

A Technical and Distributed Management Basis for an Environmentally Clean and Sustainable Energy Supply

Horst F. Wedde¹, Sebastian Lehnhoff¹, Edmund Handschin², Olan Krause²

Abstract

Rising market prices for energy, an apparent future shortage in fossil fuels, and alarming reports on pollution through CO₂ are causing a world-wide trend towards renewable and ecologically clean forms of energy. We report about ongoing work in the R&D project DEZENT establishing renewable electric energy supply and eventually replacing fossil energy sources. Producers are at the same time also consumers. Their production and consumption are largely unpredictable. With our combined expertise in Real-Time systems and Electric Power Distribution we developed price negotiations which are pursued by consumer/ producer agents on a P2P basis and are governed by tough end-to-end deadlines (< 0.5 sec) dictated by EE constraints. The strategies used for periods of 0.5 sec are designed for fast convergence while we may at the same time assume a constant demand/ supply situation. Malicious users will not succeed, and customers pay considerable less than under conventional management policies or structures. In this paper we allow the negotiation strategies themselves to be adaptive across periods thus achieving a most flexible bargaining for each individual customer involved. For this purpose we have defined distributed learning algorithms derived from *Reinforcement Learning*. While maintaining all benefits from the earlier stage of development we demonstrate that we obtain a much better performance across periods than the initial static algorithms. To our knowledge we have presented and investigated the first distributed learning algorithm in the area of Adaptive Real-time Systems. Since the electric distribution management can be equally finalized within each period we have laid the ground for a thorough provision with sustainable and clean electric energy.

1. Introduction

Rising market prices for energy, an apparent future shortage in fossil fuels, and alarming reports on pollution through CO₂ are causing a world-wide trend towards renewable and ecologically clean forms of energy. This trend has been steadily growing over the past decade. Private investments have been encouraged and heavily subsidized in most of the European countries, through tax deductions, and even more through a very favorable refund program for feeding electric power from renewable sources into the public network.

In 2005 the European Union for the Coordination of Transmission of Electricity (UCTE) demanded to impose an obligation on grid operators to reduce integration costs for renewable energy capacities and to enable wind turbines to actively contribute to grid stability [UCTE05]. Enabling renewable power capacities to serve as reserve and balancing power capacities could obviously meet these demands. This was an encouraging incentive for our research.

Technological Basis. Different from traditional power production, while environmentally clean and sustainable technologies are based on solar or wind power, or on other renewable energy sources, they are typically realized through highly distributed small or mid-size facilities. All of them are absolutely environmentally clean, or climate-neutral as in the case of vegetable oil based BHPPs. Fuel cells and combined block heat and power plants can even be put into the basement of private homes. Their start-up and shut-down times are very short (ca. 1 min.).

Through electrolytic processes hydrogen and oxygen can be produced from excess energy thus allowing for a stable *long-term storage* of electric energy. Recent progress in lithium-ionic accumulator

¹ School of Computer Science, University of Dortmund, 44221 DORTMUND, Germany

² College of Electrical Engineering, University of Dortmund, 44221 DORTMUND, Germany

technology in combination with ultra- or super-capacitors allows for an efficient *short-term storage strategy*, with extremely high power gradients. These combined technologies allow for efficient and scalable power storage, providing high peak power within 0.5 sec when required [BuB05].

Instead of costs for raw materials, or of their transportation to large power plants, renewable energy largely comes for free (or at most at low production and minor transportation costs), except for the initial investment for the installation and continued maintenance of power facilities, for the growth of rape and its minor local transportation costs. The sources are inexhaustible, and through combining electric and heat production the technical efficiency is well over 90 %.

Since windcraft, photovoltaics and BHPPs are typically *widely distributed* and dispersed, their combined effect may well be used to guarantee a stable supply.

Decentralized power distribution and management. Traditionally, electric power production and distribution are handled in a centralized manner. It is a *top-down* procedure, from the 500 kV, 110 kV, 20 kV levels, down to the 0.4 kV grids. On the other hand, the lack of timely prediction about local or regional consumption peaks requires a very conservative planning of reserve capacities. Also, due to technical constraints in large power plants (like long start-up and shut-down times with extensive maintenance and decreased life times) the generators would run continuously, thereby creating a considerable reserve capacity that may never be used: a built-in waste of energy. Finally power failures and energy balancing in large grids, if globally handled, are hard to manage as e.g. recently proven through the catastrophic black-outs in the Eastern US and Canada or, lately, the crashing of dozens of huge power line pylons resulting from heavy icy rain in Germany. The centralized control concept causes high overhead costs, high inflexibility and a lack of scalability and fault tolerance. Even under of highly efficient computer control, timely reactions and negotiations could only be achieved through a *decentralized management*.

Instead, in this paper we present a *bottom-up* principle of power distribution and balancing, as part of a completely decentralized management of renewable electric energy production and consumption. For the sake of higher fault tolerance it even *exploits the widely distributed renewable source structure as a basis for efficient fault control*: Failures would have a limited local or regional impact only, and, while also a consumer, every energy producer represents a potential back-up/ reserve facility.

We have introduced a staged management of the electrical power grid with 4 operating stages. Due to page limitations we only describe the negotiation phase among the involved agents. The other 3 phases the safe distribution and management of the electric power. For more details we refer the reader to [HKW+06].- Our DEZENT algorithm is, to the best of our knowledge, the first completely decentralized solution for these problems.

2. Distributed Agent Negotiations in DEZENT

2.1 The Model

2.1.1 General assumptions

Under the assumption that the overall power needs in a region can be covered through renewable sources (which is already realized in quite a number of towns in Southern Germany) the customers should negotiate the prices for electric power themselves, even more so since consumers are (potentially) also producers. Due to their common interests or double roles, respectively, we consider as their dominant attitude:

- to satisfy their needs under minimal investments;
- to rely on excess power in case of failures or shortcomings while in turn providing excess production to neighbors on demand (*balancing-in-the-small or bottom-up*).

As renewable energy comes for free or at moderate prices (e.g. bio-gas) and covering is secured, unused excess production would not be an issue except for user investment while balancing power is a local or at most regional business. While balancing may proceed on several levels negotiations on each level run in parallel on each level. In the worst case an extra local or regional reserve capacity (regenerative or not) may come into the picture (see Figure 1).

We follow basic requirements of fairness:

- to negotiate and distribute every portion of consumer demand, or of produced power;
- to take into account their unpredictable variation.

We respect these by choosing the smallest perceivable action possible as *negotiation period*. This will be 0.5 sec (e.g. the latency of a light switch). During this time we assume the demand and supply situation to be constant. In other words: *Changes occurring in the meantime will be accumulated until the next period*. This imposes *narrow and hard end-to-end deadlines both on the negotiation and power distributions processes*. The algorithms we present below have been designed according to these constraints, i.e. as *real-time adaptive solutions* of the problem phrased.

Negotiation processes on behalf of actors (human or technical) will be carried out through *distributed software agents* since they will take place well below the level of actor perception or reaction.

As common in Electrical Engineering, electric energy will be partitioned into arbitrary portions, according to needs and supply. Since the actor latency (e.g. a switch action) will be not less than 0.5 sec until the requested action is in effect we will assume that during this interval the need and supply situation is constant. All energy is available in the whole network. The underlying electric grid structure is supposed to be free of failures. This opens the door for participants acting under their own responsibility yet poses particular novel challenges on an appropriate handling of unpredictable consumer requests and producer offers, under fine-grained time-critical and stringent fault tolerance constraints.

2.1.2 Agent Negotiation Structure

In DEZENT distributed agent negotiations take place on multiple levels within subdivisions of the total agent population. Within these subdivisions (balancing groups) negotiations are carried out through balancing group managers (BGMs). While monitoring bids and offers, BGMs will arrange for contracts on power quantities on the basis of “close” matches of bids and offers (see Figure 1).

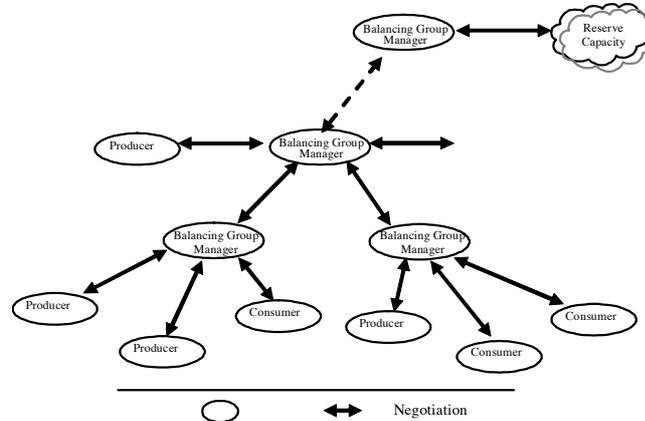


Figure 1: Negotiation Topology

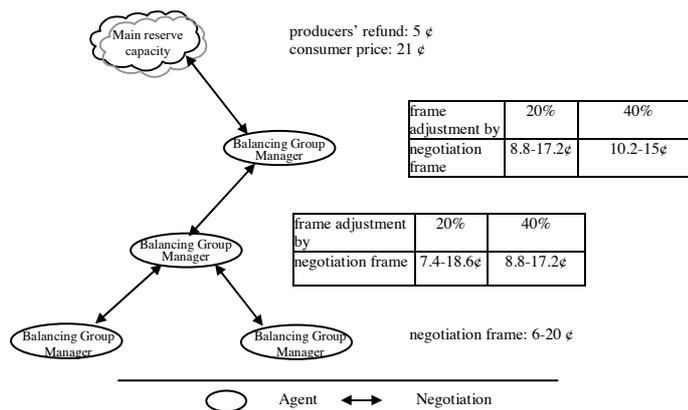


Figure 2: Negotiation Frames and Adjustment

Negotiations will start independently for the groups on the lowest level (each corresponding to a balancing group on the lowest (0.4kV) voltage level). If a balance cannot be found for all processes in a group the negotiation scope will be extended to the other groups on the same level, or higher up, under the control of the next-higher BGM. The purpose is to accommodate the unsatisfied processes. Only in the worst case will the back-up services be utilized (2.1.1).

Since an actor may be a producer and a consumer at the same time negotiations are initialized as follows:

The customer agent, after having computed the difference $current_needs - current_production$ acts as a producer agent if the difference is negative, as a consumer agent if the difference is positive, and it does nothing if the difference is zero.

During each negotiation period consumers issue bids for energy quantities they need, producers offer rates to sell such quantities. All quantities are limited to the next negotiation period. Since we assume the need and supply situation not to change during the period under discussion the price for a quantity will not depend on its size, in other words: According to the spirit of the approach **there are no long-term negotiations or discounts.**

As costs for producers arise just for amortization and maintenance a limited negotiation range is deemed appropriate. Within the given range consumers will tend to issue bids on the low side, producers will try to offer power for relatively high rates, each group according to their interests. As the negotiations proceed and unless a deal has been closed producer/ consumer rates are lowered, or raised, respectively, from step to step, in order to be finished before a negotiation cycle is finished. The urge is motivated by the fact that for the next cycle the yet unsatisfied processes would face a narrower negotiation range and additional charges that account for estimated power losses calculated from the supply configuration of the previous negotiation period. Thus both sides are put to a disadvantage.

2.2 The Base Negotiation Algorithm

2.2.1 Negotiation period

As just explained there are producer/ consumer agents and balancing group agents. The latter conduct negotiations between producers and consumers on various levels (see Figure 2). On each level negotiations are performed in cycles of 10 steps each. *For the purpose of simplicity we assume that in the*

model presented, based on synchronized clocks, negotiations in each cycle under a BGM start at the same time, and the duration of a step is 1 ms. (This still allows balancing group managers to process the requests of a large number of customer agents (up to 10^4)). After reaching the highest level (level 3 in Figure 2) negotiations will be finished since the remainder needs and power quantities will be handled by the main reserve facility. No new customers will be admitted during this period. We call this a **negotiation period**. Customers who have been satisfied during the negotiation period do no longer participate. As a consequence, a producer cannot act as a consumer during a negotiation period, and vice versa.

2.2.2 Price frames and adjustments

Negotiations on each level are held within fixed price frames. Frames on the same level have identical sizes. Customers that are unsatisfied after a cycle of one level will continue negotiations on the next higher level, however, the negotiation frames are shrunk by a fixed shrinking value Sr for all levels (See the example in Figure 2 with a shrinking value of 20% and 40%, respectively), lowering or raising the upper and lower limits, respectively, by half of the percentage. We do not only finalize on matching pairs of bids and offers but also consider bids and offers for contracting that are *similar* as specified by preset limits for their differences. The finalized energy price on level k is then calculated by adding $A_k - A_{k-1}$ to the arithmetic mean value between the similar bid and offer for the consumer and by subtracting $B_{k-1} - B_k$ from the mean value for the producer (for $k \geq 1$).

Let a current frame at a negotiation level k be denoted by $[A_k, B_k]$; $k = 0, 1, 2, \dots$. For a producer/consumer the minimum offer/maximum bid will be A_k/B_k , respectively. The opening bid bid_0 has to be chosen from $[A_k, \frac{1}{2}(B_k + A_k)]$, the opening offer $offer_0$ is taken from $[\frac{1}{2}(B_k + A_k), B_k]$.

Each agent also specifies a device-specific urgency urg_0 and strategy parameters s_{1C} and t_{1P} . They characterize the gradient of the bidding and offer curves, respectively. When after step n ; $n \in [0, 9]$ the unsatisfied agents adjust their bids/offers this will be done according to:

$$bid_C(n) = -\frac{1}{e^{\frac{urg_0 \cdot n}{s_{1C}} + s_{2C}}} + B_k \quad (1)$$

$$offer_P(n) = \frac{1}{e^{\frac{urg_0 \cdot n}{t_{1P}} + t_{2P}}} + A_k \quad (2)$$

The s_{2C} and t_{2P} are determined by the opening bid ($bid_C(0) = bid_0$) or offer ($offer_P(0) = offer_0$), respectively.

$$s_{2C} = -\log(B_k - bid_0), \quad (3)$$

$$t_{2P} = -\log(offer_0 - A_k), \quad (4)$$

Bidding and offering curves, after starting from their opening values, asymptotically approach B_k and A_k , respectively, with increasing n .

2.2.3 Negotiation period

At the end of each step unsatisfied consumers are identified and sorted by the BGM for current level k , according to their current bids. Then the consumers are processed top-down starting with the highest

bidding consumer. Offers similar to the bid of the first consumer are identified and sorted by price. Offers are processed top-down as well.

For closing a contract between the first-listed consumer and the first-listed producer the needs of the consumer will be fulfilled as far as possible, within the following constraint: Only up to X Wh (Watt-hours) will be granted to the consumer at a time. (This prevents any consumer to purchase a very high amount of energy thus leaving other consumers out in the cold!) After purchasing X Wh from one or more producers the current consumer's negotiation is interrupted, and the algorithm proceeds with the next-listed consumer. After processing the last-listed consumer, the algorithm starts again with first interrupted consumer (from the top of the list), allowing it to continue its negotiation for up to another X Wh. Going through the customer procedure again it proceeds until no match can be found in the current step any more. The algorithm stops and proceeds with the next step (and the afore mentioned bid/offer adjustments). (This approach is quite similar to the Round Robin mechanism found in process-scheduling to maximize CPU-utilization and to prevent starvation of late or low-priority processes).

When a contract between a similar bid and offer has been closed (*on a maximum of X Wh!*) we distinguish the following cases for handling the total quantities:

- A. The needed quantity is only a fraction of the offered size. The offer of the producer is adjusted to the difference of the need and the current size of the offer, the highest bidder is deleted. The algorithm proceeds with the next consumer.
- B. The offered quantity matches the needed one exactly. Producer and consumer are deleted, and the algorithm proceeds with the next consumer.
- C. The needed quantity is not completely covered by the offer. The need of the bidder is adjusted to the difference of his current need and the given offer size, the producer is deleted, and the algorithm proceeds to identify the next similar producer.
- D. If the need of the consumer is not yet satisfied but no similar offers are identified or left, the algorithm proceeds with the next consumer.

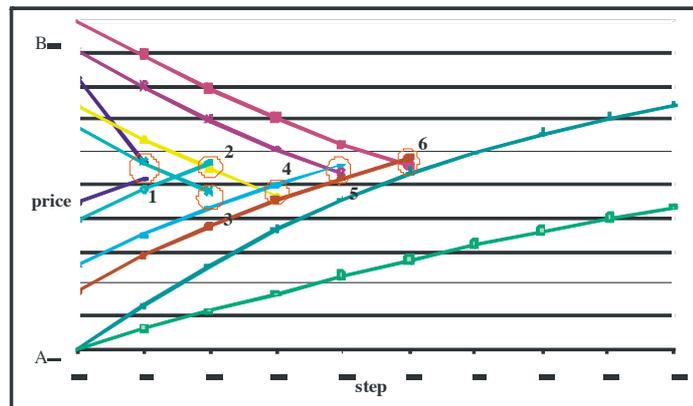


Figure 3: Contracting for Energy Quantities

Figure 3 illustrates the progression of the negotiation algorithm under a BGM during one cycle. In this example there are 6 consumers (ascending curves) and 5 producers (descending curves) participating. Encircled bid/offer pairs (of similar values) and numbers correspond to the order in which contracts are closed. According to the first three cases of the afore mentioned algorithm, on contracting either the

consumer curve ends (contract 2) due to needed quantities smaller than offers, or the producer curve ends (contracts 3, 4) due to offers smaller than needed quantities.

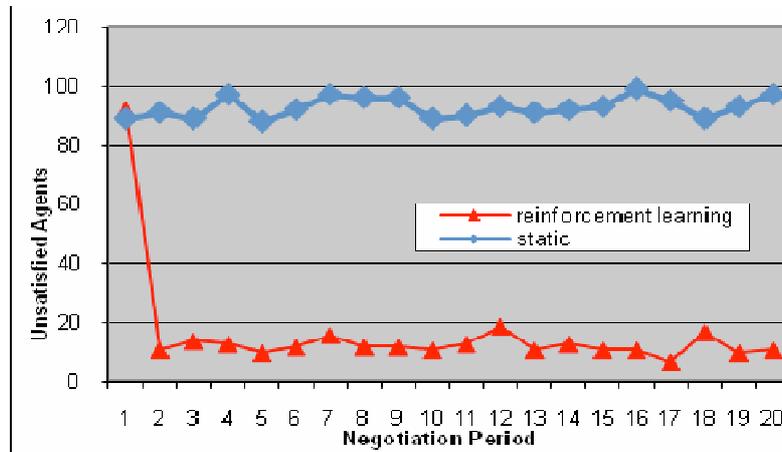


Figure 4: Unsatisfied Agents

Finally both curves end since needed and offered quantities match exactly (contracts 1, 5, 6). In this example two consumers remain unsatisfied by the end of the tenth step. They start with the lowest possible bid and do not adapt fast. Conversely, the curves that are contracted at 1 start rather high (consumer) or low (Producer), respectively and they adapt their values very fast.

The algorithm proceeds from level k to level $(k+1)$ if unsatisfied users are left on level k . The BGM on level $(k+1)$ starts with collecting all these users. Then BGM checks for each opening offer or bid from level k whether or not it fits into $[\frac{1}{2}(B_{k+1} + A_{k+1}), B_{k+1}]$ or $[A_{k+1}, \frac{1}{2}(B_{k+1} + A_{k+1})]$, respectively.

If the check is positive the values remain unchanged for level $(k+1)$. Otherwise the value would be outside of $[A_{k+1}, B_{k+1}]$, and the opening bid/ offer will be adjusted to A_{k+1}/ B_{k+1} , respectively (in this case s_{2C}/ t_{2P} have to be recalculated according to (3) and (4)).

Calculated power losses within a cycle on the k^{th} level are contracted out on level $(k+1)$ by a consumer agent acting on behalf of the BGM of that cycle, obeying all of the above mentioned constraints.

For further studies on the favorable cost behavior in DEZENT please see [WHL+06a, WHL+06b]. Also, for the results that the algorithm is *robust* against a large class of security attacks we refer the reader to the same references, due to space limitations.

3. Periodic Reinforcement Learning in DEZENT

As explained in 2.2 negotiations are organized in cycles, and the strategies within a cycle as well as strategy adjustments (of the negotiation frames) between cycles are automated. Typically only 3-4 cycles (of 30-40 msec total duration) are needed for finalizing the negotiation process, resulting in covering the consumer needs with regenerative power, without accessing traditional (reserve) power sources (see Figure 3).

Between periods, i.e. every 0.5 sec (2.2.1), a different form of adaptation has been established. It is based on distributed learning concepts which do not require (globally organized) training for the agents. Instead, we have derived novel AI techniques from the methods of *Reinforcement Learning* [SB98].

In the sequel, we will briefly outline the basic ideas of our approach. A technical representation is subject of a forthcoming publication.

Reinforcement Learning is a computational approach for understanding and automating goal-directed learning and decision making. It focuses on individual learning from direct interaction with the individual's environment. This is different from *supervised learning*, the traditional form of learning studied in most forms of Machine Learning, Statistical Pattern Recognition, and Artificial Neural Networks. These AI approaches are important examples of learning, but not really adequate for agent based learning and learning from interaction [SB98]: In interactive problems it is often impossible to obtain examples of a desired behavior that are both correct and representative of global situations and requirements where agents have to act. Under mostly unpredictable interaction between distributed agents each agent is left with learning *from its own experience*.

For a typical learning problem an agent is faced repeatedly with a choice among different actions. If for an action a numerical reward is to be received which depends directly on this action the purely *evaluative* feedback indicates how good a chosen action is, but not whether it is the best or worst possible action. Purely *instructive* feedback, on the other hand, indicates the correct action to choose, unaffected by the action actually chosen.

One of the challenges that arise in reinforcement learning and not in other kinds of learning is the trade-off between *exploration* and *exploitation*. To obtain a high reward, an agent will prefer actions that it has tried in the past and found to be efficient for producing a reward. Yet for discovering highly rewarded actions, it may have to try actions that it has not selected before. So the agent has to *exploit* the past but it also has to *explore* actions in order to make better action selections in the future. The dilemma is that neither exploration nor exploitation can be pursued exclusively without failing at the task at hand. Each agent will try a variety of actions *and* progressively favor those that appear to be best. On a stochastic task, each action must be tried many times to gain a reliable estimate of its expected reward [SB98]. *Exploitation* is the right thing to do for maximizing/minimizing the expected reward during the next period, but *exploration* may produce the higher total reward in the long run.

In DEZENT each combination of an opening bid or offer, $bid_{C_i}(0)$ or $offer_{P_i}(0)$ respectively, and an element of a *finite set* of strategy parameters s_{1C_i} or t_{1P_i} , respectively, is called a *strategy* and corresponds to a possible action to be chosen by an agent for the next negotiation period. (A strategy corresponds to a producer or consumer curve.)

For a consumer C_i 's strategy $strategy_{C_i}$, we have

$$\begin{aligned} s_{1C_i} &\in [\underline{s}_{1C_i}, \bar{s}_{1C_i}] =: \mathbf{S}_{C_i}, \\ bid_{C_i}(0) &\in [A_0, \frac{1}{2}(B_0 + A_0)] =: \mathbf{O}_{C_i}, \\ strategy_{C_i} &\in [\mathbf{S}_{C_i} \times \mathbf{O}_{C_i}] =: StrategySpace_{C_i}. \end{aligned} \quad (5)$$

For a producer P_i 's strategy $strategy_{P_i}$, we have

$$\begin{aligned} t_{1P_i} &\in [\underline{t}_{1P_i}, \bar{t}_{1P_i}] =: \mathbf{T}_{P_i}, \\ offer_{P_i}(0) &\in [\frac{1}{2}(B_0 + A_0), B_0] =: \mathbf{O}_{P_i}, \\ strategy_{P_i} &\in [\mathbf{T}_{P_i} \times \mathbf{O}_{P_i}] =: StrategySpace_{P_i}. \end{aligned} \quad (6)$$

After each period an agent considers the contracted rate which partially depends on the chosen strategy, otherwise on the producer strategies involved. In order for consumers (producers) to keep their power costs low (their reimbursement for investment and maintenance covered) they will assume as a main idea pursued in reinforcement learning that strategies resulting in low energy rates (for consumers, high rates for producers) should more likely recur than those followed by high prices. But how can an agent determine whether a price is high or low? (Please remember that the individual demands or supplies

may vary unpredictably.) A judgment regarding a good negotiated energy price during the next period will be made by comparison against a *reference price*. A simple estimate of the *reference price* is an average of previously negotiated energy prices. This method is called *reinforcement comparison*. We weigh recent energy prices higher than long-past ones, by means of a constant *step-size* parameter. Then the reference price $\bar{r}_{C_i}(t)$ for a Consumer C_i in period t is calculated by C_i as follows:

$$\begin{aligned}
\bar{r}_{C_i}(t) &= \bar{r}_{C_i}(t-1) + \alpha(r_{C_i}(t-1) - \bar{r}_{C_i}(t-1)) \\
&= \alpha r_{C_i}(t-1) + (1-\alpha)\bar{r}_{C_i}(t-1) \\
&= \alpha r_{C_i}(t-1) + (1-\alpha)\alpha r_{C_i}(t-2) + (1-\alpha)^2 \bar{r}_{C_i}(t-2) \\
&= (1-\alpha)^t \bar{r}_{C_i}(0) + \sum_{n=1}^{t-1} (1-\alpha)^{t-(1+n)} r_{C_i}(n)
\end{aligned} \tag{7}$$

where $\alpha, 0 < \alpha \leq 1$ is the step-size parameter and $r_{C_i}(t)$ the energy price negotiated in period t . The *reference price* $\bar{r}_{C_i}(0)$ is initialized with the current price frame's mean energy price, suggesting a balanced supply situation:

$$\bar{r}_{C_i}(0) = \frac{B_0 + A_0}{2} \tag{8}$$

Calculation of a reference price is done individually by every consumer C_i and every producer P_i . Reinforcement comparison is then used by every agent to update the new estimated energy price or *strategy preference* $p_{C_i, P_i}(t+1, a_i)$ of a strategy a chosen in period t , resulting for C_i in

$$p_{C_i}(t+1, a_i) = p_{C_i}(t, a_i) + \beta [\bar{r}_{C_i}(t) - r_{C_i}(t)] \tag{9}$$

On this basis chose between two *exploitation* actions will be derived, by modifying both the strategic parameters (flattening or steeping the curves) and the opening bid ($bid_{C_i}(0)$ or offer $offer_{P_i}(0)$, respectively) for the next period.

The *exploration* action is a simple trial-and-error approach where an agent, starting from its last strategy, randomly selects a strategy in its near neighborhood. Two consumer strategies $strategy_{C_i}'$ and $strategy_{C_i}$ are *neighbors* in our current model inf.:

$$|s_{1C_i}' - s_{1C_i}| \leq \frac{1}{4} |S_{C_i}| \quad \text{and} \quad |bid_{C_i}(0)' - bid_{C_i}(0)| \leq \frac{1}{4} |O_{C_i}|$$

The neighborhood between producer strategies $strategy_{P_i}'$ and $strategy_{P_i}$ is defined accordingly. The selection among the conflicting *exploitation/ exploration* actions is done by utilizing a common randomized action selection method.

This learning method has been verified to greatly enhance the situation of customer agents compared to various elementary strategy adjustment procedures. A typical chart displaying the difference between the static strategy and one that utilizes Reinforcement Learning is in Figure 4. (Please not that "unsatisfied" agents in our model will still be served by the main reserve capacity (Figure 1,2)). Due to page limitations, we have to refer the reader to [WHL+07b] for further details.

4. Conclusion

We have defined a novel distributed real-time negotiation procedure for agents taking care of producers and consumers of renewable energy on a large scale. Producer and consumer agents act in their own responsibility, and so do the actors behind, although the agent actions are far below the level of human perception. The base negotiation algorithm is very flexible, and Reinforcement Learning methods added considerably to the adaptive performance of the agents thus creating the basis for a superior form of innovative distributed power grid management. Different from traditional power distribution, while balancing in DEZENT is arranged bottom-up, local failures do not cause global blackouts. Since the electric distribution management can be equally finalized within each period [HKW+06] we have laid the ground for a thorough provision with sustainable and clean electric energy.

Bibliography

- [BuB05] Burnett, M.B. Borle, L.J.: “A power system combining batteries and super-capacitors in a solar/hydrogen hybrid electric vehicle”; Proc. of the Vehicle Power and Propulsion Conference, September 7-9, 2006; IEEE Computer Society Press
- [HKW+06] E. Handschin, O. Krause, H.F. Wedde, S. Lehnhoff: „The Emerging Communication Architecture in Electrical Energy Supply and its Implications”. Proceedings CRIS Workshop 2006, Influence of the Distributed and Renewable Generation of Power System Security, Magdeburg, Germany, 2006-6.-8. Dec.
- [SB98] R.S. Sutton, A.G. Barto, “Reinforcement Learning”; MIT Press, Cambridge 1998
- [UCTE05] Union for the Coordination of Transmission of Electricity; “Seven Actions for a Successful Integration of Wind Power into European Electricity Systems. Press Release May 17, 2005. URL: <http://www.ucte.org>
- [WLH+06a] H.F. Wedde, S. Lehnhoff, E. Handschin, D. König, O. Krause: “Real-Time Multi-Agent Support for Decentralized Management of Electric Power”; Proc. of the ECRTS’2006, July 5-7, 2006. Dresden/Germany; IEEE Computer Society Press
- [WLH+06b] H. F. Wedde, S. Lehnhoff, E. Handschin, O. Krause: “A Distributed Real-Time Multi-Agent System for Decentralized Management of Electric Power”; Technical Report #809, Computer Science, University of Dortmund, Feb. 2006; URL: <http://ls3-www.cs.uni-dortmund.de/Publications/pdf/techreport809.pdf>
- [WLH+07a] H. F. Wedde, S. Lehnhoff, E. Handschin, O. Krause: “Distributed Learning in Adaptive Real-Time Systems”; Submitted to ECRTS’07
- [WLH+07b] H. F. Wedde, S. Lehnhoff, E. Handschin, O. Krause: “Autonomous Real-Time Management of Unpredictable Power Needs and Supply”; Technical Report #810, Computer Science, University of Dortmund, Jan. 2007; URL: <http://ls3-www.cs.uni-dortmund.de/Publications/pdf/techreport810.pdf>