The Power Trading Agent Competition

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

ERIM Report Series *Research in Management*	
ERIM Report Series reference number	ERS-2011-011-LIS
Publication	May 2011
Number of pages	42
Persistent paper URL	http://hdl.handle.net/1765/23263
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 2011 (Power TAC 2011). Agents are simulations of electrical power brokers, who must compete with each other for both power production and consumption, and manage their portfolios.
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 Power Trading Agent Competition

Wolfgang Ketter, John Collins[a], Prashant Reddy[b], Christoph Flath[c]

May 2011
ERS-2011-011-LIS

[a]Computer Science Department, University of Minnesota, USA
[b]Machine Learning Department, Carnegie Mellon University, USA
[c]Karlsruhe Institute of Technology, GER

Department of Decision and Information Sciences
Rotterdam School of Management
Erasmus University
3062 PA Rotterdam, The Netherlands

Abstract

This is the specification for the Power Trading Agent Competition for 2011 (Power TAC 2011). Agents are simulations of electrical power brokers, who must compete with each other for both power production and consumption, and manage their portfolios.

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

Contents

1 Background and motivation

We know how to build "smart grid" [1] components that can record energy usage in real time and help consumers better manage their energy usage. However, this is only the technical foundation. Variable energy prices that truly reflect energy scarcity can motivate consumers to shift their loads to minimize cost, and for producers to better dispatch their capacities [14]. This will be critical to the effort to develop a more sustainable energy infrastructure based on increasing proportions of variable-output sources, such as wind and solar power. Unfortunately, serious market breakdowns such as the California energy crisis in 2000 [3] have made policy makers justifiably wary of setting up new retail-level energy markets.

The performance of markets depends on economically motivated behavior of the participants, but proposed retail energy markets are too complex for straightforward game-theoretic analysis. Agent-based simulation environments have been used to study the operation of wholesale energy markets [22], 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 [11], 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.

The *Power Trading Agent Competition* is an example of a Trading Agent Competition (TAC)[1] applied to energy markets. Earlier successful examples of TAC include the Trading Agent Competition for Supply-Chain Management (TAC SCM) [7] and the Trading Agent Competition for Ad Auctions (TAC AA) [13].

2 Competition overview

The major elements of the Power TAC scenario are shown in Figure 1. Competing teams will construct trading agents to act as self-interested "brokers" that aggregate energy supply and demand with the intent of earning a profit. In the real world, brokers could be energy retailers, commercial or municipal utilities, or cooperatives. Brokers will buy and sell energy through contracts with retail customers (households, small and medium enterprises, owners of electric vehicles), and by trading in a wholesale market that models a real-world market such as the European or North American wholesale energy markets. Brokers compete with each other trying to attract customers by offering *tariff* contracts to a population of anonymous small customers (households, small businesses), and by negotiating individual contracts with larger customers (such as major manufacturing facilities, or greenhouse complexes with many Combined Heat and Power (CHP) units). Contract terms may include fixed or varying prices for both consumption and production of energy, along with other 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

[1] see http://www.tradingagents.org

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.

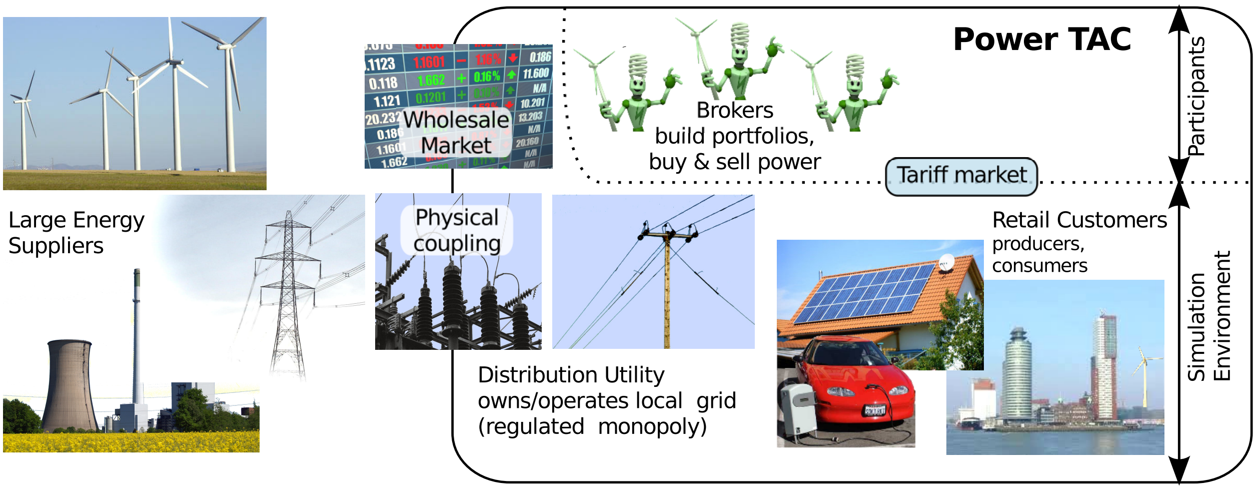

Figure 1: Major elements of the Power TAC scenario.

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.[2] 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. 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 un-

[2]Individual contract negotiation is not implemented for the 2011 pilot competition.

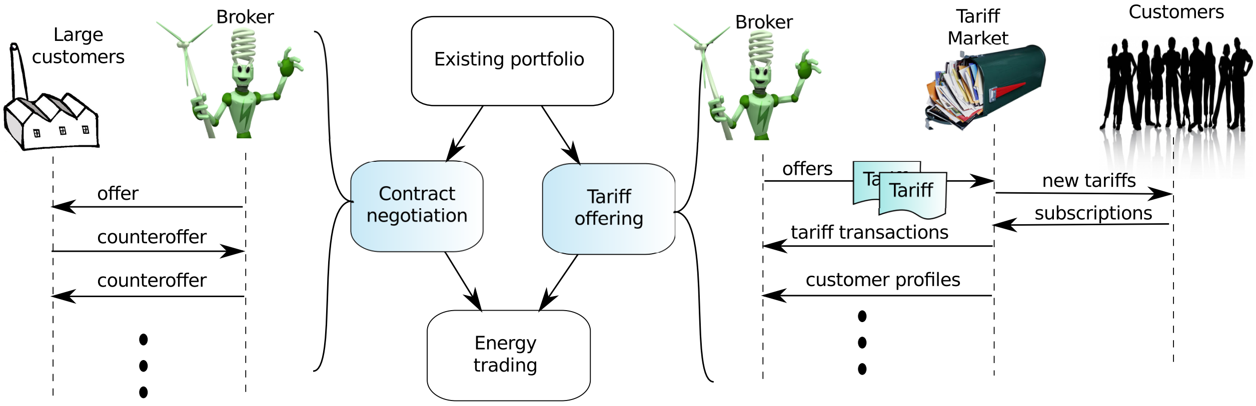

Figure 2: Portfolio management process. Tariff offerings proceed in parallel with individual contract negotiation.

correlated sources and loads that can be expected to balance each other in real time, by acquiring storage capacity, by acquiring flexible consumption and generation capacities (balancing capacity), by selling variable-price contracts, and by trading future energy supply contracts on the wholesale market.

2.1 Simulation time

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

2.2 Customer market

Broker agents try to acquire energy generation capacity from local producers and sell energy tariffs to local consumers. Brokers can buy and sell energy through two different mechanisms. For most customers, such as households, small businesses, and small energy producers, brokers may offer tariffs that specify pricing and other terms. For large producers or consumers (for example, a large industrial facility 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 4 times/day.

Power TAC supports fairly complex tariff specifications in-line with current developments in real-world electricity markets. Brokers can offer base fees, time-of-use tariffs with hourly intervals, tiered rates, sign-up and exit fees as well as dynamic pricing where the realized rate can continuously be adjusted by the broker. These tariff design elements allow the brokers to shape and control their portfolio.

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 [27]. With no common standard in place to build on for Power TAC, we use with the work of Tamma et al. [24], an ontology that describes a negotiation process including (i) the involved parties, (ii) the object to negotiate on, and (ii) the negotiation process, i.e. the economic mechanism itself.

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

Time: including points in time, time intervals, periodicity (days, weeks, months, etc.), and temporal relationships (before, after, during, etc.). These terms can be used to specify contract duration as well as other time-related contract terms.

Energy: including amounts of energy produced or consumed, and rate of production or consumption (power). Some contracts or tariffs will also need to specify amounts of energy that can be remotely controlled, for example by shutting off a domestic water heater for 15 minutes every hour during peak demand periods. Such remotely-controllable sources or loads are called "balancing capacity."

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

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

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

2.2.1 Tariff selection

The tariff market allows brokers to directly interact with the customer population (households, businesses, small power plants, electric vehicle owners) provided by the simulation environment. Brokers can publish tariffs (see Figure 8) from which customers will self-select the ones which are most appropriate for them. Therefore, broker to broker interaction is indirect in this market as competing tariff offers are evaluated against each other. Details on the Customer's tariff evaluation process are given in Section 4.3.

2.2.2 Contract negotiation

Individual contracts are negotiated through a request-for-quote (RFQ) process, initiated by large customers (producers and consumers of power), and proceeding through one or more rounds with as many brokers as continue to be interested. The process ends when any party accepts the current contract, or when either the RFQ originator or all brokers choose to withdraw. The smallest entities that will engage in this process will have capacities of at least 100 times the mean demand of individual households. The specifics of the negotiation process are undefined at this point, but there are many examples in the literature, such as [12].

Contract negotiation will be included for the first official competition in 2012, but not for the pilot competition in 2011.

2.3 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[3]. The brokers can buy and sell power contracts for future timeslots to optimize their portfolio. In the wholesale market brokers interact with each other directly as well as with generation companies (GenCos) and other wholesale market participants as described below in Section 2.3.3.

2.3.1 Trading and timeslots available for trade

Brokers can submit market orders or "shouts" (limit buy or sell) for the upcoming 24 timeslots. The timeslots available for trading are marked as "enabled"; changes in timeslot status are communicated to brokers at the beginning of each timeslot. Shouts submitted for non-enabled (disabled or not yet enabled) timeslots are silently discarded. Depending on the market configuration brokers may also be able to delete submitted shouts from order books. The market collects the submitted orders and stores them in the corresponding 24 order books.

2.3.2 Market clearing

When the simulation clock is advanced to the next timeslot, all orderbooks in the wholesale market are cleared. In the clearing process demand and supply curves are constructed from bids and asks to determine the clearing price of each orderbook at the intersection of the two — i.e. the price maximizing turnover. If there are multiple prices generating the same maximum turnover the market clears at the mean of the lowest bid and the highest ask price supporting this turnover. All bids above the clearing price and all asks below are executed. Figure 3 depicts the supply and demand curves for 5 market participants (A – E) offering energy contracts and 5 (F – J) looking to buy. In this example prices between 20 and 25 realize the maximum turnover of 30 [19]. In Power TAC the realized clearing price would be 22.5.

After the market is cleared the following steps are performed:

- Clearing price and volume are publicly broadcast (public information).
- Brokers are informed about their own executed transactions (private information).
- The necessary updates of cash and market positions are performed and communicated to brokers (private information).
- Depending on the market configuration all uncleared orders may be deleted.
- The cleared timeslot is disabled while the timeslot 24 hours ahead is enabled.

[3]see http://www.nordpoolspot.com, http://www.ferc.gov, or http://www.eex.com/en.

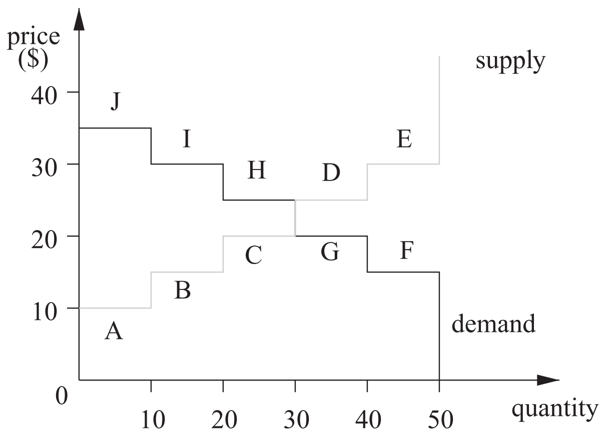

Figure 3: Illustrative supply and demand curves in the wholesale market [6]

2.3.3 Wholesale suppliers and buyers

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

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

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

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

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

In addition to the GenCos, there will be a wholesale buyer who is willing to buy unlimited quantities of power at a discount below the current lowest ask price from the set of GenCos for the current timeslot.

2.4 Balancing market

In electricity markets, supply and demand have to be balanced perfectly in real time. A major task of the Independent Systems Operator (ISO)[4] 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 gets more challenging as more renewable energy sources, such as solar and wind, are connected to the grid [25]. Many of these sources (e.g. wind) are only partially predictable. The grid balancing problem has been studied on various levels (wholesale vs. retail) and with different approaches [16].

In Power TAC, the DU needs to determine the price that energy brokers need to pay in case their profile of consumption and production of electricity is imbalanced in the current timeslot (see Figure 4).

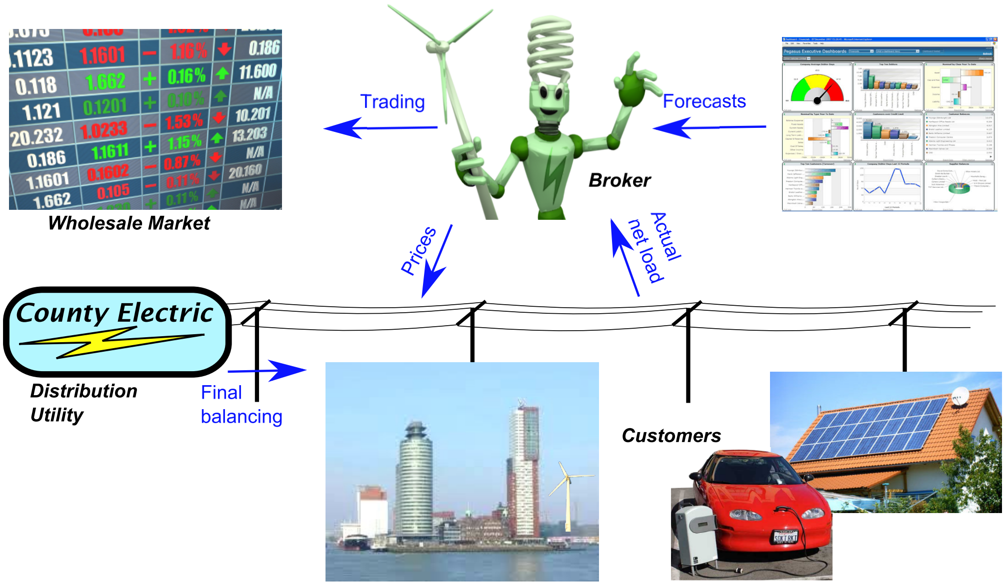

Figure 4: Entities and activities during balancing.

Brokers have the possibility to control some portion of customer production and consumption, by offering price concessions in exchange for the ability to remotely interrupt loads or sources for limited periods of time. Controllable loads are devices that are installed on the customer side and can be turned on or off by the broker or DU for a certain time period dependent on the type of contract the broker has with its customer. Examples of such controllable loads are remotely controlled CHPs (Combined Heat and Power – increased or reduced production) and domestic water heaters that can be remotely turned off (reduced consumption).

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

2.4.1 Adjusting energy demand and supply

Here we explain more formally how the simulation computes balance, and how brokers can act to avoid imbalance in the current timeslot. The total energy consumption for broker b in timeslot s is

$$e_c(b,s) = e_{\mathrm{ex}}(b,s) + \sum_{i=1}^{|\mathcal{C}_b|} e_i(s) \tag{2}$$

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

$$e_g(b,s) = e_{\mathrm{im}}(b,s) + \sum_{j=1}^{|\mathcal{G}_b|} e_j(s) \tag{3}$$

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

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

$$\forall s \in \mathcal{S},\ e_g(b,s) - e_c(b,s) = 0 \tag{4}$$

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

For each timeslot s, each broker b should balance expected supply and demand closely enough that the DU can achieve exact balance without requiring regulating services. Expected demand is the total expected load, or the sum of committed power exports and the expected loads $e'_i(s)$ of each consumer i in the broker's consumer portfolio $\mathcal{C}_b$ during timeslot s (see Equation 2):

$$e'_c(b,s) = e_{\mathrm{ex}}(b,s) + \sum_{i=1}^{|\mathcal{C}_b|} e'_i(s) \tag{5}$$

Expected supply is committed power imports plus total expected production capacity of all generators g within the broker's portfolio $\mathcal{G}_b$ during timeslot s (see Equation 3):

$$e'_g(b,s) = e_{\mathrm{im}}(b,s) + \sum_{j=1}^{|\mathcal{G}_b|} e'_j(s) \tag{6}$$

These values are maximum values in case some customers in the broker's portfolios have agreed to external control, presumably in exchange for better prices. For example, a combined heat and power generator with a nominal output of 50kW can be adjusted by an external control so that its real production is within certain boundaries, e.g., [40kW – 50kW]. Similarly, a domestic water heater may be configured to permit remote shutoff for up to 15 minutes every hour. The total controllable load for a broker b during timeslot s is $\epsilon_c(b,s)$, and the total controllable production capacity is $\epsilon_g(b,s)$. As long as $e_g(b,s) - \epsilon_g(b,s) \leq e_c(b,s)$ and $e_c(b,s) - \epsilon_c(b,s) \leq e_g(b,s)$, then supply and demand during timeslot s is expected to be in balance. Within this range, the DU will either reduce load or reduce output as needed to achieve exact balance.

The dispatching of balancing power (or load) by the DU is done only during the current simulation timeslot s_n. In Figure 5(a), we can see in the current slot s_n that both the actual observed supply and demand have deviated from the forecasted overall supply and demand for broker b. But as the difference between $e_c(b,s_n)$ and $e_g(b,s_n)$ was smaller than $\epsilon_g(b,s_n)$, the controllable production capacity of broker b in this slot, the DU was able to automatically reduce supply such that overall demand and supply for timeslot s_n was rebalanced.

For timeslot s_{n+1} in Figure 5(a), expected overall demand is forecasted to be within range of the available production capacity, but the uncertainty envelope (grey boxes) shows that this is not certain. In other words words $e'_g(b,s_{n+1}) - \epsilon_g(b,s_{n+1}) \leq e'_c(b,s_{n+1})$. After 2τ simulation time has elapsed (Figure 5(b)), this slot is now designated s_{n-1}, and we can see that the real consumption $e_c(b,s_{n-1})$ in this timeslot turned out to be lower than $e_g(b,s_{n-1}) - \epsilon_g(b,s_{n-1})$. This means that even after the simulation environment reduced the broker's production capacity to its minimum level, the overall production still exceeded the overall consumption. In this case the DU either reduced imports through the regulating market, or matched the surplus with a shortage of power from some other broker, to absorb the excess generated energy.

In slot s_{n+2} in Figure 5(a), a significant difference between overall production and overall consumption is forecast. Internal balancing capacity is likely to be insufficient for leveling the expected difference. In order to avoid the (expensive) utilization of external balancing power, broker b can either sell some of its surplus energy on the wholesale market, or use its contracted pricing power to try to encourage (i) some or all of its consumers to increase their demand, or (ii) some or all of its producers to reduce their production.

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

$$\hat{e}_c(i,s_{n+2}) = e'_c(i,s_{n+2},p_c(i,s_{n+2})) \tag{7}$$

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

$$\hat{e}'_c(i,s_{n+2}) = e'_c(i,s_{n+2},p'_c(i,s_{n+2})) \tag{8}$$

The ratio of demand change to price change

$$PE_i = \frac{\hat{e}_c(i,s,p) - \hat{e}_c(i,s,p')}{p - p'} \tag{9}$$

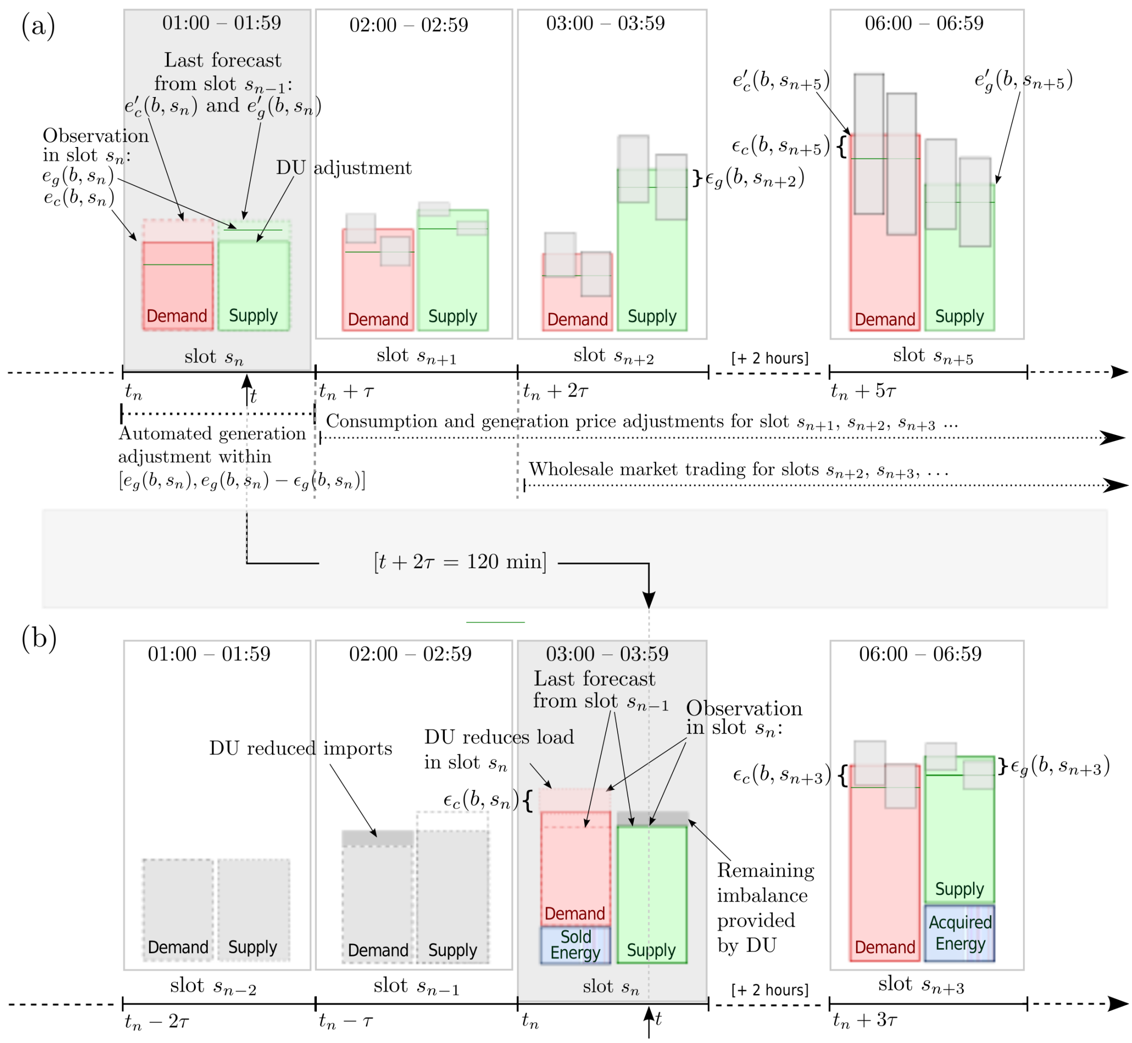

Figure 5: Broker's expected and actual energy supply and demand at two points in time.

is called the "price elasticity" for consumer i. Price elasticities will to be modeled within the different consumer agents provided by the competition environment following empirical findings on price elasticity as described for example in [23, 20].

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

2.4.2 Market-based balancing mechanisms

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

Scenario I: no controllable loads This will be implemented for the pre-pilot release.

Scenario II: static with controllable loads This will be implemented for the pilot competition.

Scenario III: dynamic with controllable loads This will be implemented for the first official competition.

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

Desirable balancing mechanism properties

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

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

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

5. Since manipulating the costs of potential controllable loads can lead to sub-optimal solutions, an additional goal in this setting is to make the strategy of declaring *the true capacities and costs of controllable loads incentive compatible*.

6. A second criterion in the case of controllable loads is a so-called participation constraint, i.e., the mechanism should benefit participating brokers, or otherwise brokers just will not declare any controllable load at all. In other words, the mechanism should be *individually rational.*

In the following two sections we discuss which mechanisms can be used for the above two scenarios (with/without controllable loads). The mechanisms for the different scenarios presented meet all the desirable properties described above.

Scenario I: no controllable loads

The relevant players in both scenarios are the N brokers denoted by $\{1, 2, \ldots, n\}$, and the system operator or distribution utility (DU), denoted by 0. In our analysis, as is commonly done, we consider energy production and consumption over a small amount of time, assuming that energy production and consumption are more or less stable during that period. In such a timeslot, each broker $i \in N$ has an expected net local energy (potentially negative) surplus of $x_i \in \mathbb{R}$. Each broker is requested to inform the DU of this surplus. Since this *declared surplus* potentially is different from the actual surplus, we denote it by $\hat{x}_i$.[5] Furthermore we use P^+ to denote the maximum market price of energy for this timeslot in the day-ahead market (over all day-ahead trade periods), P^- to denote the minimum market price over all day-ahead trade periods, and $P^* \geq P$ the all-time highest price possible.

The DU has the potential to import or export energy on the wholesale market (or apply its own ancillary services such as spinning reserves) to arrive at a perfectly balanced energy production and consumption. This comes at a cost of $c_0 : \mathbb{R} \to \mathbb{R}$ per unit. See Figure 6 for an example, illustrating that the cost of buying additional energy is higher than the benefit of selling additional energy at this last instance.

To make the brokers' strategy of declaring the true imbalance x_i incentive compatible (requirement 1), we apply a mechanism with verification [17]. These mechanisms define payments after the selected outcome (in our case the planned energy production and consumption) has been realized, and there is a possibility of verifying the agents' declaration in this realization; in our case this is the net energy production. To obtain incentive compatibility here, we charge a penalty correlated with the difference between the declared and true net energy production, which is guaranteed to be higher than the maximum profit that can be made with the same amount of energy in the day-ahead market.

Given the declared imbalances, the DU can compute the net declared imbalance $\hat{x} = \sum_{i \in N} \hat{x}_i$ over all brokers, and then apply its cost function c_0 to the opposite to determine the (expected) total costs for balancing, i.e., $-\hat{x} \cdot c_0(-\hat{x})$. Since in this first scenario there is no other way to recover from imbalances, this meets our requirement (2) of an efficient solution in case all imbalances are declared truthfully.

In addition to the penalty p_1 set by the mechanism with verification, payments need to be set such that also the third (regarding incentive compatibility) and fourth (regarding budget balancedness) requirement are met. We denote these payments by p_2. Since these payments are computed after the timeslot has passed, we can base them upon the real imbalances x_i. The third requirement in fact implies that the payment for an imbalance should always be higher than resolving the imbalance against the maximum market price P in the day-ahead market, i.e., $p_{2,i} \geq -x_i \cdot P^+$

[5] In Power TAC, the declared surplus does not necessarily need to be modeled, since the DU does not actually arrange resources to recover from the net imbalance of all brokers based on the declared surplus.

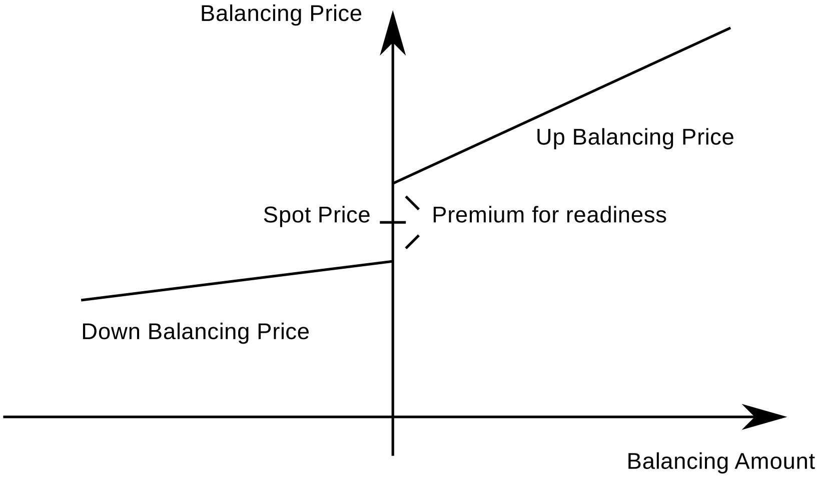

Figure 6: Usually the price (to be received) for a reduction of the production is significantly lower than the price on the spot market, and the price for producing more is significantly higher than the price on the spot market.

if $x_i < 0$, or otherwise $p_{2,i} \geq -x_i \cdot P^-$. Finally, the fourth requirement just says the payments from the brokers (who consume more than they produce) should be more than the payments to the brokers (who produce more than they consume) and the costs for recovering from the imbalance together, i.e., $\sum_i p_{2,i} + x \cdot c_0(-x) \geq 0$.

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

$$
\begin{array}{rll}
\text{minimize} & \sum_{i \text{ if } x_i<0}(p_{2,i} + x_i \cdot P^+)^2 + \sum_{i \text{ if } x_i \geq 0}(p_{2,i} + x_i \cdot P^-)^2 & \\
\text{subject to} & p_{2,i} \geq -x_i \cdot P^+ \quad \text{if } x_i < 0 & \\
& p_{2,i} \geq -x_i \cdot P^- \quad \text{if } x_i \geq 0 & \\
& \sum_i p_{2,i} + x \cdot c_0(-x) \geq 0 &
\end{array}
\tag{10}
$$

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

According to this definition over-consuming ($x_i < 0$) brokers always have to pay a positive amount. In the program in Equation 10 the distribution of the costs of balancing is defined by the minimization criterion, which expresses that each broker should pay an equal portion above the minimum amount defined by the constraints. However, this minimization criterion can be chosen differently (e.g., never let over-producing brokers pay a positive amount); the results below rely solely on the given constraints.

Scenario I will be used for the pilot competition in summer of 2011, using data from the econometric analysis of [21] for parameter settings, scaled by the size of the customer population.

Scenario II: static with controllable loads

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

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

$$\begin{array}{rl} \text{minimize} & \sum_{i \in N \cup \{0\}} \delta_i \cdot c_i\left(\delta_i\right) \\ \text{subject to} & \delta_i \in \left[c_i^-, c_i^+\right] \\ & \sum_{i \in N \cup \{0\}} \left(x_i + \delta_i\right) = 0 \end{array} \tag{11}$$

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

In case the correct information is used, Equation 11 meets the first requirement (of efficiency). Furthermore, if all the functions c_i are monotonically increasing, this problem is convex and can therefore again efficiently be solved [4].

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

1. The same mechanism with verification (with payments p_1) as in scenario I is used to make declaring the true imbalance a dominant strategy.

2. The same mechanism as in scenario I (with payments p_2) is used to make having no imbalance a dominant strategy.

3. An additional payment p_3 is introduced in this section to make declaring the true cost function a dominant strategy.

The utility of a broker i for a solution δ is not just defined by these payments, but also by the costs of load control, i.e., $u_i(\delta) = -c_i(\delta_i) - p_{1,i} - p_{2,i} - p_{3,i}(\delta)$. Regarding the incentives, note that the first payment is completely independent of the others, and that the analysis of the incentives from the previous section thus automatically transfers. The focus of this section will be on the third payment.

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

$$\begin{array}{lll} \delta_i \geq 0 & \text{if } \sum_{i \in N} x_i < 0 & \text{(under-production)} \\ \delta_i \leq 0 & \text{if } \sum_{i \in N} x_i > 0 & \text{(over-production)} \end{array} \tag{12}$$

With these conditions and step-wise cost functions, the mechanism is similar to a multi-unit auction (in the case of over-production) or a reverse multi-unit auction (in the case of under-production) [15]. With any type of cost functions, [10] mechanisms are the only mechanisms to achieve both an efficient allocation and a truthful declaration, in our case of the cost function of the controllable loads. Within this class, the VCG mechanism [26, 5] ensures that brokers always have a nonnegative utility for participating (i.e., individual rationality) under two conditions that hold in this domain: (i) choice-set monotonicity, which says that removing an agent never increases the set of alternative solutions, and (ii) no negative externalities, which says that every agent has zero (or more) utility for any choice that is made without its participation. When VCG is applied in the above setting, given the optimal production vector δ, the payment for each broker i is defined as follows (note that the sign is flipped because VCG is defined on the maximum social welfare, not on the minimal costs).

$$p_{3,i}(\delta) = -\sum_{j \neq i} \delta_j^{-i} \cdot c_j \left(\delta_j^{-i}\right) + \sum_{j \neq i} \delta_j \cdot c_j \left(\delta_j\right), \tag{13}$$

where δ^{-i} denotes the optimal solution to Equation 11 in a situation where the controllable loads of broker i cannot be used, i.e., $\delta_i = 0$. In case of over-production, all payments are positive, and thus the VCG mechanism meets all our requirements.

2.5 Accounting

Provides accounting services in support of other components. Cash accounting including customer billing and payments, tariff publication fees, market settlements, interest on debt, and credits and debits related to taxes and incentives. Market position accounting tracks the current commitments in the forward market for each broker. This is needed by the Distribution Utility to run the balancing process in the current timeslot.

2.6 Weather reports

Provides weather forecasts and current-hour weather conditions. Some customer models will use this information to influence energy consumption (temperature, for example), and production (wind speed, sky 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.

3 Brokers

3.1 Actions available to brokers

Figure 7 illustrates the timeline and information exchange between a broker and the simulation environment in each timeslot. Note that the specific order of events is more flexible than what is shown.

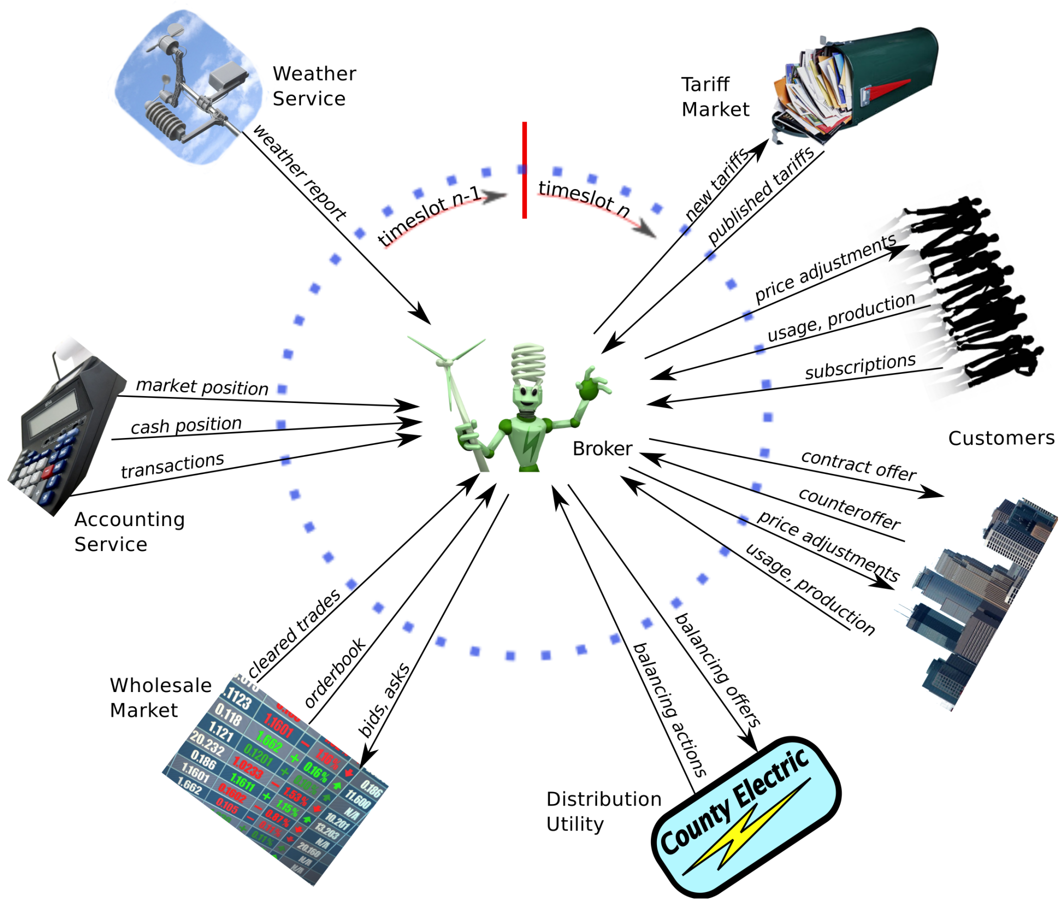

Figure 7: Illustration of Power TAC activities within one timeslot, where the broker receives information from the accounting and weather services, enters contracts with customers, trades in the wholesale market, and interacts with the distribution utility. The complete circle represents a 1-Hour timeslot; therefore each day consists of 24 timeslots.

In the following list we describe the minimum set of actions that a competition broker has to implement. The interaction component is listed in parentheses, e.g. Tariff Market.

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

Revoke tariffs (Tariff Market): Revoke tariffs from customers.

Price adjustments (Customers): Adjust prices in a current tariff.

Contract negotiation (large Customers): Agree on a contract through bilateral negotiation.

Balancing offer (Distribution Utility): Design contracts and tariffs that include controllable loads for real-time balancing.

Create asks and bids (Wholesale Market): Create asks and bids to sell (procure) energy.

3.1.1 Design, offer and revoke tariffs

To manage their portfolios, brokers have to design, offer and revoke tariffs. In the following we describe how tariffs in Power TAC can be constructed. Figure 8 shows a representation of a constrained version of this ontology, slightly denormalized to simplify relational mapping:

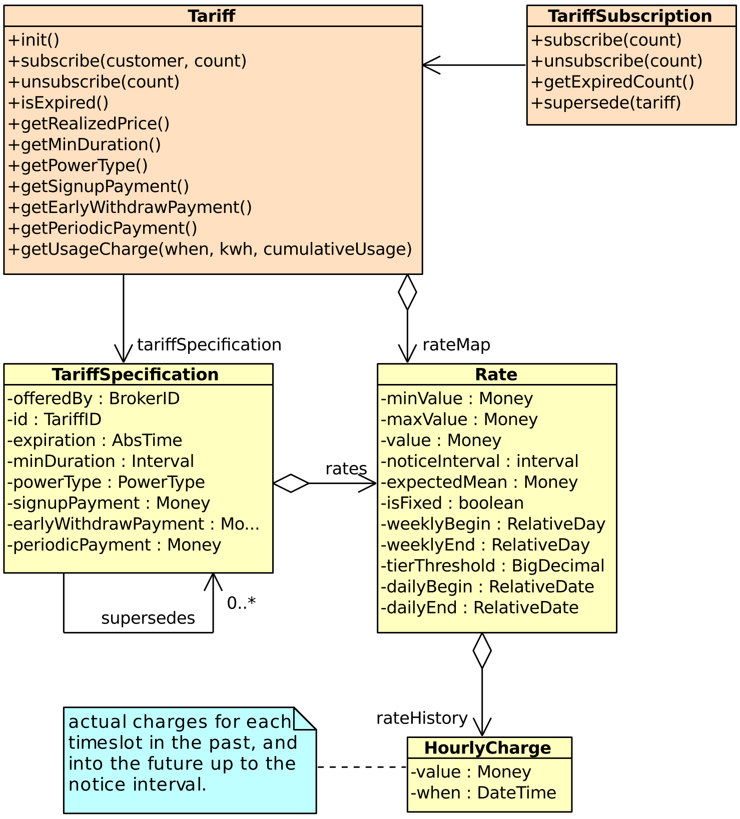

Figure 8: Tariff structure.

This structure allows a number of features within a simple, compact object graph. Many of the ontology concepts are collapsed into the Tariff 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 for unused features: null references. Here are some tariff features that can be represented:

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

- early withdrawal penalties;
- variable rates with minimum and maximum values, estimated mean values, and notice intervals.

This model is implemented in the `powertac-common` plugin. The classes in yellow, `TariffSpecification` and its associated `Rates` and `HourlyCharges` are serializable value objects that are used to communicate between Broker and Server. Once a `TariffSpecification` arrives at the server, it is wrapped with a Tariff, which remaps the Rates of a Tariff into a 2-D array, indexed by usage tier and hour. Customer subscriptions to the Tariff, are represented by `TariffSubscriptions`, which are responsible for keeping track of the number of customers in a population model that are currently subscribed, and for computing credits and debits for power usage and production.

A major feature of the full ontology that is not included here is the ability 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.

3.1.2 Dynamic pricing decisions

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

There are several environmental features that factor into the prices that the broker may want to settle on. At a basic level, a broker may want to forecast the load and supply from the customers for the target timeslot. Two major factors in the determination of this load and supply are (i) the estimated or realized load and supply for timeslots preceding the target timeslot, and (ii) the estimated weather conditions for the target timeslot. Correspondingly, in order to forecast, the weather conditions for the target timeslot, the broker may want to consider current weather conditions and historical averages for the time of the year.

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

3.1.3 Wholesale market trading

Dynamic adjustment of prices for consumers and producers who are on variable-price tariffs and the advance reservation of producer capacity as balancing power reserves are two possibilities to balance a broker's portfolio over time. Besides these two options, a third one is to buy missing, or to sell excess, capacities on the wholesale market.

Table 1 shows a sell order before its submission to the wholesale market (left) and after its clearing (right). In this example Broker 1 wants to sell 10 MWh of power within the timeslot May 01, 2011 03:00 – 03:59 at a minimum price of 40 €/MWh. At 01:00 the order is matched on the energy market at a price of 45 €/MWh and a quantity of 6.2 MWh. Broker 1 managed to sell 6.2 MWh of energy for this timeslot, which helps him to better balance his portfolio's overall energy supply and demand in this timeslot.

(a) Submitted Sell Order

Order ID:	1
Order Owner:	Broker 1
Timeslot:	May 01, 2011 03:00 – 03:59
Order Type:	Sell
Order Status:	Open
Quantity:	10 MWh
Limit:	40 €/MWh
Created at:	May 01, 2011 00:00
Clearing Date:	–
Clearing Quantity:	–
Clearing price:	–

(b) Matched Sell Order

Order ID:	1
Order Owner:	Broker 1
Timeslot:	May 01, 2011 03:00 – 03:59
Order Type:	Sell
Order Status:	Matched
Quantity:	10 MWh
Limit:	40 €/MWh
Created at:	May 01, 2011 00:00
Clearing Date:	May 01, 2011 01:00
Clearing Quantity:	6.2 MWh
Clearing price:	45 €/MWh

Table 1: Asks submitted to and later matched on the wholesale market.

3.1.4 Portfolio management

The primary goal of a broker is to acquire access to power sources and customers that result in a portfolio that is profitable and balanced, at least in expectation, over the period of the next execution activities and timeslots. 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. 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 sources 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 sources that can be used as needed (balancing sources), and by trading futures contracts on the regional exchange.

At any given time, a broker will have some number of active contracts, having negotiated them earlier. Such contracts have expiration dates beyond the current date. Also, tariffs offered earlier may remain active; customers who have agreed to a tariff in the past may or may not have an opportunity to opt out and choose a different tariff, and if they have the opportunity they may not choose to exercise it.

Acquire power sources Power source commitments are obtained by three different methods:

- Large local sources (large wind turbines, wind farms, large CHP plants, etc.) are traded in the local market through the RFQ process as described in Section 2.2.2.
- Small local sources (household and small-business sources) are obtained by offering tariffs in the local market, as described in Section 2.2.1.
- Power from the regional grid is obtained by trading in the wholesale market.

Power sources can be continuous or intermittent, and local continuous sources may have a non-zero balancing component as discussed in Section 2. Continuous sources include power obtained from the regional exchange, as well as the continuous portion of the output from many CHP and hydro plants. Intermittent sources include most renewable sources such as wind and solar plants.

Acquire storage capacity 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 local market through the tariff or the RFQ process. For example, individual owners of plug-in electric vehicles (PEVs) could sign up for 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.

Acquire loads Loads may be contracted through both the local market and the wholesale market, as is the case with power sources.

- Large local loads (industrial facilities and large office parks, for example), could negotiate rates through the RFQ process.
- Small local loads (households and small businesses, for example) must choose tariffs in the local market.
- Agents may choose to sell future power capacity in the wholesale market for periods when it expects to have a surplus. Such advance sales are hard commitments; the sold quantity of power will be transferred out of the system during that interval at a constant rate.

3.2 Information available to brokers

Tables 2, 3, 4 and 5 summarize the information available to brokers at various times during the game. Specifically, Table 2 lists information that is made available at **Startup**. Table 3 lists the information made available once per **Tariff Period**, which is typically once every 6 simulation hours. Table 4 lists the information made available once per **Timeslot**, which is typically once every 1 simulation hour. Table 5 lists the information available at **Anytime** during the game.

Each item in these tables is annotated with three attributes:

1. *Public* vs. *Private*: Is the item visible to all brokers; i.e., all competition participants, or only the relevant broker?
2. *Free* vs. *Cost*: Is there a cost associated with acquiring this information? The cost, if any, may be a fixed game parameter or change during the course of the game. Please refer to the detailed documentation for further specification.
3. *Receive* vs. *Request*: Is the information automatically received by the broker as a message or series of messages or does the broker need to initiate a request to the competition server to

obtain this information in response? Please refer to the detailed technical documentation for the APIs and message formats.

In the tables, values for the three attributes are indicated with a check mark (✓) for *Public*, *Free* and *Receive* while *Private*, *Cost* and *Request* are indicated with a cross (×).

Table 2: Information available to a broker at **Startup**. See Section 3.2.

Item	**Attributes**			**Description**
	Pub?	Free?	Rcv?	
Game parameters	✓	✓	✓	The parameters used to configure or instantiate the specific game. See Section 5.1 for details.
Initialization context	✓	✓	✓	Additional information, besides the competition parameters, *e.g.*, identities of participating brokers, available after the login process has completed. See Section 5.1 for details.
Broker identities	✓	✓	✓	The identities of the participating brokers in the current game. A particular competition participant maintains the same identity over the different rounds of the competition.
Default tariffs	✓	✓	✓	At game initialization, the tariff market consists only tariffs offered by the Default Broker.
Weather report	✓	✓	✓	The current weather and the forecast for the next 48 hours.
Wholesale market prices	✓	✓	✓	Prices for each of the 23 open trading hours on the wholesale market, based on offers from large generating stations and bids from the Distribution Utility since brokers are not yet participating in the wholesale market.

Table 3: Information available to a broker per **Tariff Period**. See Section 3.2.

Item	**Attributes**			**Description**
	Pub?	Free?	Rcv?	
Tariff updates	✓	✓	✓	New tariffs, revoked tariffs and superseding tariffs where applicable.
Portfolio changes	×	✓	✓	New and dropped customer subscriptions.
Tariff transactions	×	✓	✓	Bonus and early-exit penalty transactions.
Customer history	×	✓	×	7-day historical consumption and production profile of customers in current portfolio.

4 Customers

The simulation includes a range of customer models, including electric vehicles, CHPs, solar panels and wind turbines, and multiple models of private households, clustered by preference similarity.

Table 4: Information available to a broker per **Timeslot**. See Section 3.2.

Item	**Attributes**			**Description**
	Pub?	Free?	Rcv?	
Wholesale market prices	✓	✓	✓	Market clearing prices for the 23 currently open trading hours in the wholesale market.
Balancing price	✓	✓	✓	Price set by the DU to clear the balancing market.
Total supply and demand	✓	✓	✓	Consumption and production metrics aggregated over all customers as seen by the DU; includes derived statistics on net exports and imports from or into the wholesale market.
Weather report	✓	✓	✓	Realized weather for the current timeslot and forecast for the next 48 hours.
Portfolio supply and demand	×	✓	✓	Production and consumption metrics for the current customer portfolio, broken down by customer subscription (individual customers or fractions of customer population models).
Wholesale market cleared trades	×	✓	✓	Execution reports for bids and offers cleared by the wholesale market.
Market position	×	✓	✓	Net import/export commitments for the 23 open trading hours on the wholesale market.
Balancing report	×	✓	✓	List of balancing actions taken by the Distribution Utility on customers in the current portfolio.
Accounting transactions	×	✓	✓	List of market transaction entries as processed by the Accounting service for the current period; includes tariff subscription-related transactions, wholesale market transactions and balancing fee transactions. Note that tariff transactions are reported separately.
Cash position	×	✓	✓	Updated cash position after all current accounting transactions have been applied.
Negotiation bids/offers	×	✓	✓	Initial or continuing updates on negotiations with contract customers. Please refer to the Negotiation Protocol documentation for further details.
Level-2 wholesale market data	✓	×	✓	Full orderbook from the wholesale market data with price and quantity of all bids/offers, from all brokers, to be considered for clearing for the 23 open trading hours.

An important feature of these models is their responsiveness to price changes. A special focus lies on modeling substitution effects between timeslots as price elasticity effects would be very limited in 60 days of simulation time. In the literature such effects have been analyzed by means of synthetic aggregate models [18] or micro-founded bottom-up models [2]. Power TAC's dynamic

Table 5: Information available to a broker at **Anytime**. See Section 3.2.

Item	**Attributes**			**Description**
	Pub?	Free?	Rcv?	
Resynchronization context	✓	✓	×	The game parameters and initialization context available during Startup can also be queried if a broker needs to recover from a crash.
Current tariffs	✓	✓	×	The full set of tariffs that have not been expired, revoked or superseded.
Reputation metrics	✓	✓	×	Realized prices on variable rate tariffs and other computed metrics quantifying the reputation of each broker in the market.
Aggregate customer profile	✓	×	×	Aggregate metrics characterizing the set of customer models instantiated in this game. (TODO: Need more detail.)
Demand forecast	✓	×	×	TBD
Wholesale market forecast	✓	×	×	TBD
Tariff market forecast	✓	×	×	TBD - Forecasts for default tariffs only

demand models extend the latter approach to describe a rich customer population.

From a technical perspective customers are realized in the form of plugins. A customer model plugin instantiates a population of a customer type. Such population models can represent larger groups of homogeneous customer which helps to reduce computational complexity. Plugins can maintain local databases to store information to guide their further activity. The plugin approach allows researchers to investigate questions related to specific consumer types or behavioral assumptions by using only relevant customer models.

In the game context customers perform three major tasks; choosing tariffs, recording meter readings and providing balancing capabilities. Choosing tariffs facilitates customer participation in the tariff market; customer balancing capabilities provide an additional balancing tool while the generation of meter readings instantiates the fundamental objective of interest of Power TAC – dispersed consumption and generation of energy that need to be balanced. Therefore, customer models are very important building blocks of the game environment.

4.1 Customer types

In the Power TAC simulation environment the `CustomerType` enumerations allows for the designation of generic customer types. These customer types characterize the general load profile of a customer. The following customers types are implemented:

- households — typical residential consumption behavior; potentially provision of generation from renewable energy sources
- offices — typical flat consumption throughout working hours, limited consumption at other times
- factories — similar to office consumption but with greater magnitudes and more variance

- electric vehicles — large loads (charging / V2G) when connected to grid, otherwise zero load

A customer's load profile is further specified by the power type it supports. A customer uses at least one of the types from the `PowerType` enumeration, but can also use multiple. The following power types are implemented:

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

4.2 Population models

Power TAC provides support for customer population models. These population models wrap multiple customers of the same generic type. This allows representation of large customer bases without compromising server performance. Population models have a total number of committable consumption which they can split up between multiple tariffs.

4.3 Tariff market interaction

The tariff market facilitates the matching of consumers and brokers. Customer models actively participate in the tariff market by choosing new tariffs through periodic evaluation of the tariffs offered by the brokers. Customers face the tariff selection problem as shown in Figure 9. That is they need to select an attractive tariff from the set of available tariffs.

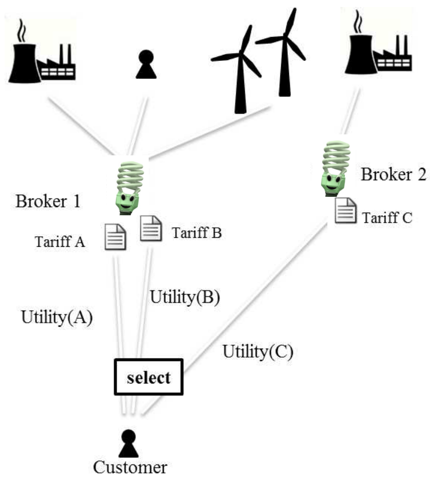

Figure 9: Tariff market interaction.

The key part of customer tariff evaluation is calculation of the expected cost (gain) over the lifetime of a contract relationship. This quantity is composed of the expected variable payments from estimated consumption (production), periodical payments as well as sign-up fees or bonuses. Especially the derivation of expected variable payments is crucial: It needs to properly reflect a customer's consumption (production) choice under the tariff to be evaluated. Therefore, tariff choice needs to be fundamentally driven by consumption choice under a tariff as described in the next subsection. Especially for complex tariffs this is a key design challenge for creating customer models. Since early exit is possible, customer models may evaluate available tariffs at any time. Clearly in this case, a proper switching evaluation has to additionally factor in the exit fees from leaving the current tariff. This monetary evaluation is complemented by an additional assessment of other tariff aspects, e.g. broker reputation, energy sources, interruptibility properties or early exit fees. The tariff comparison is therefore described by a utility value for each available tariff. This value moderates costs and other factors. The tariff utility function and the corresponding tariff choice logic are the key characteristics of customer model actions in the tariff market. Elicitation of these tariff preferences is thus a major aspect of a successful broker strategy.

From the currently available tariff list customers need to select a suitable one (see Figure 10). 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

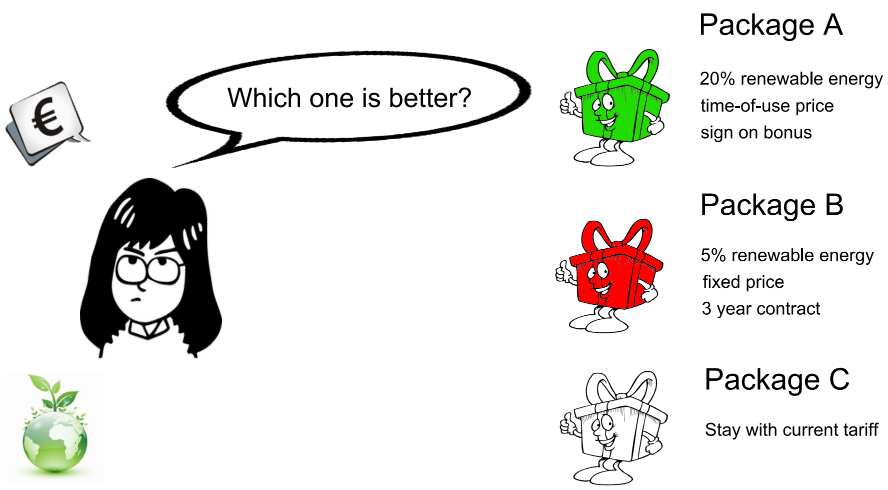

Figure 10: Tariff selection problem.

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

4.3.1 Deriving tariff utility

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

$$u_i = -(c_v + c_f)\alpha_{cost} - e_i\alpha_{energy} - r_i\alpha_{risk} - I_i\alpha_{inertia}. \tag{14}$$

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

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

Fixed tariff payments c_f Fixed tariff payments consist of sign-up fees/ bonuses of the new tariff, $c_{\text{sign-up}}$, daily periodic payments c_{daily} as well as exit fees of current tariff c_{exit}. These costs need to be normalized to a one day time span. While this is trivial in case of the periodic payment, it requires the expected tariff life $\tilde{t}$ for the other payments.[6] The normalized values of the fixed payments are summed to obtain the fixed other payments value, $c_f = c_{\text{daily}} + \frac{c_{\text{sign-up}} + c_{\text{exit}}}{\tilde{t}}$.

Energy content e_i e_i is the percentage of non-renewable energy sources in the tariff considered.

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

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

$$I_i = \begin{cases} 1 & \text{if } i \neq j \\ 0 & \text{if } i = j. \end{cases} \tag{15}$$

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

4.3.2 Choosing from a list of tariffs

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

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

$$\mathbb{P}_i = \frac{e^{\lambda u_i}}{\sum_{t \in \mathbb{T}} e^{\lambda u_t}} \tag{16}$$

The parameter $\lambda \geq 0$ is a measure for how rationally a customer chooses tariffs: $\lambda = 0$ represents random, irrational choice, while $\lambda = \infty$ represents perfectly rational customers always choosing the tariff with the highest utility. Depending on the customer model type this choice probability can be used in two ways — either to represent randomized and 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.

[6] The derivation of $\tilde{t}$ may be customer-specific.

4.4 Provide balancing capacity

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

Interruptible consumption

Certain types of appliances (water heaters, heat pumps) can support remote interruption by the broker. If a broker has such interruption under contract it can, when used, be turned off for the broker's balancing actions.

Pledged energy from storage

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

Controllable micro generation

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

4.5 Record meter reading

Customer models ultimately represent the entities connected to the grid. As such the game implications of their actions can be represented as timeslot meter readings for both consumption (positive reading) and generation (negative reading). Serving and balancing the meter readings of their customer population is the brokers' main task in Power TAC. Hence the determination of meter readings is a crucial task performed by customer models.

The meter readings generated by a customer will typically depend on model primitives (number of household members, work shift hours), randomization, the current season and weather conditions (e.g. turning on A/C, output from solar panels), the selected tariff (time-of-use pricing induces customers to shift consumption) as well as balancing capability actions (respond to current or previous load interruption).

The relationship between a customer's tariff and the meter reading is described by an economic consumption or generation logic. In the following sections typical implementations for these consumption/ generation logics are described. Note, that the other consumption influences (weather, balancing actions) influence this logic by changing the origin meter reading.

Fully static consumption

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

Flexible consumption amount, static timing

This type of customer model implements a simple demand behavior: for each timeslot the optimal consumption amount is decreasing in the timeslot electricity price. Such customer models reflect synthetic consumption profiles determined in a top-down approach considering aggregate electricity consumption as a continuous good with positive and decreasing marginal utility. Controllable generation with well-defined cost functions (e.g. micro-CHP) are also captured by this modeling approach. These models are especially helpful for economic analyses as their behavior can be described in compact mathematical form.

Static consumption amount, flexible timing

Customer models who can change the timing of their loads (e.g. through automatic appliance scheduling) will not change their consumption amount under a given tariff but will try to minimize their cost by scheduling the activities appropriately. Such household models are typically bottom-up models where consumption originates from the activity/ appliance level [2].

Fully dynamic consumption

Fully dynamic consumption features both flexible consumption amounts as well as flexible timing. Such models can be both top-down as well as bottom-up. Top-down models specify cross-price elasticities between timeslots [18]. Fully dynamic bottom-up models endogenize price for activity occurrence and scheduling.

5 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.

5.1 Competition initialization and Default Broker

The standard competition parameters can be found in Table 6. Values for these parameters are sent to a broker at the start of every game. For details see the software documentation.

To create a fair start of each game Power TAC initially allocates all tariffs to an incumbent *Default Broker*, a non-competition broker. After a few timeslots, the Default Broker will change all existing tariffs. This will force customers to evaluate new tariffs, and therefore stimulate tariff switching to competition brokers.

5.2 Competition ending

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

Table 6: Parameters used in the Power TAC game.

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
Nominal length of game	E	60 days
Probability that there are k timeslots after timeslot 1320 (start of day 55) before end of game	$[p_{min}, p_{max}]$	$[p_\omega, 1]$
Probability for each timeslot after timeslot 1320 (start of day 55) of end of game	p	$\frac{1}{121}$
Minimum game length	Min(TS)	1320
Expected game length	E(TS)	1440
Timeslot length	x	60 minutes
Time compression ratio	ρ	720 (5 seconds/timeslot)
Open timeslots on wholesale market	x	24
Market closing time	x	1 timeslot ahead
Distribution fee	x	[0.01 - 0.3]/kWh
Balancing price basis	P	mean price over final three market clearings
Balancing cost	c_0	[0.02 - 0.06]/kWh
Tariff publication fee	x	1000
Tariff revocation fee	x	1000
Tariff publication interval	x	6 timeslots
Annual bank debt interest rate	$[\beta_{min}, \beta_{max}]$	4.0 – 12.0%
Annual bank deposit interest rate	$[\beta'_{min}, \beta'_{max}]$	0.5β
Weather report interval	x	1 hour
Weather forecast interval	x	24 hours
Weather forecast horizon	x	48 hours

$E(K) = (1-p)/p$. The cumulative probability distribution that the game ends after at most k extra timeslots is equal to:

$$P(K \leq k) = 1 - (1-p)^{k+1}, \qquad \text{for } k = 0, 1, \ldots \tag{17}$$

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

$$P(K > k') \leq \omega \quad \Leftrightarrow \quad (1-p)^{k'+1} \leq \omega \tag{18}$$

$$\Rightarrow \quad k' \leq \frac{\ln \omega}{\ln(1-p)} - 1 \tag{19}$$

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

$$P(K > k') \leq \omega \quad \Rightarrow \quad p \geq 1 - \sqrt[k'+1]{\omega} \tag{20}$$

If the probability that the game ends after 60 days (timeslot 1440 - timeslot 1320), $k' = 120$, is to be no more than 1%, $\omega = 0.01$, then the timeslot ending probability should be set at $p \geq 1 - \sqrt[121]{0.01} = 0.037$. The choice of p will be operationalized as a random drawing from a uniform distribution defined on the domain $[p_\omega, 1]$, where p_ω refers to the probabilities calculated before; for example, $p_{0.01}$ would be 0.037. Given the random end of game and that each Power TAC day lasts 120 seconds in real time, an average Power TAC game will last around 2 hours overall.

5.3 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

5.4 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. Consequently this section describes the performance criteria used to rank order the Power TAC participants; albeit with quite some differences compared to previous TAC competitions. 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.

5.4.1 Performance criteria

Each Power TAC participant (broker) is assessed and an overall rank order of all participants is created based on overall profits p^{profit}, calculated as the (monetary) payments, p^{pay}, minus costs, p^{cost}, minus fees, p^{fee}:

$$p^{\text{profit}} = p^{pay} - p^{\text{cost}} - p^{\text{fee}} \tag{21}$$

- **Payments** are monetary transfers from customers (consumer)to brokers and are based on the agreed contract conditions and the actual (ex-post) measured energy consumptions of the respective customer (consumer) as described in Section 2.4.1. Other payments for instance include sales in the wholesale market, and possible payments from external balancing.
- **Costs** are monetary transfers from brokers to customers (producers) and are based on the agreed contract conditions between the respective customer (producer) and broker and the actual (ex-post measured) energy produced as described in Section 2.4.1. Other costs for instance include procurement in the wholesale market.
- **Fees** are (i) the cost for external balancing power (see Section 2.4) used, and (ii) 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.

5.4.2 Final ranking algorithm

After each competition round ends, i.e. at the end of the finals, z-scores of the accumulated profits for each broker can be straightforwardly 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 a_{bc}, the average accumulated profits of all brokers in the competition as $\overline{a}_c$ 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:

$$z_{bc} = \frac{a_{bc} - \overline{a}_c}{S_c} \tag{22}$$

After all competitions C have ended, an overall measure of relative broker performance will be obtained by taking a simple average of the standardized broker performance per competition:

$$z_b = \frac{1}{C} \sum_{c=1}^{C} z_{bc} \tag{23}$$

where C is the number of competitions.

5.4.3 Tournament structure

A typical Power TAC tournament consists of several rounds:

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 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 tournament.

5.5 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.

- Note that 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 game 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 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.
- 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.
- Collusion is not allowed between the different agents, i.e. their teams.

6 System architecture

6.1 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 11.

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

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

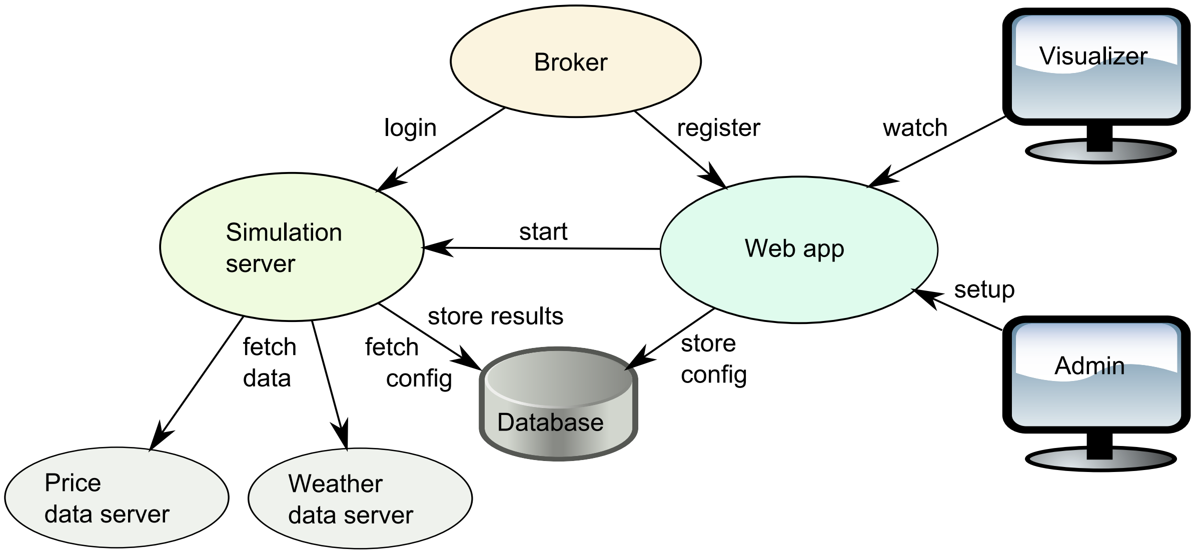

Figure 11: Competition systems architecture.

services, hosted on their own databases. The shared database will hold summary information for completed games, including access information for retrieving game logs.

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

6.2 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 12. 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 12 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.

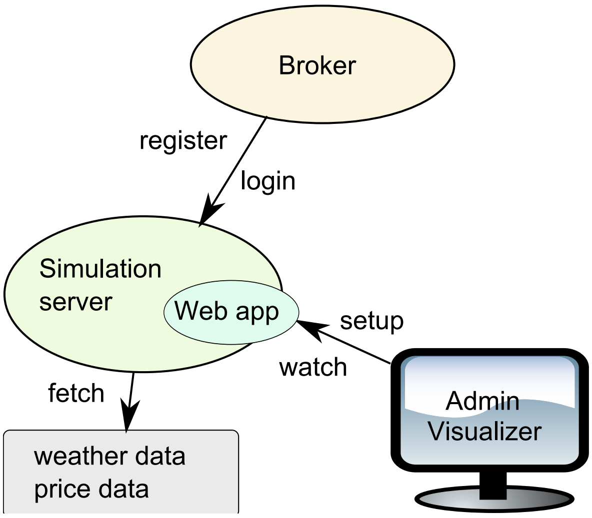

Figure 12: Research systems architecture.

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] C. Block, J. Collins, S. Gottwalt, W. Ketter, and C. Weinhardt. Modeling household energy consumption under fixed and variable pricing. In *Workshop on Information Systems and Technology*, St. Louis, Missouri, USA, December 2010.

[3] S. Borenstein, J. B. Bushnell, and F. A. Wolak. Measuring market inefficiencies in California's restructured wholesale electricity market. *The American Economic Review*, 92(5):1376–1405, 2002.

[4] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

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

[6] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, Hewlett-Packard Laboratories Bristol, UK, 1997.

[7] J. Collins, W. Ketter, and N. Sadeh. Pushing the limits of rational agents: the trading agent competition for supply chain management. *AI Magazine*, 31(2):63–80, 2010.

[8] M. de Weerdt, W. Ketter, and J. Collins. Pricing mechanism for real-time balancing in regional electricity markets. In *Workshop: Trading Agent Design and Analysis at the Twenty-second Int'l Joint Conf. on Artificial Intelligence*, pages –, Barcelona, July 2011.

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

[10] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[11] C. Hirsch, L. Hillemacher, C. Block, A. Schuller, and D. Möst. Simulations in the smart grid field study MeRegio. *it - Information Technology*, 52(2):100–106, 2010.

[12] C. M. Jonker and J. Treur. An agent architecture for multi-attribute negotiation. In *International joint conference on artificial intelligence*, pages 1195–1201, 2001.

[13] P. R. Jordan, B. Cassell, L. F. Callender, and M. P. Wellman. The ad auctions game for the 2009 trading agent competition. Technical report, University of Michigan, Department of Computer Science and Engineering, 2009.

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

[15] V. Krishna. *Auction Theory.* Academic Press, 2002.

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

[17] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2):166–196, 2001.

[18] S. S. Oren, S. A. Smith, and R. B. Wilson. Multi-product pricing for electric power. *Energy Economics*, 9(2):104–114, 1987.

[19] S. Parsons, J. A. Rodriguez-Aguilar, and M. Klein. Auctions and bidding: A guide for computer scientists. *ACM Comput. Surv.*, 43:10:1–10:59, February 2011.

[20] A. Siddiqui, E. Bartholomew, and C. Marnay. Empirical analysis of the spot market implications of price-elastic demand. Berkeley Lab Publications LBNL-56141, Lawrence Berkeley National Laboratory, July 2004.

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

[22] 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.

[23] K. Spees and L. Lave. Impacts of responsive load in PJM: Load shifting and real time pricing. *The Energy Journal*, 29(2):101–122, 2008.

[24] V. Tamma, S. Phelps, I. Dickinson, and M. Wooldridge. Ontologies for supporting negotiation in e-commerce. *Engineering Applications of Artificial Intelligence*, 18(2):223–236, March 2005.

[25] 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.

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

[27] D. von Dollen. Report to NIST on the smart grid interoperability standards roadmap. Technical Report SB1341-09-CN-0031, Electric Power Research Institute (EPRI), June 2009.

A Assumptions

In particular we make the following *assumptions*:

1. Within the simulated region, grid constraints (line capacity limitations) are assumed to be non existent, i.e. power flows within the region are unconstrained. Local distribution grids are typically overdimensioned with respect to their line capacities, thus this assumption is not a strong restriction but may have to be rethought in future once much more distributed generators and storage facilities are installed.

2. The point of common coupling (PCC) between the simulated distribution grid and the higher level transmission grid has a maximum capacity for power inflow and outflow. A specialized agent that serves as a "liquidity provider" on the regional energy spot market, and is able to arbitrage with the national energy spot market, has to obey these technical limits.

3. 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.

4. Power distribution and transformation losses are ignored. In Germany these losses are estimated at 3%; for North America they are estimated at 5,5% [9]. These losses can be considered as being more or less constant within a distribution grid and identical for all grid participants. Thus the validity of the simulation results is not affected.

5. 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.

6. 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.

7. The simulation will model time as a series of discrete "timeslots" rather than as continuous time. This models the trading intervals in the regional wholesale market, and enables the simulation to model a period of days rather than minutes or hours.

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

9. 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

We would like to thank Carsten Block, who was one of the initiators (besides Wolfgang Ketter and John Collins), for his involvement in the early stages of the project, and the great discussions we had in Karlsruhe, and Rotterdam.

We would like to thank further the members of the TAC, E-Energy, 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 goes to Antonios Chrysopoulos, Jan van Dalen, Travis Daudelin, David Dauer, Clemens van Dinther, Joseph D'Costa, Sebastian Gottwalt, Yixin Lu, Adis Mustedanagic, Nguyen Nguyen, Vedran Podobnik, Kang Puthyrak, Daniel Schnurs, Andreas Symeonidis, Manuela Veloso, and Mathijs de Weerdt.

Publications in the Report Series Research* in Management

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

2011

Sequencing Heuristics for Storing and Retrieving Unit Loads in 3D Compact Automated Warehousing Systems
Yugang Yu and René B.M. De Koster
ERS-2011-003-LIS
http://hdl.handle.net/1765/22722

A Local Search Algorithm for Clustering in Software as a Service Networks
Jelmer P. van der Gaast, Cornelius A. Rietveld, Adriana F. Gabor, and Yingqian Zhang
ERS-2011-004-LIS
http://hdl.handle.net/1765/22723

Implementing Standardization Education at the National Level
Henk J. de Vries
ERS-2011-007-LIS
http://hdl.handle.net/1765/22812

The Power Trading Agent Competition
Wolfgang Ketter, John Collins, Prashant Reddy, and Christoph Flath
ERS-2011-011-LIS
http://hdl.handle.net/1765/23263

---

* A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship