Parallel Multi-agent Smart Grid Simulation

Christian Kuschel*, Ulrich Rüde

Friedrich-Alexander-Universität Erlangen-Nürnberg, Lehrstuhl für Systemsimulation (Informatik 10), Cauerstr. 11, 91058 Erlangen, Germany; *christian.kuschel@fau.de

SNE 27(1), 2017, 17-23, DOI: 10.11128/sne.27.tn.10363
Received: February 10, 2017, (Selected ASIM STS 2016 Post-conf. Publ.), Accepted: February 25, 2017
SNE - Simulation Notes Europe, ARGESIM Publisher Vienna, ISSN Print 2305-9974, Online 2306-0271, www.sne-journal.org

Abstract. The approach in this paper renders it possible to simulate large-scale smart grids by efficient parallel computations. This permits a detailed analysis of the consumption behaviours, efficiency and impact of green energies, and self-sustainability of a smart grid. The smart grid is modelled as a multi-agent system. Each agent represents a building which is optimally controlled. That is, an agents meets its prescribed energy demand by trading energy or applying devices, e.g. solar panels and fuel cells, minimising its costs. A cooperative bargaining game is devised in which the agents participate to obtain a global optimal solution. In this paper, this inherently serial bargaining game is parallelised. The parallelisation is necessary to be able to deal with the large amount of data and computations which need to be performed. In the experiments the validity of the presented approach is shown and as a proof of concept a large smart grid of over 40 million agents is simulated.

Introduction

With technologies for integrating energy generation and storage in residential buildings, the notion of smart grid was derived [1, 2]. Smart grids have been in the focus of research because of various aspects. Among others, the energy generation and storage devices themselves, e.g. combined heat and power generation and energy storages, and the coordination of them pose many challenges [3, 4, 5, 6, 7]. The focus of this paper, however, is on the interactions of participants of the smart grid.

The smart grid is modelled by a multi-agent system. Each agent optimally controls a residential buildings and is able to generate and store energy with the appropriate devices [8, 9]. Based on the models of [10, 11], the agents are able to communicate with each other to participate in a cooperative bargaining game which possesses a unique Nash equilibrium. In addition, the pricing scheme from [13] is applied to obtain fair energy prices during bargaining.

The efforts of this paper go beyond the previously mentioned approaches of simulating smart grids: The smart grid model used in this paper extends the model [10] by incorporating a more detailed agent model which was published in [11]. In both publications a serial, weak coupling approach for solving the bargaining game is realised. The goal of this paper is enable simulations of realistically sized, large smart grids. Therefore the serial approach is parallelised for a distributed memory architecture, because a single computer can no longer deal with the intended complex and memory consuming simulation.

With this parallelisation, detailed information about the optimal, efficient usage of energy can be obtained. In particular, the efficient usage of green energies and the self-sustainability of the smart grid can be examined. To that end, a smart grid of more than 40 million agents, which is of equivalent size of Germany [12], is simulated in parallel. To the knowledge of the authors, this is the largest smart grid simulation performed until today.

In Section 1, a brief overview about the applied model and the serial algorithmic approach is given. The parallelisation of the cooperative bargaining game is presented in Section 2. Finally, the experimental validation and performance results are presented in Section 3. The test results validate the approach and show the parallel performance in strong scaling experiments.

1 Smart Grid Model

The smart grid model and the solution approach are essentially based on [10, 11, 13]. In the following a brief summary of the method published in [11] is given. Consider a smart grid as depicted in Figure 1.

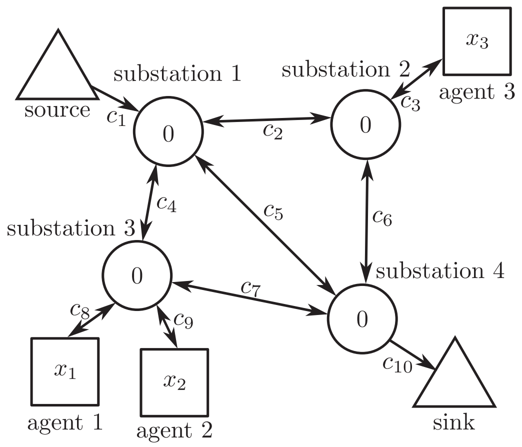

Figure 1: Schematic illustration of a smart grid with line capacities. The differently shaped nodes indicate different roles. The values within nodes specify the required net load.

In this paper the line capacities c_i of Figure 1 are assumed to be infinite, so that the grid does not pose any restrictions and can be neglected.

A smart grid consists of agents and a grid operator which are capable of communicating with each other. The grid operator uses power plants to generate energy and substations to distribute the energy. Since the line capacities are infinite, the substations can be neglected.

Buildings are optimally controlled by agents by a solving optimisation problem. The optimisation problem takes devices, like e.g. central heating, fuel cells, refrigerators, into account, and enforces that a prescribed energy demand is met. This is incorporated in constraints of the optimisation problem. The constraints define a non-empty, convex feasible set Ω_n for each agent n. The details about the feasible set are not relevant for the considerations below, but can be found in [11]. It is sufficient to assume that the set is non-empty (to ensure feasibility) and that the choice of the optimal solution is restricted in one way or the other (excluding a trivial solution).

The agents' optimisation problems are solved to minimise their respective costs. Since the goal is to simulate millions of agents, it is assumed that the agents form a market which determines the price with respect to demand and supply. Therefore, the agents can adapt their demand to the prices and influence the prices by changing the demand. This adaption is modelled by a cooperative bargaining game to minimise the agents' costs which is presented below.

The total incurring costs $\mathscr{C}$ of the budget balanced grid operator to supply the agents is defined by

$$\mathscr{C} := \left(\sum_{n=1}^{N} x_n \right)^2 ,$$

where N is the number of agents and x_n is the net load of agent n. The net load subsumes an agent's demand and supply in one variable. If $x_n > 0$ it represents a demand, if $x_n < 0$ it represents a supply. The cost function $\mathscr{C}$ can be understood as the squared deviation of a self-sustaining smart grid. If $\mathscr{C} \neq 0$ the grid operator must intervene and costs incur which must be covered by the agents.

Based on the fair pricing scheme suggested by [13], the total costs are split among the agents proportionally to their contribution to the total costs. This yields the individual incurring costs $\mathscr{C}_n$ of agent n with respect to all other agents' loads x_{-n}:

$$\mathscr{C}_n(x_n; x_{-n}) = \overbrace{(x_n + x_{-n})}^{\text{price}} \cdot \overbrace{x_n}^{\text{net load}} , \qquad (1)$$
$$\text{with } x_{-n} := \sum_{\substack{j=1 \\ j \neq n}}^{N} x_j .$$

So each agent minimises (1), the local objective function, in an optimisation problem. In particular, $\mathscr{C} = \sum_{n=1}^{N} \mathscr{C}_n(x_n; x_{-n})$.

In order to obtain a global optimum, a bargaining game is devised:

- Players: All agents in the smart grid.
- Strategies: Player n computes its best response $x_n = \arg\min \mathscr{C}_n(x_n, x_{-n})$ s.t. $x_n \in \Omega_n$.
- Costs: $\mathscr{C}_n$ for agent n.

For the considered game it was shown in [10] that a unique Nash equilibrium exists. Therefore, the locally computed optima by the agents lead to a global optimum. This Nash equilibrium is characterised by each player obtaining its optimum. Moreover, if one player deviates from the optimum solution, the costs increase for that player.

Since the costs of one agent are dependent on the net loads of all other agents, as indicated in (1), the game is played in rounds. In each round, all agents adjust their played net load to their respective current best responses after another. This is done until no adjustments are carried out by all agents. This procedure of the game is summarised in Algorithm 1. In the subsequent section, this algorithm is parallelised.

$k \leftarrow 0$
Set x^k to initial total grid load
repeat
  $\delta \leftarrow 0$
  $y^k \leftarrow x^k$
  for all agent n in the smart grid **do**
    $x^k_{-n} \leftarrow y^k - x^k_n$
    $x^{k+1}_n \leftarrow \arg\min \mathscr{C}_n(x^{k+1}_n; x^k_{-n})$ s.t. $x^{k+1}_n \in \Omega_n$
    $y^k \leftarrow x^{k+1}_n + x^k_{-n}$
    $\delta \leftarrow \delta + |x^k_n - x^{k+1}_n|/N$
  end for
  $x^{k+1} \leftarrow y^k$
  $k \leftarrow k+1$
until $\delta < \varepsilon$

Algorithm 1: Serial bargaining algorithm.

2 Parallel Bargaining Game

In order to parallelise Algorithm 1, the best response computation is distributed on processes which run in parallel. In general, more agents than processes are used, therefore one process is assigned multiple agents for computation. Thus each process computes the best responses of its agents serially as in Algorithm 1. This can be interpreted as computing a single best response for each process, independent of the actual number of agents associated with the process. So instead of referring to the agents' best responses, below the processes' best responses are considered.

An essential property of the game in Algorithm 1 is the successive best response computation. Evidently, when processes compute the best responses in parallel as described above, this property is violated. Consequently, for the parallel approach a synchronisation scheme needs to devised which is applied after the processes' best response computation to compensate the lack of the successive best response computations.

Consider again Algorithm 1. Since the best response computation is the solution to the agents' optimisation problems, this entails that the process' costs never increase. This non-increase property is expressed as

$$\mathscr{C}_n(x^{k+1}_n; x^k_{-n}) \leq \mathscr{C}_n(x^k_n; x^k_{-n}) \tag{2}$$

and must hold for every iteration k. This also implies that $|x^{k+1}_n| \leq |x^k_n|$ for all $n = 1, \ldots, N$. In fact, the objective value monotonically decreases until a global optimum has been found [10, 14].

The monotonic decrease in (2) holds for each process individually. However it must also hold for the total incurring costs in the smart grid (as in the serial case), therefore

$$\mathscr{C}^{k+1} \leq \mathscr{C}^k := \sum_{n=1}^{N} \mathscr{C}_n(x^k_n; x^k_{-n}). \tag{3}$$

In the parallel approach, a monotonic decrease in the total costs is obtained by computing the best response with respect to an auxiliary term $\bar{x}^k$, which represents the total net load in the grid. Let

$$\bar{x}^k_{-n} := \bar{x}^k - x^k_n.$$

So instead of minimising $\mathscr{C}_n(x^{k+1}_n; x^k_{-n})$, $\mathscr{C}_n(x^{k+1}_n; \bar{x}^k_{-n})$ is minimised.

The auxiliary term is defined as

$$\bar{x}^{k+1} = \frac{N-1}{N^2} \cdot \sum_{n=1}^{N} \left(x^{k+1}_n + \bar{x}^k_{-n} \right) + \frac{1}{N} \cdot \sum_{n=1}^{N} x^{k+1}_n. \tag{4}$$

It can be shown that (3) holds for this choice of $\bar{x}^k$. To show that $(\bar{x}^k)^2 \geq (\bar{x}^{k+1})^2$, the reduction γ is introduced to yield $(\bar{x}^k - \gamma)^2 = (\bar{x}^{k+1})^2$. Applying the definition of $\bar{x}^{k+1}$ from (4) reads

$$\left(\bar{x}^k - \gamma\right)^2 = \left(\bar{x}^k - \frac{1}{N}\bar{x}^k + \frac{N-1}{N^2} \cdot \sum_{n=1}^{N} \left(\frac{2 \cdot N - 1}{N-1} \cdot x^{k+1}_n - x^k_n \right) \right)^2.$$

Obviously,

$$\gamma = \frac{1}{N}\bar{x}^k - \frac{N-1}{N^2} \cdot \sum_{n=1}^{N} \left(\frac{2 \cdot N - 1}{N-1} \cdot x^{k+1}_n - x^k_n \right) \tag{5}$$

must be within $0 \leq \gamma \leq 2 \cdot \bar{x}^k$ for $\bar{x}^k \geq 0$ or $0 \leq -\gamma \leq 2 \cdot \bar{x}^k$ for $\bar{x}^k \leq 0$ to yield a decrease in the objective value.

In the following the case of $\bar{x}^k \geq 0$ is elaborated. The case $\bar{x}^k \leq 0$ can be derived analogously. To show that $\gamma \geq 0$, (5) is rewritten as

$$\frac{1}{N} \cdot \left(\bar{x}^k - \sum_{n=1}^{N} x^{k+1}_n \right) - \frac{N-1}{N^2} \cdot \sum_{n=1}^{N} \left(x^{k+1}_n - x^k_n \right).$$

Both terms in parenthesis represent the change in the objective value from iteration k to $k+1$. The term in the first parenthesis is the change with respect to the auxiliary term $\bar{x}^k$ whereas the term in the second is the change of the best responses $\sum_{n=1}^{N} x_n^k$. Both expressions in parenthesis are negative, but due to the averaging of $\bar{x}^k$ the term in the first parenthesis is more inertial to change than the second term. Thus the absolute value of the term in second parenthesis is larger so that $\gamma \geq 0$ holds.

(5) is inserted to $\gamma \leq 2 \cdot \bar{x}^k$ to show this inequality also holds:

$$(2 \cdot N - 1) \cdot \bar{x}^k \geq \frac{2 \cdot N - 1}{N} \sum_{n=1}^{N} x_n^{k+1} - \frac{N-1}{N} \sum_{n=1}^{N} x_n^k.$$

Recall that $|x_n^{k+1}| \leq |x_n^k|$ from (2) is implied, so

$$(2 \cdot N - 1) \cdot \bar{x}^k \geq \left(\frac{2 \cdot N - 1}{N} - \frac{N-1}{N} \right) \sum_{n=1}^{N} x_n^k = \sum_{n=1}^{N} x_n^k$$

is obtained. In the course of iteration, $\bar{x}^k$ approximates $\sum_{n=1}^{N} x_n^k$. Consequently, both terms are within the same order and the estimate

$$\bar{x}^k \leq \sum_{n=1}^{N} x_n^k \leq 2\bar{x}^k < (2N-1) \cdot \bar{x}^k \Rightarrow \left| \bar{x}^k - \sum_{n=1}^{N} x_n^k \right| \leq |\bar{x}^k|$$

holds. $|\bar{x}^k - \sum_{n=1}^{N} x_n^k|$ is generally small and decreases with increasing number of iterations k, since the x_n^k are computed as solution to optimisation problems and used to construct $\bar{x}^k$. Therefore (4) yields a monotonic decrease.

With the monotonic decrease shown in each iteration, it yet needs to be validated that the sequence of $\bar{x}^k$ converges to the optimum. This is done by showing convergence of $\bar{x}^k \to \sum_{n=1}^{N} x_n^k$ for $k \to \infty$. By expressing $\bar{x}^k$ with respect to the best responses x_n^{k+1} and $\bar{x}^k$ the following recurrence equation for a given initial value $\bar{x}^0 = \sum_{n=1}^{N} x_n^0$ is obtained:

$$\bar{x}^{k+1} = \left(\frac{N-1}{N} \right)^{k+1} \cdot \bar{x}^0 + \frac{2 \cdot N - 1}{N^2} \cdot \sum_{n=1}^{N} x_n^{k+1} + \frac{N-1}{N^2} \cdot \sum_{m=1}^{k} \left(\frac{N-1}{N} \right)^{k-m+1} \cdot \sum_{n=1}^{N} x_n^m. \quad (6)$$

Since the sequence $\bar{x}^k$ is monotonically decreasing, i.e. converging to a finite limit, the subsequence x_n^m is also convergent. So the sum in (6) containing x_n^m can be bound by a geometric series (since $|1/N| < 1$) by replacing x_n^m with x_n^k. The geometric series can be rewritten in a closed form. Using $x^* = \sum_{n=1}^{N} x_n^k$ for $k \to \infty$, from (6)

$$\lim_{k \to \infty} \bar{x}^{k+1} = 0 + \left(\left(\frac{N-1}{N} \right)^2 + \frac{N-1}{N^2} + \frac{1}{N} \right) \cdot x^* = x^*$$

is obtained. Consequently, using an auxiliary term as defined in (4) implies convergence to the optimum. Thus, the serial bargaining game can be parallelised by this method.

Concluding this section, the parallel bargaining is summarised in Algorithm 2. The new variable P in line 15 denotes the total number of agents in the entire smart grid. This number is required to check the termination criterion in line 17.

1: Set X to the initial net load
2: $X \leftarrow \text{allReduce}(X, +)$
3: $k \leftarrow 0$
4: **repeat**
5: $\quad Y, \delta \leftarrow 0$
6: $\quad$ **for all** agent j of process n **do**
7: $\quad\quad X \leftarrow X - x_j^k$
8: $\quad\quad x_j^{k+1} \leftarrow \arg\min \mathscr{C}_j(x_j^{k+1}; X)$ s.t. $x_j^{k+1} \in \Omega_j$
9: $\quad\quad X \leftarrow X + x_j^{k+1}$
10: $\quad\quad Y \leftarrow Y + x_j^{k+1}$
11: $\quad\quad \delta \leftarrow \delta + |x_j^k - x_j^{k+1}|$
12: $\quad$ **end for**
13: $\quad X \leftarrow \text{allReduce}(X, +)$
14: $\quad Y \leftarrow \text{allReduce}(Y, +)$
15: $\quad \delta \leftarrow \text{allReduce}(\delta, +)/P$
16: $\quad X \leftarrow (N-1) \cdot X/N^2 + Y/N$
17: $\quad k \leftarrow k + 1$
18: **until** $\delta < \varepsilon$ and $|X - Y|/N < \varepsilon$

Algorithm 2: Parallel bargaining algorithm.

Each process performs the computations for its assigned agents independent of other processes on a distributed memory architecture. The "allreduce"-operation is the only operation which requires interprocess communication. It acts as barrier and all processes wait for each other at this point. When it is called, each process collects the passed value from the other processes and sums up the values, respectively. The return value of this function is the same for all processes.

Although in Algorithm 2 it is described that three communication steps in lines 14 to 16 are required, all reductions can be fused in only one reduction by writing all three variables in one buffer and perform one reduction on the buffer instead of one per variable.

3 Simulation Results

Experiments were conveyed to validate the approach of Section 1 and also analyse the parallel performance of the presented bargaining algorithm. The algorithms were implemented in C++ using the GNU linear programming kit (glpk) [15] as central solver for optimisation problems and Open MPI [16] for the parallelisation and interprocess communication.

In Section 3.1, the validity of the approach from Section 2 is shown. The parallel performance is analysed with respect to strong scaling in Section 3.2.

3.1 Validation of the approach

The validation was carried out for a reference smart grid with 10,000 agents. The smart grid is simulated for 100 time steps, which represents roughly one day in 15 minutes intervals. 200 bargaining iterations were performed in each case.

As shown in Figure 2, all plots converge to the reference result, which was computed serially by the approach of [11]. So the approach as described in Algorithm 2 allows to compute the optimum in parallel. It can be observed that, although convergent, the more processes are used, the slower the convergence. This can be traced back to the coefficients dependent on N in (6) which tend to 1 for large N. Therefore the more processes are used, the slower the convergence rate.

Although the plot shows that the number of required iterations to reach the optimum increases with increasing number of processes, the total runtime does not increase. These details of the performance are presented in the next section.

3.2 Parallel performance

The efficiency of the implementation is considered in the following. Therefore the strong scaling behaviour of the parallel bargaining approach is presented. The termination criterion was set to a maximal average change of net load per agent per time step of 10^{-2} kWh. Furthermore, the agents are uniformly distributed on the processes in a round robin fashion.

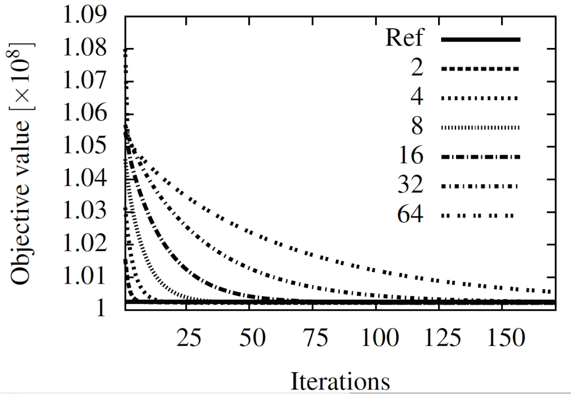

Figure 2: Convergence of parallel bargaining for 2 to 64 processes, and the reference solution.

The strong scaling test scenario is the same as in Section 3.1. It was computed on the local compute cluster of the department of computer science 10 of Friedrich-Alexander-Universität Erlangen-Nürnberg [17]. It has eight compute nodes each of which consists of four Intel Xeon E7-4830 (eight cores each) at 2.4 GHz and 256 GB RAM per node. The serially computed scenario took 42.5 minutes to finish.

As already mentioned, in Figure 2 it can be observed that the more processes are used the more iterations are required to attain the optimum. This, however, does not entail a longer total runtime, as can be seen in Table 1. In fact, since the optimisation problems of the agents are uniformly distributed, the time per iteration approximately halves when using twice as much processors. The number of iterations hardly increases with increasing number of processes. Consequently, the parallelisation has a greater impact than the number of increased iterations, therefore the total runtime decreases.

In Table 1, it is stated that there is close to perfect speed-up up to eight processes. Beyond that the speed-up of the computations increases less. The reason for that are the increased communication in addition to the fluctuating times to solution of an agent's optimisation problem. For the considered scenario the latter times are between 2 s and 0.05 s with a median of 0.07 s. Obviously, most of the agents' optimisation problems are solved quickly and few slowly. If the optimisation problems are distributed among the processes in a way that the average time to solve all optimisation problems per process is approximately equal, a good speed-up can be expected.

For the considered scenario more than 1,250 agents

# Processes	1	2	4	8	16	32
# iterations	3	4	5	5	6	6
it. time $[s]$	798	349	181	81	37	21
Speed-up	1	1.7	2.8	6.0	10.9	20.1

Table 1: Strong scaling results for 10,000 agents, showing the number of iterations, the averaged iteration time and the speed-up with respect to the serial computation.

per process yielded excellent scaling. Fewer agents caused some processes to wait for other, slower processes reducing the effects of parallelisation. This effect is expected to be alleviated by using dynamic instead of the static load balancing used in the experiments. Nevertheless the parallel efficiency does not drop below 60% in the experiments, which indicates good scalability.

In addition to the strong scaling computations, one large simulation as proof of concept for the realisation of a nation-sized smart grid was carried out. Algorithm 2 was executed on the Emmy cluster of Regionales Rechenzentrum Erlangen [18] simulating 40 million agents on 1280 processes. This was the largest simulation with respect to the physical limitations of the compute nodes and it completed within 16 hours. This scenario corresponds to a smart grid in the order of Germany, which is the country with the most private households in the EU [19]. To the knowledge of the authors, this is the largest simulated smart grid until today.

This carried out experiment merely hints at the potential of this approach. Being able to simulate smart grids representing countries, analysis of efficiency and sustainability can be performed. Especially if investments pay off can be analysed. From the simulation aspect, the weak scaling has yet to be examined. This is necessary to determine the efficiency of this approach with increasing number of agents on an increasing number of processes. This is analysed in the current work in progress.

4 Conclusion and Future Work

In this paper, a cooperative bargaining game was successfully parallelised. Since the serial algorithm relies on successive best response computation, it is necessary to introduce an auxiliary term to ensure convergence of the approach.

Almost perfect speed-up is reached in the strong scaling scenarios, in which 10,000 agents are considered, if the average time to solution per process is almost equal. This was commonly the case for more than 1,250 agents per process in the experiments.

In a next step the parallel efficiency needs to be evaluated for weak scaling scenarios. When simulating large smart grids, the presented algorithm is required to scale well with the problem size. In conclusion, the largest run optimised more than 40 million agents, a problem size equivalent to a country. Being able to simulate such a large grid, estimates on a realistic scale about self-sufficiency, demand and supply and efficiency of green energies can be fleshed out.

References

[1] Biswas MdM et al. *Towards Implementation of Smart Grid: An Updated Review on Electrical Energy Storage Systems*. In: Smart Grid and Renewable Energy 4.4 (2013), pp. 122–132.

[2] Monti A, Ponci F. *Power Grids of the Future: Why Smart Means Complex*. In: Complexity in Engineering, 2010. COMPENG '10. 2010, pp. 7–11.

[3] Houwing M, Negenborn R, De Schutter B. *Demand response with micro-CHP systems*. In: Proceedings of the IEEE 99.1 (Jan. 2011), pp. 200–213.

[4] Arif M, Oo A, Ali A. *Role of Energy Storage on Distribution Transformer Loading in Low Voltage Distribution Network*. In: Smart Grid and Renewable Energy 4.4 (2013), pp. 237–251.

[5] Atzeni I et al. *Demand-Side Management via Distributed Energy Generation and Storage Optimization*. In: Smart Grid, IEEE Transactions on 4.2 (2013), pp. 866–876.

[6] Giuntoli M, Poli D. *Optimized Thermal and Electrical Scheduling of a Large Scale Virtual Power Plant in the Presence of Energy Storages*. In: IEEE Transactions on Smart Grid 4.2 (2013), pp. 942–955.

[7] van de Ven P et al. *Optimal Control of End-User Energy Storage*. In: IEEE Transactions on Smart Grid 4.2 (2013), pp. 789–797.

[8] Karl J. *Dezentrale Energiesysteme: Neue Technologien im liberalisierten Energiemarkt*. Oldenbourg, 2006.

[9] Ekanayake J et al. *Smart Grid: Technology and Applications*. Wiley, 2012.

[10] Mohsenian-Rad AH et al. *Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid*. In: IEEE Transactions on Smart Grid 1.3 (2010), pp. 320–331.

[11] Kuschel C, Köstler H, Rüde U. *Multi-Energy Simulation of a Smart Grid with Optimal Local Demand and Supply Management*. In: Smart Grid and Renewable Energy 6 (2015), pp. 303–315.

[12] Bundesamt für Statistik. *Staat & Gesellschaft – Haushalte & Familien – Statistisches Bundesamt (Destatis)*. www.destatis.de/DE/ZahlenFakten/GesellschaftStaat/Bevoelkerung/HaushalteFamilien/HaushalteFamilien.html. Feb. 2016.

[13] Baharlouei Z et al. *Achieving Optimality and Fairness in Autonomous Demand Response: Benchmarks and Billing Mechanisms*. In: IEEE Transactions on Smart Grid 4.2 (2013), pp. 968–975.

[14] Berninghaus S, Ehrhart KM, Güth W. *Strategische Spiele*. Springer-Lehrbuch. Springer-Verlag Berlin Heidelberg, 2006.

[15] GLPK. *GNU Linear Programming Kit*. www.gnu.org/software/glpk. July 2015.

[16] The Open MPI Project. *Open MPI: Open Source High Performance Computing*. www.open-mpi.de. Mar. 2016.

[17] Computer Science 10, Friedrich-Alexander-University Erlangen-Nuremberg. *HPC Cluster*. www10.cs.fau.de/en/research/hpc-cluster. Feb. 2016.

[18] Regionales Rechenzentrum Erlangen. *RRZE – Emmy Cluster*. www.rrze.fau.de/dienste/arbeiten-rechnen/hpc/systeme/emmy-cluster.shtml. Feb. 2016.

[19] Bundesamt für Statistik. *EU – Privathaushalte in den Mitgliedsstaaten 2014*. de.statista.com/statistik/daten/studie/74720/umfrage/anzahl-der-privathaushalte-in-den-laendern-der-eu. Mar. 2016.