature of that customer's capacity.

4.3 Generating meter readings

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

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

5 Wholesale market

The wholesale market in Power TAC operates as a periodic double auction (PDA) and represents a traditional energy exchange like NordPool, FERC, or EEX[6]. The brokers can buy and sell power contracts for future time slots 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 time slots available for trade

Brokers can submit orders to the wholesale market for delivery between one and 24 hours in the future. The time slots available for trading are marked as "enabled"; changes in time slot status are communicated to brokers at the beginning of each time slot. Orders submitted for non-enabled (disabled or not yet enabled) time slots are silently discarded. 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 time slot s, an amount of energy e in megawatt-hours, and optionally a limit price per megawatt-hour p. Energy and price quantities are treated as proposed debits (negative values) and credits (positive values) to the broker's energy and cash accounts. So an order $(b_1, s_{12}, 4.2, -21.0)$ represents a bid (a buy order) from broker b_1 to acquire 4.2 MWh of energy in time slot s_{12} for at most 21 €/MWh. Orders that specify a limit price p are called "limit orders", while orders that do not specify a limit price are called "market orders."

5.2 Market clearing

When the simulation clock is advanced to a new time slot, the wholesale market clears the orderbook for each of the enabled time slots. Note that at the beginning of the clearing process an updated list of enabled time slots 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 time slots from the broker's viewpoint differs from the set of enabled time slots from the market's viewpoint.

In the clearing process, as shown in Figure 8, demand and supply curves are constructed from bids and asks to determine the clearing price of each orderbook (one for each enabled time slot) 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.

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

---

[6] See http://www.nordpoolspot.com, http://www.ferc.gov, or http://www.eex.com/en.

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

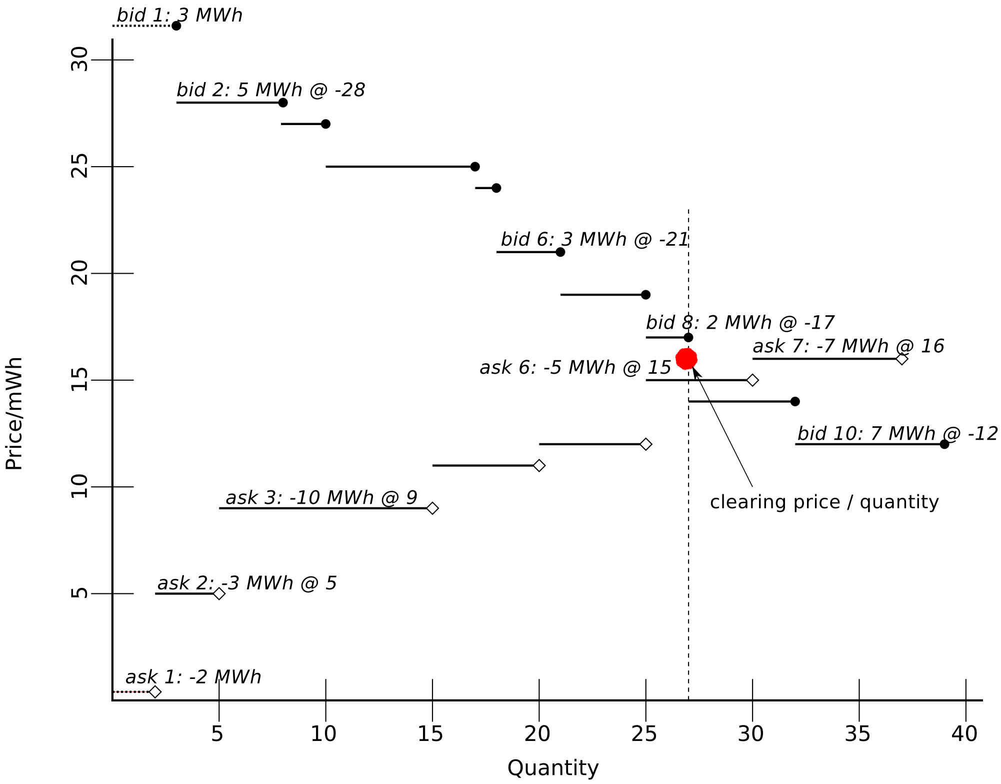

Figure 8: Market clearing example: bid 8 and part of ask 6 are the last to clear.

After the market is cleared the following steps are performed:

- Clearing price and volume are publicly broadcast (public information). In the example of Figure 8, this would be $(27, 16)$.
- Post-clearing orderbooks are published for each cleared time slot, giving the un-cleared bids and asks, without broker information. In the example, the orderbook would include two asks $((-3, 15), (-7, 16))$, and two bids $((5, -14), (7, -12))$.
- Brokers are informed about their own executed transactions (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 time slot 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

$$C_{g,s} = C_{g,s-1} + v(2\nu - 1)\hat{C}_g + vm(\hat{C}_g - C_{g,s-1}) \tag{3}$$

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 time slot that is farther in the future than its commitment leadtime τ_g. Once it has sold at least some power for a given time slot, it is committed, and will attempt to sell the remainder by continuing to submit asks in each enabled time slot, including those closer to the current time than its commitment leadtime. If it fails to sell at least some power in a given time slot by its commitment time, then it will withdraw its capacity from the market for that time slot.

Once a Genco has sold power for a given time slot, 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 $\nu \in [0, 1]$, it computes a price $p_{b,s} = -p_b \ln(1-\nu)$ for each time slot s and places a bid $(b_b, s, q_b/p_{b,s}, p_{b,s})$ in each open time slot. This exponential distribution produces large numbers of low-priced high-quantity bids, and a few higher-priced low-quantity bids.

6 Balancing mechanisms

In electricity markets, supply and demand have to be balanced almost perfectly in real time. A major task of the Independent Systems Operator (ISO)[7] on the wholesale (transmission) level and 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 [15]. 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 [12]. In Power TAC, brokers accumulate credits and debits to their energy

---

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

budgets for each time slot by selling (exporting) power or buying (importing) power in the wholesale market, and by the power consumption and production activities of their contracted customers. The total net energy budget for a time slot s and a broker b is denoted by $x_{b,s}$. To carry out its responsibility to balance supply and demand in each time slot, 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 9).

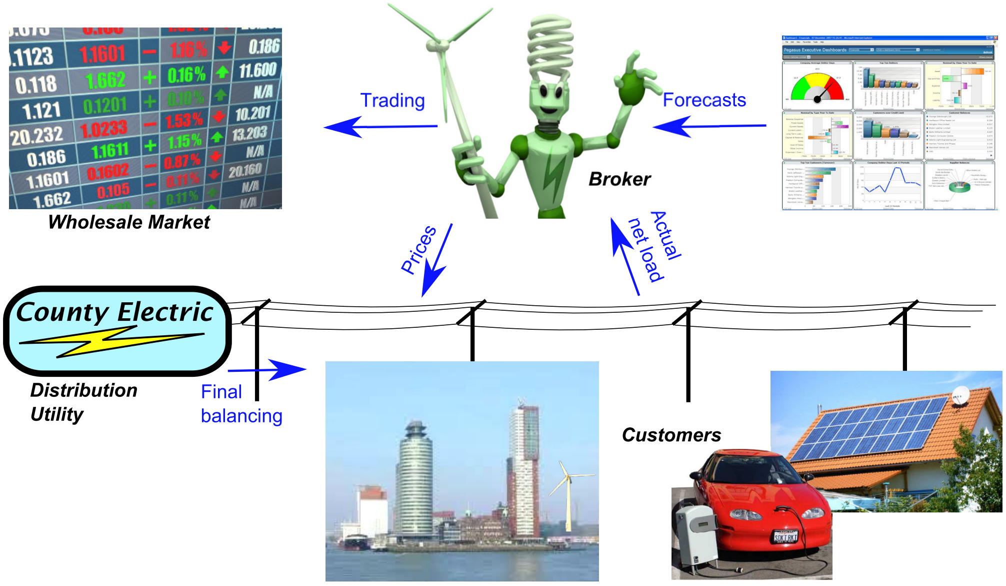

Figure 9: Entities and activities during balancing.

Brokers acquire balancing capacity by offering price concessions in exchange for the ability to remotely curtail loads (or sources) for limited periods of time. These loads are connected to controllers installed at a customer site that allow the DU to interrupt or modulate power flow for 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. In Power TAC regulating is currently (2012) limited to curtailment of loads.

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 is implemented for the 2012 competition.

Scenario III: dynamic with controllable capacities This may be implemented as an option in the 2013 competition.

More detailed background and examples on the balancing market can be found in [7]. In

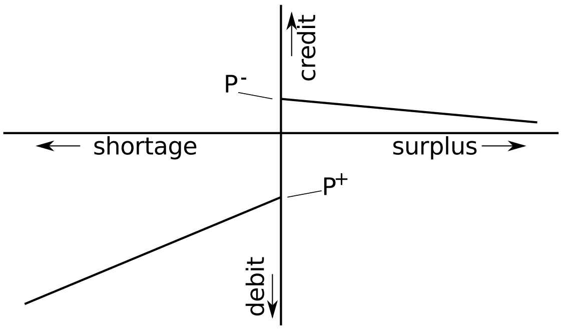

Figure 10: The price paid for an energy surplus is always lower than the lowest price observed in the day-ahead market, while the price charged for an energy deficit is always higher than the highest price observed in the day-ahead market.

the following two sections we discuss the mechanisms that are used for the first two scenarios (with/without controllable capacities).

6.1 Scenario I: no controllable capacities

Scenario I was used in the pilot competition in the summer of 2011, using data from the econometric analysis of [13] for parameter settings, scaled by the size of the customer population. This scenario is similar to a number of current real-world implementations.

In this scenario, the charge for an individual broker b is set equal to the costs (or profits) the DU incurs for balancing its individual imbalance quantity $x_{b,s}$. These $x_{b,s}$ values are the individual imbalance quantities, positive if the broker's market position is greater than the actual usage of its subscribed customers, and negative otherwise. In general the costs for the DU for balancing an amount x are given by $c_0(x)$. This represents procurement or sale of power in the regulating market. In Power TAC we let $c_0(x) = -x \cdot P^+(s)$ when $x < 0$ (for up-regulation) and $c_0(x) = -x \cdot P^-(s)$ when $x \geq 0$ (for down-regulation), where $P^+(s)$ is the highest price at which energy has been traded for time slot s, and $P^-(s)$ is the lowest. The charge of $c_0(x_{b,s})$ for broker b on time slot s ensures that it is never better to let your imbalance be resolved by the DU, and thus provide an incentive to focus on trading in the day ahead market.

6.2 Scenario II: static with controllable capacities

With controllable capacity, brokers can submit balancing orders that allow the DU to exercise capacity control in order to achieve balance. Each balancing order specifies a tariff, a ratio, and a price, and allows the DU to curtail subscribers to the tariff up to the specified ratio of their actual usage, for the stated price/kWh. Note, however, that there are several constraints on the amount of power available for curtailment: the rate currently in effect specifies a maximum curtailment ratio, and an economic control may have already been exercised against a particular tariff. Therefore, the available curtailment is the product of the unexercised ratio and the actual power usage of the customers against this tariff.

Brokers must submit their balancing orders before the customer models run, and the DU runs its balancing process after customer consumption is known for the current time slot. At this point, the DU can determine the actual quantities available for curtailment for each balancing order.

The DU acts to resolve the net imbalance over all brokers at minimal cost. To achieve this, given a set of balancing orders, the DU

1. discards the orders that cannot contribute to the solution; if overall balance is negative (up-regulation needed), then only consumption curtailment is used, and if overall balance is positive, then only production curtailment is used.
2. includes "dummy" orders with essentially infinite capacity that represent procurement or sale of power in the regulating market at costs of $c_0(x)$, a linear function of quantity. For up-regulation, $c_0(x) = P^+(s) \cdot x + \cdot \phi^+ x^2$, and for down-regulation $c_0(x) = P^-(s) \cdot x + \cdot \phi^- x^2$, where ϕ^+ and ϕ^- are the slopes of the cost functions for up-regulation and down-regulation respectively. Note that in case there are balancing orders with prices above P^+ or P^-, the dummy orders will be split around such balancing orders. Competition values for ϕ^+ and ϕ^- are given in Table 2.
3. Sorts the remaining orders by price, with the lowest first.
4. In price order, the DU selects the cheapest orders up to the required capacity. Note that in general, the most expensive order selected may only be partially exercised.
5. The price for a broker b depends then on both its own imbalance, as well as on its balancing orders. This computation is the sum of a VCG payment p_{vcg} [16, 5], and an imbalance payment p_{imb} as defined in more detail below.

The payment for brokers consists of two parts: a payment for the use of its controllable capacity p_{vcg}, and a payment for its imbalance p_{imb}. Both payments typically are negative (the broker pays) in case of being short or when selling downwards controllable capacity (e.g., curtail production), and positive (the broker is paid) when it has a surplus, or can curtail its consumption. In the 2012 competition, as mentioned above, only curtailment of consumption is supported.

The setting for choosing controllable capacity is very similar to a one-sided auction, and for this part the VCG payment is used. The VCG payment for controllable capacity is defined to be the marginal contribution of broker b: the difference in (declared) balancing costs for the other brokers for the remainder of the balance, and the balancing cost of the complete net imbalance without using b's controllable capacity. To compute this for a broker b, we compute the optimal combination of bids while leaving out broker b's bids, and compare this to the costs to the other brokers of the optimal combination using the orders of all brokers including b. Additionally, we resolve the following issues by the second part of the payment p_{imb}.

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

The cost for the DU is the sum of the VCG payments and the costs of exercising any dummy orders.

The idea of the imbalance payments is to let the brokers that contribute to the imbalance pay for both the costs of the DU as well as for the opposite imbalance other brokers may have (since that also reduces the balancing costs). Similarly to VCG, we remove the part that a broker can influence (in this case the costs of its own controllable capacity) from the equation. Denoting the set of orders for controllable capacity by C, and that of a broker b by $C_{b,s}$, the costs of the DU for a given net imbalance X (following from the VCG payments and possibly some dummy orders) is denoted by $DU_{costs}(C \setminus C_{b,s}, X)$. A broker b with a non-zero imbalance $x_{b,s}$ that does not contribute to the imbalance (i.e., $x_{b,s} \cdot X \leq 0$) then "pays" $\frac{DU_{costs}(C \setminus C_{b,s}, X)}{X} x_{b,s}$. The payment for a broker b that contributes to the imbalance is defined the same in case there are no non-contributing brokers with controllable capacity. However, when such non-contributing brokers (denoted by B) do exist, we must make sure that the payment for contributing brokers such as b is sufficient to cover the payment for all non-contributing brokers. To guarantee this, we exclude all balancing orders of brokers B in computing the costs for contributing brokers, so broker b pays $\frac{DU_{costs}(C \setminus \{C_{b,s} \cup_{k \in B} C_{k,s}\}, X)}{X} x_{b,s}$.[8]

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

The VCG prices ensure for a single (isolated) time slot that brokers cannot gain from pricing their orders higher or lower than their real costs (if nobody else changes its bids), and that they often gain (and never lose) from placing orders for curtailment if they have any. In other words, myopic brokers should bid their (estimated) actual costs for balancing capacity in the balancing market. With the second payment, if you expect other brokers to be (almost) balanced, it is better to be balanced as well.

The following example, graphically depicted in Figure 11, illustrates the balancing mechanism described above.

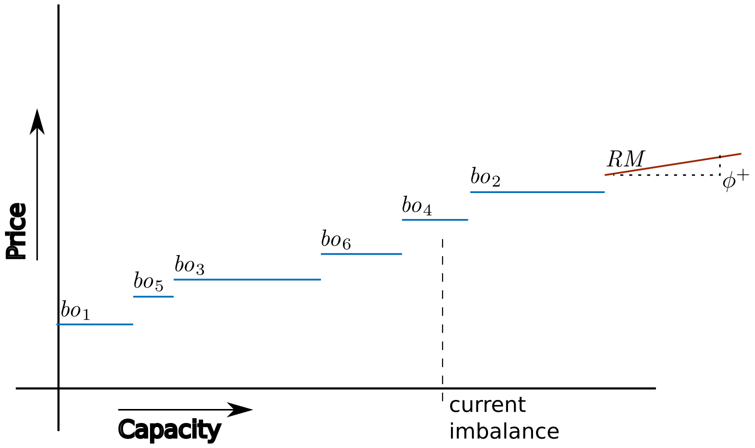

Figure 11: The balancing orders are ordered on price. Only capacity up to the current net imbalance is used.

Example 1. *Assume brokers A_0, A_1, A_2, and A_3 have imbalances of 0, +40, -80, and -140 kWh,*

---

[8] In the future, we hope to introduce variants of this mechanism that are a bit cheaper on the brokers, with the same guarantees.

respectively, for a total imbalance $X = -180kWh$*. We have six balancing orders* bo_1 *through* bo_6*, and a dummy order* RM*. The total imbalance falls within the range of one of the orders,* bo_5*. All orders with lower prices will be exercised, and* bo_5 *will be partially exercised. The signs in this example are from the standpoint of the brokers. This means that negative cash values represent payments from brokers to the DU, and negative energy values represent amounts the brokers have sold but not acquired, or amounts the brokers can consume by curtailing production.*

The next step is to set prices for each broker's balancing orders, using the VCG mechanism. For each broker that has orders to be exercised, we must discover the price that would have to be paid for its capacity if its orders were not in the mix. To see how this works, assume the orders are as follows: bo_1 *is (*A_0*, 35 kWh, 0.003/kWh);* bo_2 *is (*A_0*, 62 kWh, .0091/kWh);* bo_3 *is (*A_1*, 67kWh, .0051/kWh),* bo_4 *is (*A_1*, 30kWh, .008/kWh),* bo_5 *is (*A_2*, 20kWh, .0042/kWh);* bo_6 *is (*A_2*, 39kWh, .0062/kWh); and* RM *is (DU, xx kWh, .01/kWh,* $\phi^+ = 0.001/kWh$*). Sorted on the cost, we thus have the following balancing orders:* $bo1, bo5, bo3, bo6, bo4, bo2$ *(See Figure 12). To balance we need all of* $bo1, bo5, bo3, bo6$ *and only 19kWh of* $bo4$*.*

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

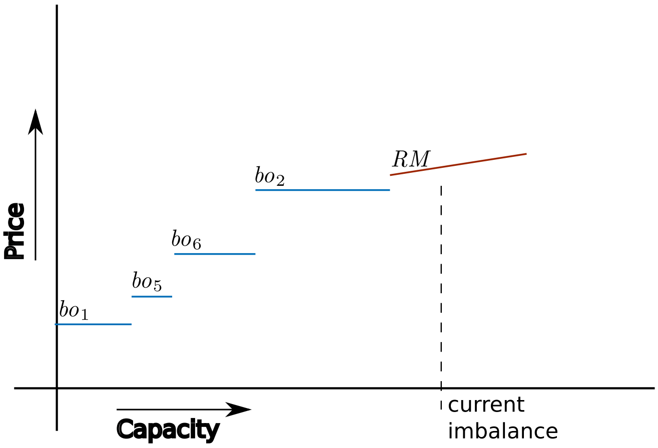

Figure 12: The price for balancing for broker A_1 is determined by the marginal cost in case all orders from A_1 are removed.

To compute the imbalance portion of the payment, we treat the contributing and non-contributing brokers separately. Since the overall balance is negative, the set of non-contributing brokers thus is just $\{A_1\}$*. Contributing brokers* A_2 *and* A_3 *pay their shares of the cost to the DU to resolve the total imbalance, assuming that there are no balancing orders from the non-contributors. For each broker, the hypothetical total cost paid by the DU is the sum of the VCG payments to the other contributors, plus the residual amount the DU would have to purchase from the regulating market; the broker's share is the product of the total cost and the ratio of its imbalance to the total imbalance. For e.g. broker* A_2 *we first compute the costs for the DU in case* bo_3 *and* bo_4 *from* A_1 *and* bo_5 *and* bo_6 *from* A_2 *are removed. For this we compute the VCG payments for all other brokers (*A_0 *and* A_4*) and sum these (*26.481 *for* A_0 *since otherwise everything is to be resolved by the RM;* A_4 *does*

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

Table 1: Broker payments for the example

not have controllable loads), add the cost of the remaining 83 kWh at the marginal rate of 0.093, i.e., 7.719, divide by the overall imbalance $X = -180$*, and multiply by the imbalance of* A_2 *of* -80*, giving a payment of 15.2. Broker* A_1 *is a non-contributing broker, so we remove* bo_3 *and* bo_4 *and compute the VCG payments for* A_0 *(15.035) and* A_2 *(6.903), and the DU cost for the extra 24 kWh from the regulating market (0.816) to get the hypothetical total DU cost (22.754) and* A_1*'s share of -5.056444.*

All payments are summarised in Table 1.

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

6.3 Scenario III: dynamic with controllable capacities

In the current (2012) version, the process of Scenario II is repeated every time slot. However, the incentive issues discussed above do not automatically transfer to a repeated game over multiple successive timeslots, because the shifting behaviors of the customers must be taken into account. There may not always be an incentive to be truthful about the price of the curtailment orders, because in the tariffs their use is typically limited. For example, when a certain load may only be used once every four time slots, and the broker expects the price of regulating services to go up, it may be wise to wait. 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 [4]. This extension is object of our current study.

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

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

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

In order to interpret the market prices in the bootstrap dataset, it is necessary to understand the bidding behavior of the default broker. The default broker estimates the net 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 time slot in a week. During the second through nth week, these cells contain the exponentially-smoothed ($\alpha = 0.3$) net consumption value for the customer in that time slot, counting from the start of a week. During the first week, it uses the actual consumption observed in the given hour h during the previous 24 hours, and during the first day it uses the usage observed in the previous time slot.

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

$$l_{prev} = \left\{ \begin{array}{lcl} l_{s,t-1} & : & \text{if order in previous time slot } t-1 \text{ did not clear} \\ l_{max} & : & \text{otherwise} \end{array} \right. \tag{4}$$

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

$$l_{s,t} = \max\left(l_{min}, 2\frac{l_{min} - l_{prev}}{s - t - 1}\right) \tag{5}$$

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

Table 2: Parameters used in Power TAC tournament games.

Parameter	**Symbol**	**Standard Game Setting**
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 time slots after time slot 1320 (start of day 55) before end of game	$[p_{min}, p_{max}]$	$[p_\omega, 1]$
Probability of game end for each time slot after time slot 1320 (start of day 55)	p	$\frac{1}{121}$
Minimum game length	Min(TS)	1320
Expected game length	E(TS)	1440
Timeslot length	τ	60 minutes
Time compression ratio	ρ	720 (5 seconds/time slot)
Open time slots on wholesale market		24
Market closing time		1 time slot ahead
Distribution fee		[0.003 - 0.03]€/kWh
Balancing price basis	P	most recent clearing price
Balancing cost	c_0	[0.02 - 0.06]€/kWh
Slope of regulating market price	ϕ^+, ϕ^-	10^{-6}, 10^{-6} €/kWh
Default broker's min and max bid order prices	$l_{min}(\text{bid}), l_{max}(\text{bid})$	-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		[1000 - 5000] €
Tariff revocation fee		[100 - 500] €
Tariff publication interval		6 time slots
Annual bank debt interest rate	$[\beta_{min}, \beta_{max}]$	4.0 – 12.0%
Annual bank deposit interest rate	$[\beta'_{min}, \beta'_{max}]$	0.5β
Weather report interval		1 hour
Weather forecast interval		1 hour
Weather forecast horizon		24 hours

7.2 Competition ending

The game ends at a random number of K time slots after day 55 (time slot 1320), $K = 0, 1, \ldots$. For each time slot, starting day 55, there is a fixed probability p that the game ends by the end of that particular time slot. As a consequence, the number of time slots in excess of day 55, K, follows a geometric distribution. The expected number of time slots 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 time slots is equal to:

$$P(K \leq k) = 1 - (1-p)^{k+1}, \qquad \text{for } k = 0, 1, \ldots \tag{6}$$

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

$$\begin{aligned} P(K > k') \leq \omega \quad &\Leftrightarrow \quad (1-p)^{k'+1} \leq \omega & (7) \\ &\Rightarrow \quad k' \leq \frac{\ln \omega}{\ln(1-p)} - 1 & (8) \end{aligned}$$

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

$$P(K > k') \leq \omega \quad \Rightarrow \quad p \geq 1 - \sqrt[k'+1]{\omega} \tag{9}$$

If the probability that the game ends after 60 days (time slot 1440 - time slot 1320), $k' = 120$, is to be no more than 1%, $\omega = 0.01$, then the time slot 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, $fee_{b,c,g}$:

$$\pi_{b,c,g} = pay_{b,c,g} - cost_{b,c,g} - fee_{b,c,g} \tag{10}$$

- **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) after curtailments are exercised. Other payments for instance include sales in the wholesale market, and possible payments from external balancing.
- **Costs** are monetary transfers from brokers to customers (producers) and are based on the agreed contract conditions between the respective customer (producer) and broker and the actual (ex-post measured) energy produced after curtailments are exercised. Other costs for instance include procurement in the wholesale market.
- **Fees** are (i) the cost for external balancing 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 π_{bc}, the average accumulated profits of all brokers in the competition as $\overline{\pi}_c$ 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_{bc}, is obtained as:

$$z_{b,c} = \frac{\pi_{b,c} - \overline{\pi}_c}{S_c}, \tag{11}$$

where

$$\pi_{b,c} = \sum_{g=1}^{N_{b,c}} \pi_{b,c,g}, \tag{12}$$

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:

$$z_b = \sum_{c=1}^{C} z_{bc} \tag{13}$$

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 13). 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 though 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.

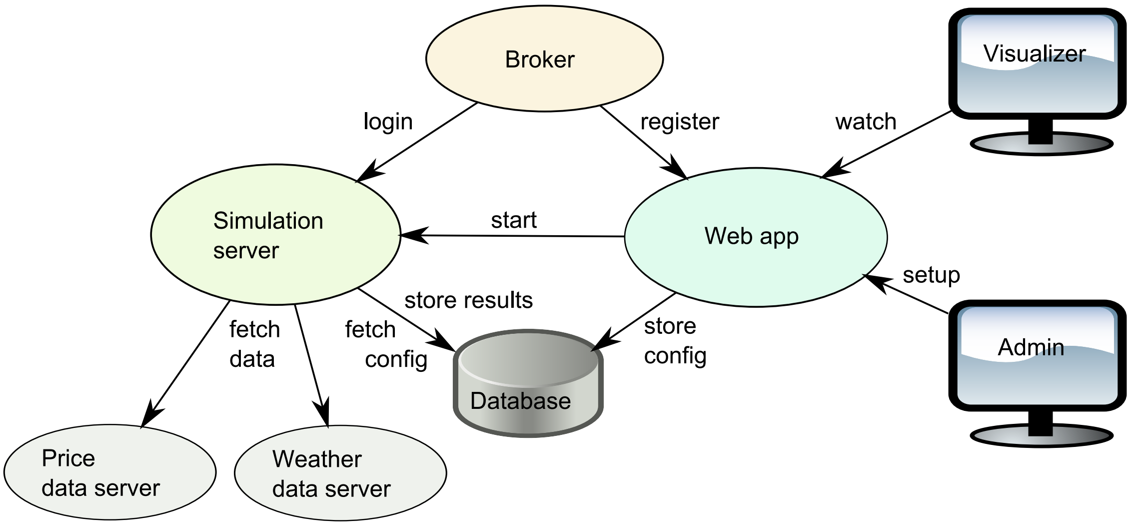

Figure 13: Competition systems architecture.

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 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.

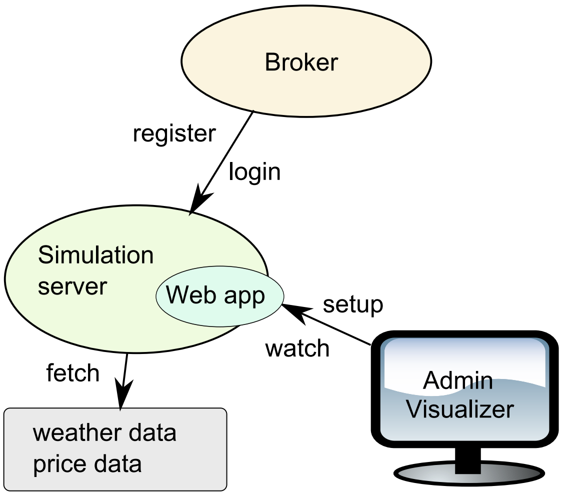

Figure 14: Research systems architecture.

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] R. Cavallo. *Social Welfare Maximization in Dynamic Strategic Decision Problems*. PhD thesis, Harvard University, 2008.

[5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[6] 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.

[7] 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 1–17, Charlotte, November 2011.

[8] Technology options for the near and long term. Technical report, U.S. Department of Energy Climate Change Technology Program, September 2006.

[9] S. Gottwalt, W. Ketter, C. Block, J. Collins, and C. Weinhardt. Demand side management - a simulation of household behavior under variable prices. *Energy Policy*, 39:8163–8174, 2011.

[10] P. 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.

[11] P. Joskow and J. Tirole. Retail electricity competition. *The Rand Journal of Economics*, 37(4):799–815, 2006.

[12] C. Möller, S. Rachev, and F. Fabozzi. Balancing energy strategies in electricity portfolio management. *Energy Economics*, 33(1):2–11, 2011.

[13] K. Skytte. The regulating power market on the Nordic power exchange Nord Pool: An econometric analysis. *Energy Economics*, 21(4):295–308, 1999.

[14] 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.

[15] 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.

[16] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

[17] 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 over-dimensioned with respect to their line capacities, thus this assumption is not a strong restriction but may have to be rethought in future once much more distributed generators and storage facilities are installed.

2. Power factor effects, i.e. phase shifts between voltage and current, are not taken into account. Modeling these effects would possibly influence the brokers' decision making on which consumers and producers to add to their portfolios but is out of scope at this time.

3. Power distribution and transformation losses are ignored. In Germany these losses are estimated at 3%; for North America they are estimated at 5,5% [8]. 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 "time slots" rather than as continuous time. This models the trading intervals in the regional wholesale market, and enables the simulation to model a period of days rather than minutes or hours.

7. The temporal distribution of energy consumption and generation *within* a time slot is not taken into account. This means for example that balancing power demand for a time slot is calculated as the difference of the sum of generation and the sum of consumption for that time slot and not as the instantaneous difference between the two timeseries.

8. Some portion of the load, including the charging and discharging of plug-in Electric Vehicles (PEVs), could be controlled by voluntary or automated means, using prospective or real-time price signals.

B Acknowledgements

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

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

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

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