

Self-Organizing Demand Response with Comfort-Constrained Heat Pumps

Christian Hinrichs, Michael Sonnenschein
Environmental Informatics

Dept. of Comp. Science, University of Oldenburg
Oldenburg, Germany

Email: {hinrichs,sonnenschein}@informatik.uni-oldenburg.de

Adam Gray, Curran Crawford

Institute for Integrated Energy Systems
Dept. of Mech. Eng., University of Victoria
Victoria, BC, Canada

Email: {adamgray, curranc}@uvic.ca

Abstract—The increasing pervasion of information and communication technology (ICT) in the power grid motivates innovative research towards intelligent system control. Especially against the background of the growing share of renewable generation, novel approaches that mitigate e.g. grid expansion costs are of great interest. In this paper, the self-organized and thus completely distributed control concept COHDA for the online scheduling of responsive loads is employed as an alternative to the traditional centralized control. By distributing the decision making process, no global state information about the participating devices has to be gathered beforehand. The proposed approach is applied to the mitigation of wind power fluctuations on a minute basis by controllable heat pumps. Simulation results indicate that even in the presence of moderate time delays between wind power measurements and the resulting control action from the responsive population (i.e. including the distributed scheduling process), a significant reduction of fluctuations can be achieved.

I. INTRODUCTION

In many countries, the electrical power system is currently undergoing a fundamental transition from hierarchical top-down supply towards a decentralized structure based on small scale renewable energy generation, such as wind energy converter systems. The integration of such intermittent energy resources poses several challenges. For instance, as the grid itself is not capable of storing energy, supply and demand must be balanced at all times to ensure the safe operation of connected appliances. Traditionally, supply was scheduled with respect to predicted demand. The volatility of renewable energy generation, however, adds an increased supply-side uncertainty. Demand response (DR) is a possibility to mitigate this problem. DR has been discussed in the scientific literature and in field projects for several decades now, originally focusing on economic benefits [1]. With the increasing pervasion of renewable generation, research on DR for environmental as well as technical benefits has gained momentum, see e.g. [2].

This paper focuses on short term fluctuations of wind power supply, i.e. the reduction of variability in the residual load on a time scale of 1 minute. Due to the natural variance of wind speeds, the feed-in of wind turbines is unpredictable on this scale. Moreover, it has been shown that wind power variability extends to whole wind farms as well: based on long-range correlations observed in winds, wind farm variability effects do not average out, but persist up to several kilometers [3]. As explained above, DR can mitigate this problem. But appropriate controllable loads with sufficient ramping capabil-

ities are required in order to achieve the desired counter effects on the relevant time scale.

Thermostatically controlled loads (TCL) have been identified as suitable appliances for this task [4]. The operation of such devices rely on the thermal energy stored in the system. Due to thermal inertia, the system can be seen as a buffer store, effectively decoupling the actual operation of the device from its designed effects (i.e. heating or cooling). This inherent flexibility allows for using such appliances as regulation resources. In [5], a control strategy for heat pumps with programmable communicating thermostats is proposed. The strategy incorporates both system-level objectives and customer-level comfort constraints in the decision making process. Based on power density distribution functions calculated by a system-level central controller, the heat pumps' thermal trajectories are modified in such a way that global as well as local goals are met. For this, state information from each participating load is transferred to the central controller prior to calculating the control signals.

Against this background, the paper at hand focuses on two objectives: First, the self-organized and thus completely distributed scheduling heuristic COHDA [6] is employed as an alternative control strategy. While originally being designed for the predictive scheduling of decentralized energy resources in power markets, see [7], the approach is applied to the online-scheduling of controllable loads here. Compared to the aforementioned centralized controller, the participating appliances adapt their thermal trajectories autonomously based on limited exchanged information. This increased autonomy on the customer side can increase the willingness to participate in DR programs. The transfer of complete state space data is avoided, thus more privacy is given. Finally, self-organizing approaches can offer increased robustness over centralized methods, e.g. by healing regional outages autonomously. However, these benefits of the approach come at the expense of being an iterative heuristic despite being able to solve the problem optimally in constant time. While the effectiveness of the heuristic in terms of solution quality as well as its robustness against different sources of variation (e.g. communication network latency) have already been shown elsewhere [8], its varying run-time may pose significant problems in an online-scheduling context. Therefore, the second objective of this paper is to evaluate the effects of different time delays between wind fluctuation and the actual regulation response from the load population.

In Section II, the formal model of the proposed system is specified. This includes the heat pump control concept, the generation model for wind power as well as an overview on the COHDA scheduling approach. Subsequently, Section III presents the evaluation of the system. This comprises simulation setup, formulation of objective functions, results and discussion. Finally, Section IV concludes the paper.

II. FORMAL MODEL

As a prerequisite for the evaluation in Section III, the formal model of the proposed system is described in the following.

A. Heat Pump Model

To be consistent with the centralized control approach proposed in [5], the same heat pump model is employed in this paper. Thus, a one-dimensional equivalent thermal parameter (ETP) model as described in [9] is used to simulate the building heat dynamics with respect to the heat input \dot{q} versus the outdoor temperature θ_o . While θ_o is taken from environmental data, the heat input is determined by the operational heating power \dot{q}_{op} of the heat pump: $\dot{q} = n \cdot \dot{q}_{op}$ with $n \in \{0, 1\}$ as operational state. We consider air-source heat pumps only, thus the value of \dot{q}_{op} in turn depends on the outdoor temperature and is estimated using third order polynomials based on manufacturer's data sheets. For further implementation details, we kindly refer the reader to [10].

In [5], this model is then equipped with the method from [11] as a means for controlling heat pump operation. Accordingly, systematic deviations u to the nominal temperature set-points θ_s of the heat pumps are used to partially synchronize the operational states of a load population. For a single heat pump, the operational state n in the next time step $k + 1$ is then determined as follows:

$$n(k+1) = \begin{cases} 1 & \text{for } \theta_a(k) \leq \theta_s - \frac{\delta}{2} + u(k), \\ 0 & \text{for } \theta_a(k) \geq \theta_s + \frac{\delta}{2} + u(k), \\ n(k) & \text{otherwise} \end{cases} \quad (1)$$

where $\theta_a(k)$ depicts the measured indoor air temperature at the conclusion of interval k and δ constitutes the allowed temperature deadband. The operational state $n = 1$ refers to active (i. e. heating), while $n = 0$ refers to inactive. To prevent customer-side disruption as well as rapid-cycling of operational states (which impacts the service life of the devices), the deviations u are constrained to the quarter deadband width:

$$|u(k)| \leq \frac{\delta}{4} \quad (2)$$

With this method, appropriate control signals u can efficiently be calculated and communicated in large populations by a central controller. As the paper at hand focuses on a self-organized scheduling heuristic, where every heat pump controller decides autonomously about its subsequent operational state (cf. Section II-C), a control signal based on u is not used directly. Instead, by calculating the possible state space of (1) for each heat pump with respect to the boundary conditions given by (2), the allowed control actions for the COHDA heuristic are determined by the respective device controller autonomously, in every time step. This way, the same operational constraints as in [5] can be maintained.

B. Wind Power Model

High-resolution wind power data is difficult to obtain. On the other hand, environmental data on this scale is easily available. Therefore, wind power time series are calculated synthetically in this paper by simulating a wind turbine model based on a transfer function between the wind speed input and power output. The following formulation is directly taken from [10]: First, the mechanical power P_m extracted by a wind turbine is derived from a given wind speed u_w as:

$$P_m = \frac{1}{8} C_p \rho \pi D^2 u_w^3 \quad (3)$$

where D is the rotor diameter, ρ the air density, and C_p the turbine's coefficient of performance. For simplification, the latter is approximated using cut-in and cut-out wind speeds u_{ci} and u_{co} as follows:

$$C_p = \begin{cases} 0, & u_w < u_{ci}, \\ C_m (1 - F_1 U^2 - F_2 U^3), & u_{ci} \leq u_w \leq u_r, \\ C_r \left(\frac{u_r}{u_w}\right)^3, & u_r \leq u_w \leq u_{co}, \\ 0, & u_w > u_{co} \end{cases} \quad (4)$$

where C_m and C_r are the maximum and rated coefficients of performance, respectively, while u_m and u_r are the according wind speeds at C_m and C_r . With $U = \left(\frac{u_m}{u_w} - 1\right)$, the empirical coefficients F_1, F_2 were determined using the boundary conditions $C_p = 0$ at $u_w = u_{ci}$ and $C_p = C_r$ at $u_w = u_r$.

Second, the system is equipped with a 3-stage gearbox that is modeled as an inefficiency factor depending on the mechanical input, cf. [12]:

$$\eta_g = \begin{cases} 0.1 & \text{for } P_m \leq \frac{3 \cdot P_{rt}}{90}, \\ 1 - \frac{0.01 \cdot 3 \cdot P_{rt}}{P_m} & \text{otherwise} \end{cases} \quad (5)$$

where P_{rt} is the rated power of the turbine.

With a final inefficiency factor η_e for the electrical losses in the whole system, the total wind power is then given as:

$$P_w = \eta_e \eta_g P_m \quad (6)$$

C. COHDA

As an alternative to a centralized DR controller, a completely distributed and asynchronous approach is given with the *Combinatorial Optimization Heuristic for Distributed Agents* (COHDA), see [6], [8]. Therein, each participating DR resource (i. e. heat pump) is equipped with a software agent. Compared to a simple communicating device controller, an agent is capable of autonomously deciding upon control actions, while taking arbitrary data such as user preferences or messages from other agents into account. Based on this, the key concept of COHDA is an asynchronous iterative approximate best-response behavior, where each agent reacts to updated information from other agents by adapting the operational state of its controlled device with respect to the global target aggregate load for the population. For a time step k , let $P_l^{(target)}(k)$ be this target, e. g. announced by the grid operator, and let $a_i \in A$ be the set of agents. Further, let $S_i(k) \subseteq \{P_{l,i}^{(active)}(k), P_{l,i}^{(inactive)}(k)\}$ be the set of feasible load values resulting from the possible operational states $n_i(k)$

in this time step, calculated as described in Section II-A, for each device. Then the task is to find a combination $\gamma(k) = \{P_{l,1}(k), \dots, P_{l,|A|}(k)\}$ of feasible load values in such a way, that the resulting aggregate load from $\gamma(k)$ matches $P_l^{(target)}(k)$ as close as possible. To simplify the presentation, we use the terms operational state and load value for a heat pump synonymously in the following.

For solving this combinatorial optimization problem efficiently, the agents are placed in an artificial communication topology (e. g. a *small world* topology), such that each agent is connected to a non-empty subset of other agents. To compensate for the resulting non-global view on the system, each agent a_i collects two distinct sets of information: on the one hand the believed current configuration $\gamma_i(k)$ of the system (that is, the believed set of currently selected operational states of all agents), and on the other hand the best known combination $\gamma_i^*(k)$ of operational states with respect to $P_l^{(target)}(k)$ it has encountered so far. All agents $a_i \in A$ initially only know their own respective set of feasible operational states $S_i(k)$, and the difficulty of the problem is given by the distributed nature of the system in contrast to the task of finding a common allocation of operational states. Thus, the agents coordinate via message exchange. Beginning with an arbitrarily chosen agent (by announcing the global target), each agent a_i executes the following three steps:

- 1) **(update)** When an agent a_i receives information from one of its neighbors (say, a_j), it imports this information ($\gamma_j(k)$ and $\gamma_j^*(k)$) into its own knowledge base by updating $\gamma_i(k)$ and, if better, replacing $\gamma_i^*(k)$ with $\gamma_j^*(k)$.
- 2) **(choose)** The agent now adapts its own operational state according to the newly received information. If it is not able to improve the believed current system configuration $\gamma_i(k)$, the agent reverts its current operational state to the one stored in $\gamma_i^*(k)$, thus following the best solution it is aware of up until now.
- 3) **(publish)** If $\gamma_i(k)$ or $\gamma_i^*(k)$ has been modified in one of the previous steps, the agent finally publishes its knowledge base ($\gamma_i(k)$, including its own selected operational state, and $\gamma_i^*(k)$) to its neighbors.

The heuristic terminates when for all agents $\gamma(k)$ and $\gamma^*(k)$ are identical. At this point, $\gamma^*(k)$ is the final solution of the heuristic and contains exactly one operational state for each agent.

Please note that this procedure constitutes a *planning phase*. Consequently, the agents do not actually configure their controlled device with each execution of the `choose` step, as this would lead to rapidly switching operational states within the planning phase. Instead, the agents wait for the process to terminate, until the heat pump is configured with the finally selected operational state. However, as termination detection in a distributed setting is a nontrivial task, termination is detected and announced centrally by the simulation environment in the present case. This is fair, as it does not impact the results of our study. For real-world scenarios, established Distributed Termination Detection (DTD) algorithms can easily be employed. For instance, the approach in [13] has successfully been implemented for the COHDA heuristic [14].

Originally, this heuristic was designed for the predictive scheduling of decentralized energy resources in power markets [7]. In such scenarios, sequences of operational states for the units are determined over whole planning horizons in advance, which allows for incorporating state dependencies between time steps quite easily, such as e. g. minimal running times to prevent rapid-cycling. In the centralized approach [5], the control signal u just shifts the state space boundaries, which then yields slightly early or delayed state switches compared to the original progress. This way, rapid cycling is naturally prevented because a device has to traverse a large part of its temperature deadband before becoming responsive to a control signal u again. For the COHDA approach, the extension of the deadband width induced by u is used directly for the calculation of feasible operational states, cf. Section II-A, thereby dropping the rapid cycling prevention mechanism. To restore this property, a secondary objective has been incorporated into the COHDA approach as follows: whenever an agent during the `choose` step is about to choose an operational state that deviates from the currently configured one (i. e. $n(k) \neq n(k-1)$), this choice will be discarded with a predefined probability P_d . Because of the iterative nature of the approach, an agent will likely be asked several times to change its operational state, thus allowing for intermediate state switches nevertheless, if required for the global target. If P_d is chosen appropriately (e. g. $P_d = 0.5$), this effectively prevents rapid cycling again.

III. EVALUATION

For the evaluation study in this paper, wind speed and temperature data from spring 2014 with a resolution of 1 minute was obtained from the NREL NWTC M2 Tower [15]. From this, wind power data was generated using the simulation model in Section II-B, configured to represent a single 1.5 MW wind turbine. Similarly, a population of 1000 heterogeneous houses equipped with air-source heat pumps was generated by drawing normally distributed parameter values for the the model outlined in Section II-A, with averages as given in [10]. For each of the experiments in the following sections, the simulation models were initially run for 1440 time steps to ensure steady-state conditions. The experiments themselves also comprise a simulation horizon of 1440 time steps, i. e. one simulated day with a resolution of 1 minute. While present power markets usually operate on a coarser scale (e. g. 5 minutes or 15 minutes), we assume a power market operating on 1 minute here in order to meet the fluctuations on this scale. For the simulation horizon, both an uncontrolled reference simulation as well as a simulation with the COHDA scheduling approach were performed. Regarding the latter, in each time step the global target aggregate load was determined first (based on the current wind power, see below), which was subsequently fed into the COHDA heuristic. After termination of the heuristic, the house simulation models were executed for one time step, incorporating the control actions determined by COHDA.

In the following sections, positive power is always depicted positive, while load is represented as negative power.

A. General Effectiveness

Referring to the first objective of this paper, the general effectiveness of the proposed control concept is evaluated. According to [16], ancillary service types can be subdivided into *load-following* and *regulation*. The former usually addresses power balancing tasks on a scale of 5 to 30 minutes, while the latter responds to rapid fluctuations on a minute scale. By separating the power profile into base-load, load-following and regulation components, respective objective functions can be defined independently for each ancillary service type. According to the considered time scale of 1 minute in this paper, we focus on the regulation aspect in the following. By assuming that base-load and load-following components are properly optimized and provided elsewhere, this allows us to formulate the regulation objective function based on ramp-rate only, which is defined as the change of load over time. So instead of tracking the actual power profile to achieve a supply-demand matching, we neglect the base-load and load-following aspects and configure the control strategy to minimize the *power gradients* of the residual load. For each time step k , given the current wind power $P_w(k)$ as well as the wind power $P_w(k-1)$ and the aggregate load of the heat pump population $P_l(k-1)$ at the previous time step, the target aggregate load for the current time step is calculated as:

$$P_l^{(target)}(k) = P_l(k-1) - (P_w(k) - P_w(k-1)) \quad (7)$$

Consequently, we used the following metric to measure the performance of the simulations based on residual gradients. Given the residual load of both the reference case and the controlled case as $res_r = P_w + P_l^{(reference)}$ and $res_c = P_w + P_l^{(controlled)}$, respectively, the performance of the control strategy for a time span $[i, j]$ is calculated as:

$$ROF = \frac{\left(\sum_{k=i}^j |res_c(k) - res_c(k-1)| \right) \cdot 100}{\sum_{k=i}^j |res_r(k) - res_r(k-1)|} - 100 \quad (8)$$

This way, ROF can be interpreted as percentage reduction of fluctuation in res_c with respect to res_r and thus gives us an easily understandable performance measure. Moreover, the metric allows for comparing different scenario configurations, which will be important in the evaluation of the time lag in Section III-B.

Figure 1 shows the results for an exemplary simulation run (the delay between wind power measurements and the response from the load population was set to 10 seconds here, cf. Section III-B). The visualization comprises three parts, each focusing on a different aspect: The upper chart depicts both wind power (dashed line) and the dispatched heat pump power (solid line). Furthermore, the maximal possible ramp-rate (both up and down) of the load population is shown for each time step as shaded area surrounding the solid line. But as wind power and dispatched heat pump power are almost indistinguishable in this representation, the middle chart emphasizes the difference between those two time series. This can be interpreted as remaining objective error of the optimization in the controlled case. Finally, the lower chart depicts the actual residual load, i.e. the difference between wind power and heat pump load, and thus the remaining stress on the grid. This is shown for both the reference case and the controlled case.

In this simulation run, a performance of $ROF = 45.79$ was achieved, meaning that the fluctuations in the uncontrolled reference case could be mitigated by 45.79% by applying the control strategy in this scenario. Interestingly, several rather large spikes are present in the visualization of the objective error (middle chart). These are almost evenly distributed during time intervals with no wind power (for instance in the morning and evening hours), and vanish in times with considerably fluctuating wind power. This effect can be explained as follows: As the control strategy defines the target load based on the gradient of the wind power, the heat pump population tries to keep its aggregate load constant during times with no wind fluctuation. In combination with the secondary optimization objective of the control strategy (minimizing the amount of state changes), each heat pump will try to stay in its current operational state during such intervals as long as possible. Naturally, the temperature trajectory of the heat pumps will eventually cross the temperature dead band boundary, forcing the device to switch its operational state. Moreover, as the employed control method tends to synchronize devices [11], these state switches occur in a synchronized manner as well. The visualized ramp capacity in the upper chart as well as the sudden drops of the residual load in the lower chart support this hypothesis. On the other hand, during intervals with considerably fluctuating wind power, the control strategy performs exceptionally well. For a closer look, a magnified part of the simulation (time steps 800–1200) is shown in Figure 2. Here, the described effects are visible in detail. During intervals with highly fluctuating wind power, the dispatched load follows the target very closely. Further, even in the presence of a constant wind power, the dispatched load is generally able to follow the target, but is regularly interrupted by synchronized forced state switches.

In summary, the results indicate that an accurate online scheduling of responsive loads without a centralized controller is possible. In its current state, the approach performs well for highly fluctuating target trajectories, but does not cope well with constant targets. However, simple extensions to the objective function should be able to overcome this problem, e.g. by simply disabling the control approach during such intervals, or by desynchronizing the load population continuously as proposed in [17].

B. Effects of Time Delays

In this paper, short term fluctuations on a time scale of 1 minute are considered. As wind is highly variable on this scale, wind power supply fluctuates unpredictable even in whole wind farms on this scale as well [3]. For a DR approach to be effective, this requires the response from a controlled load population to be as quick as possible. With a central controller such as [5], response times primarily depend on the technical setup and thus can be estimated quite precisely. For a heuristic like COHDA, this is rather difficult, as the approach is designed to terminate autonomously in a self-organized manner, yielding varying run-times without any timeliness guarantees. While this is acceptable in the originally targeted predictive scheduling use case, it poses a significant problem in the present online scheduling use case. Therefore, the second objective of this paper is to evaluate the effects of different time delays between wind fluctuation and the actual regulation response from the load population.

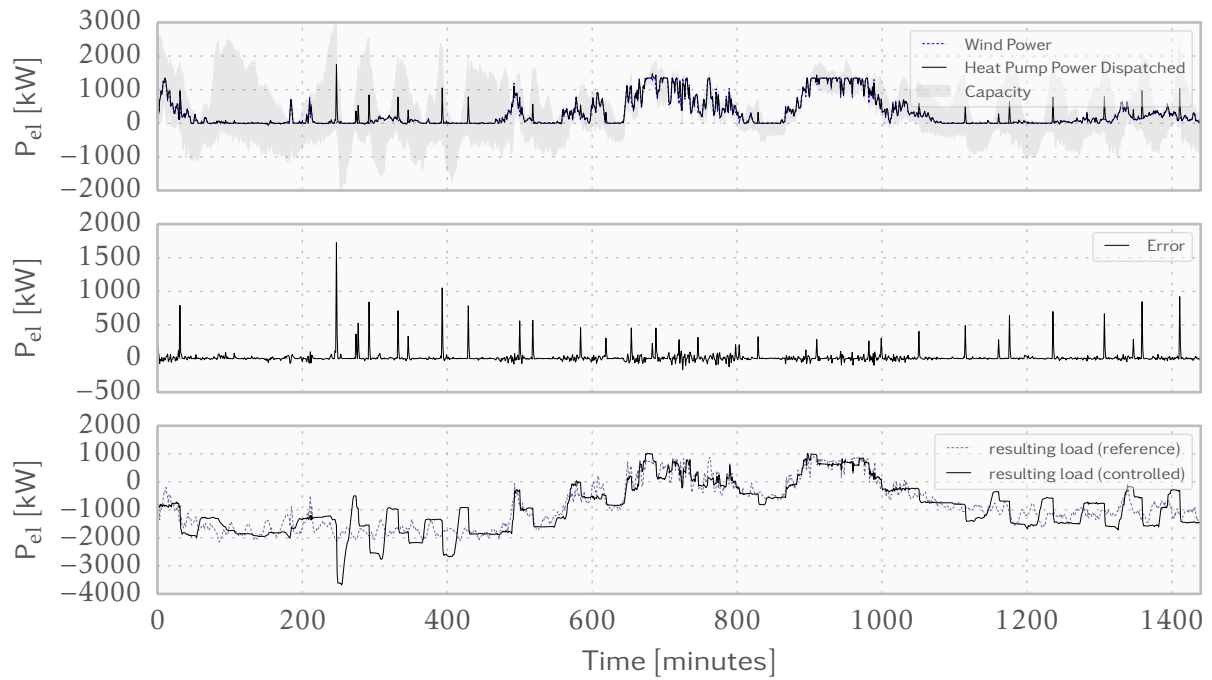


Fig. 1. Results of a single simulation run: wind power, dispatched heat pump power and ramp-rate (upper chart), objective error (middle chart), residual load (reference and controlled, lower chart).




Fig. 2. Results of a single simulation run, magnified part (time steps 800–1200): wind power, dispatched heat pump power and ramp-rate (upper chart), objective error (middle chart), residual load (reference and controlled, lower chart).

First of all, we artificially assume that the COHDA scheduling process will complete *instantly*, i. e. that it consumes no time at all. Then the simulation proceeds as follows: For each time step k , a wind power value $P_w(k)$ is measured (i. e. calculated from the wind power model). Based on this, the target aggregate load is determined according to (7), and COHDA is executed. The resulting heat pump operational states $n_i(k)$ are then fed into the heat pump model, which then yields the according response from the load population in the form of the aggregate load $P_l(k)$ for the current time step. $P_w(k)$ and $P_l(k)$ can then be evaluated directly using (8). However, as it is technically impossible for the scheduling to complete in zero time, this only serves as a reference case in our study. Now, to evaluate the effects of different time delays, we simply shift the time series P_l by fractions of minutes with respect to P_w . This is done by upsampling P_l to a resolution of 1 second, shifting it for a defined number of seconds and finally downsampling it again by averaging over every 60 intervals. This way, we can simulate arbitrary delays between wind power measurements and the according response from the load population, thus being able to analyze the effects of such delays on the quality of the response from the viewpoint of the power market.

Figure 3 shows the results of a series of experiments with increasing time delays, ranging from 0 seconds up to a maximum of 60 seconds with a step size of 1 second between the experiments. Each experiment was repeated ten times using the same environmental data (wind/temperature) and the same population characteristics, but with varied seeds for the random number generators in the simulation models and the heuristic. This leads to slightly different usage characteristics of the heat pumps as well as varying progress of the scheduling heuristic. In the artificial case of no delay at all, the achieved reduction of fluctuations is larger than 50% (median = 57.57%), but drops almost linearly in the crucial range of up to 30 seconds. At this point, a reduction of less than 20% is achieved (median = 16.85%). Starting with a delay of 37 seconds, the effects from the control actions significantly interfere with the ongoing fluctuations in wind power supply, actually leading to *increased* fluctuations in the residual load. With a delay of 52 seconds and more, no single simulation run was able to achieve a positive effect any more.

As a consequence, for the considered scenario the system should be designed in such a way that a control response can be expected after 30 seconds at the latest. Depending on the time required to obtain wind power measurements, this leaves even considerably less time for the scheduling process. The experiments in this paper have been performed in an artificial environment using a simulated communication infrastructure, therefore actual time measurements would not be meaningful. Instead, the number of sequential message exchanges (i. e. the amount of simulated time steps in the artificial communication backend) can be taken as an indicator for the total run-time. Due to the rather small solution space (the considered device models can either be switched on or off, thus having only two alternatives per time step), only 8 of these time steps were necessary on average until termination in our scenario. If we assume that a single message is transferred in at most 100 milliseconds, and add a rather large amount of 900 milliseconds of computational time per agent per time step, this still leads to a total scheduling time of less than 10 seconds. As shown

in [18], scheduling time increases logarithmically with larger population sizes. But the COHDA heuristic can be augmented with an early termination mechanism quite easily: If the agents are configured to stop creating new solution candidates after a predefined amount of time, the process terminates naturally after a few time steps from this point on.

In this consideration, potential communication link failures are not accounted for. For the centralized control method in [5], packet loss has been evaluated with respect to DR control accuracy [19]. Such an evaluation is not possible for the COHDA approach, as convergence in distributed algorithms cannot be guaranteed in the presence of packet loss (cf. the *FLP impossibility proof* by Fischer et al. [20]). Therefore, the task of handling link failures has to be deferred to the control layer of the communication protocol (e. g. as in [21]). While we do not consider packet loss in our study for this reason, we believe the heuristic would be able to self-heal automatically in many cases, e. g. due to redundant paths in the communication topology, or if the failure is only temporary.

IV. CONCLUSION

In this paper, the self-organized and thus completely distributed scheduling heuristic COHDA was employed as an alternative to traditional centralized control concepts in the Smart Grid. Originally designed for the predictive scheduling of energy resources with respect to market integration, the approach was applied to the online scheduling of responsive loads here, targeting the mitigation of wind power fluctuations on a minute-by-minute basis. After a recapitulation of the underlying simulation models, the results of a simulation study for this task have been discussed. Being a proof of concept, the study focused on a single type of deferrable load only. Thus, the integration of further models as well as other device classes such as electric vehicles is subject to future work. Moreover, while the approach performed exceptionally well for periods with highly variable wind power, performance significantly deteriorated for intervals of constant wind power. Thus the objective function in the COHDA process has to be revised in order to improve the decision making process towards such periods.

The second objective of this paper focused on the effects of time delays between announcing a control signal (i. e. the target aggregate load in the considered use case) and the actual response from the load community. Results indicate that, even with moderate delays, reasonable response effects can still be observed. As the employed COHDA heuristic constitutes an iterative approach, extensively based on communication, this is especially encouraging for this type of control method.

In summary, this paper demonstrated the application of the distributed scheduling heuristic COHDA in an *online scheduling* context. But the outlined approach is not restricted to this use case. As COHDA essentially constitutes a distributed combinatorial optimization heuristic, applications regarding this class of problems are generally possible. For example, in [7] the approach is employed in a *predictive scheduling* context within virtual power plants. On the other hand, the approach is employed in a *reactive scheduling* use case in [22, pp. 69–76], i. e. compensating for incidents during the operation of a virtual power plant.

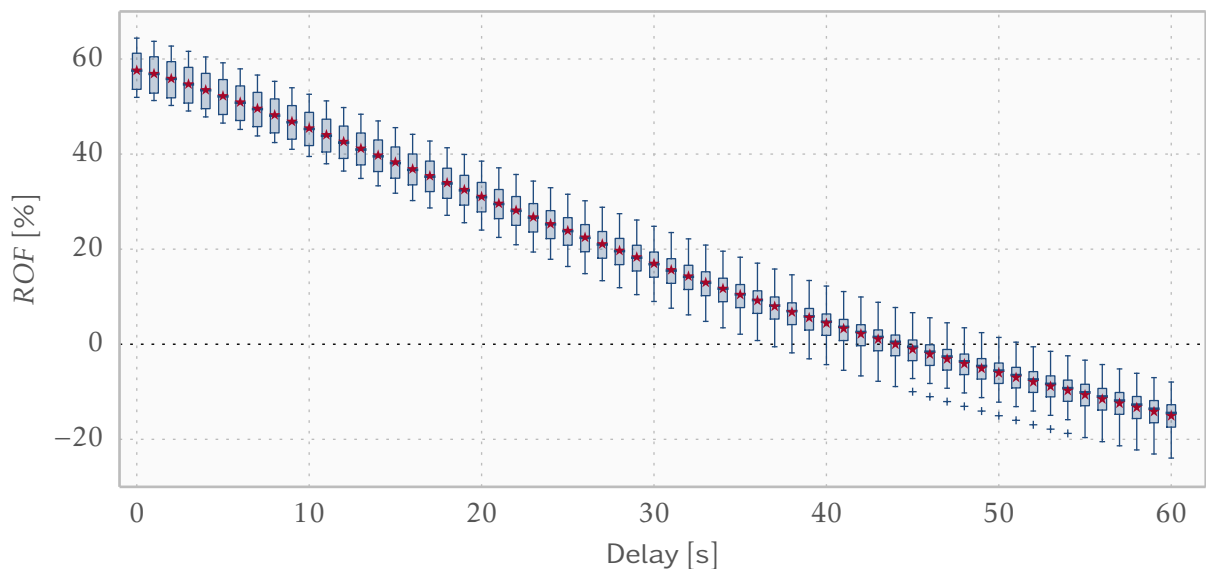


Fig. 3. Reduction of fluctuation (ROF) for increasing delays, visualized as boxplots for 10 simulations each, with additional star markers for the averages.

A trend towards decentralized control in the power grid as well as an increased autonomy of the involved agents is discussed for several year now, cf. [23], [24], [25], [26], [27]. As a result, the traditional static grid evolves to an “open system of systems” [28], shaped by a large number of dynamically interacting self-interested parties. Self-x properties such as self-configuration, self-healing, self-protection and more [29], [30] are of particular interest here, in order to facilitate the maintainability of a complex system like this. We believe that self-organized approaches like COHDA, which allow preservation of autonomy as well as data privacy of the participating agents, will play an important role in the transition towards a sustainable, stable future energy grid.

ACKNOWLEDGMENT

The authors thank the German Academic Exchange Service for partly funding this work. The National Renewable Energy Laboratory kindly provided the environmental data used in the study. The simulations were performed at the HPC Cluster HERO, located at the University of Oldenburg (Germany) and funded by the DFG through its Major Research Instrumentation Programme (INST 184/108-1 FUGG) and the Ministry of Science and Culture (MWK) of the Lower Saxony State.

REFERENCES

- [1] C. Gellings and J. Chamberlin, *Demand-Side Management: Concepts And Methods*. The Fairmont Press, Inc., 1988.
- [2] B. Shen, G. Ghatikar, C. C. Ni, J. Dudley, P. Martin, and G. Wikler, “Addressing energy demand through demand response : International experiences and practices,” Lawrence Berkeley National Laboratory; ENERNOC, INC., Tech. Rep. LBNL-5580E, 2012. [Online]. Available: <http://drcc.lbl.gov/publications/addressing-energy-demand-through>
- [3] P. Milan, M. Wächter, and J. Peinke, “Turbulent character of wind energy,” *Phys. Rev. Lett.*, vol. 110, p. 138701, Mar 2013.
- [4] S. Parkinson, D. Wang, C. Crawford, and N. Djilali, “Wind integration in self-regulating electric load distributions,” *Energy Systems*, vol. 3, pp. 341–377, 2012.
- [5] —, “Comfort-constrained distributed heat pump management,” in *Energy Procedia*, vol. 12, 2011, pp. 849–855.
- [6] C. Hinrichs, S. Lehnhoff, and M. Sonnenschein, *COHDA: A Combinatorial Optimization Heuristic for Distributed Agents*, ser. Communications in Computer and Information Science. Springer, 2014, vol. 449, pp. 23–39.
- [7] A. Nieße, S. Beer, J. Bremer, C. Hinrichs, O. Lünsdorf, and M. Sonnenschein, “Conjoint Dynamic Aggregation and Scheduling Methods for Dynamic Virtual Power Plants,” in *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems*, ser. Annals of Computer Science and Information Systems, M. Ganzha, L. A. Maciaszek, and M. Paprzycki, Eds., vol. 2. IEEE, 2014, pp. 1505–1514.
- [8] C. Hinrichs and M. Sonnenschein, “The Effects of Variation on Solving a Combinatorial Optimization Problem in Collaborative Multi-Agent Systems,” in *Multiagent System Technologies*, ser. Lecture Notes in Computer Science, J. P. Müller, M. Weyrich, and A. L. Bazzan, Eds. Springer International Publishing, 2014, vol. 8732, pp. 170–187.
- [9] H. Madsen and J. Holst, “Estimation of continuous-time models for the heat dynamics of a building,” *Energy and Buildings*, vol. 22, pp. 67–79, 1995.
- [10] S. C. Parkinson, “Managing sustainable demand-side infrastructure for power system ancillary services,” M.A.Sc., University of Victoria, BC, Canada, 2011. [Online]. Available: <http://www.uvic.ca/research/centres/iesvic/publications/theses/>
- [11] D. S. Callaway, “Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy,” *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.
- [12] G. L. Johnson, *Wind Energy Systems*. Prentice Hall PTR, 1985.
- [13] F. Mattern, “Experience with a new distributed termination detection algorithm,” in *Distributed Algorithms*, ser. Lecture Notes in Computer Science, J. van Leeuwen, Ed. Springer Berlin Heidelberg, 1988, vol. 312, pp. 127–143.
- [14] S. Martens, “Grundlagen der Selbstorganisationsmechanismen Agent Discovery und Termination Detection im Kontext eines agentenbasierten Einsatzplanungsverfahrens,” B.Sc., University of Oldenburg, Germany, 2013. [Online]. Available: <http://oops.uni-oldenburg.de/1605/>
- [15] National Renewable Energy Laboratory (NREL), “NREL: MIDC/National Wind Technology Center M2 Tower (39.91 N, 105.235 W, 1855 m, GMT-7):” [Online]. Available: http://www.nrel.gov/midc/nwtc_m2/
- [16] E. Hirst and B. Kirby, “Separating and measuring the regulation and load-following ancillary services,” *Utilities Policy*, vol. 8, no. 2, pp. 75–81, 1999.
- [17] C. Hinrichs, U. Vogel, and M. Sonnenschein, “Modelling and evaluation

- of desynchronization strategies for controllable cooling devices,” in *6th Vienna International Conference on Mathematical Modelling (Mathmod 2009)*, I. Troch and F. Breitenecker, Eds., 2009.
- [18] C. Hinrichs, “Selbstorganisierte Einsatzplanung dezentraler Akteure im Smart Grid,” Ph.D., University of Oldenburg, Germany, 2014. [Online]. Available: <http://oops.uni-oldenburg.de/1960/>
- [19] L. Zheng, S. Parkinson, D. Wang, L. Cai, and C. Crawford, “Energy efficient communication networks design for demand response in smart grid,” in *2011 International Conference on Wireless Communications and Signal Processing (WCSP)*, 2011.
- [20] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *Journal of the ACM*, vol. 32, no. 2, pp. 374–382, 1985.
- [21] D. Geelen, G. van Kempen, F. van Hoogstraten, and A. Liotta, “A wireless mesh communication protocol for smart-metering,” in *Computing, Networking and Communications (ICNC), 2012 International Conference on*. IEEE, 2012, pp. 343–349.
- [22] L. Hofmann and M. Sonnenschein, Eds., *Smart Nord: Final Report*, Hannover, Germany, 2015. [Online]. Available: <http://smarnord.de/downloads/SmartNordFinalReport.pdf>
- [23] F. F. Wu, K. Moslehi, and A. Bose, “Power System Control Centers: Past, Present, and Future,” *Proceedings of the IEEE*, vol. 93, no. 11, pp. 1890–1908, 2005.
- [24] M. D. Ilić, “From Hierarchical to Open Access Electric Power Systems,” *Proceedings of the IEEE*, vol. 95, no. 5, pp. 1060–1084, 2007.
- [25] S. D. J. McArthur, E. M. Davidson, V. M. Catterson, A. L. Dimeas, N. D. Hatziargyriou, F. Ponci, and T. Funabashi, “Multi-Agent Systems for Power Engineering Applications—Part I: Concepts, Approaches, and Technical Challenges,” *IEEE Transactions on Power Systems*, vol. 22, no. 4, pp. 1743–1752, 2007.
- [26] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings, “Putting the “Smarts” into the Smart Grid: A Grand Challenge for Artificial Intelligence,” *Communications of the ACM*, vol. 55, no. 4, 2012.
- [27] H.-J. Appelrath, H. Kagermann, and C. Mayer, Eds., *Future Energy Grid*, ser. acatech STUDIE. Springer, 2012.
- [28] G. Anders, A. Schiendorfer, F. Siefert, J.-P. Steghöfer, and W. Reif, “Cooperative Resource Allocation in Open Systems of Systems,” *ACM Transactions on Autonomous and Adaptive Systems*, vol. 10, no. 2, 2015.
- [29] R. Sterritt, M. Parashar, H. Tianfield, and R. Unland, “A concise introduction to autonomic computing,” *Advanced Engineering Informatics*, vol. 19, no. 3, pp. 181–187, 2005.
- [30] H. Schmeck, C. Müller-Schloer, E. Çakar, M. Mnif, and U. Richter, “Adaptivity and self-organization in organic computing systems,” *ACM Transactions on Autonomous and Adaptive Systems*, vol. 5, no. 3, 2010.