A Distributed Control Architecture for a Loosely Coupled Virtual Power Plant

Florian Heimgaertner*, Uwe Ziegler†, Bernd Thomas‡, and Michael Menth*

* Chair of Communication Networks, University of Tuebingen, Tuebingen, Germany
Email: {florian.heimgaertner,menth}@uni-tuebingen.de
† Ingenieurbuero Ganssloser, Tuebingen, Germany, Email: uwe.ziegler@avat.de
‡ Reutlingen University, Faculty of Engineering, Reutlingen, Germany
Email: bernd.thomas@reutlingen-university.de

Abstract—The increasing share of renewable energy with volatile production results in higher variability of prices for electrical energy. Optimized operating schedules, e.g., for industrial units, can yield a considerable reduction of energy costs by shifting processes with high power consumption to times with low energy prices. We present a distributed control architecture for virtual power plants (VPPs) where VPP participants benefit from flexible adaptation of schedules to price forecasts while maintaining control of their operating schedule. An aggregator trades at the energy market on behalf of the participants and benefits from more detailed and reliable load profiles within the VPP.

I. INTRODUCTION

The large-scale integration of renewable energy sources induces new challenges for electrical power grids and the energy market. An increasing share of power generation is weather-dependent, leading to more volatile production. At the energy market, this results in large intraday variability of wholesale energy prices. Industrial units can benefit from this variability by shifting operation times of energy-intensive facilities to times when energy prices are low. However, most small and medium-sized enterprises (SMEs) do not have the opportunity to actively participate in the energy market.

The background of this work is a joint research project of the engineering company Ganssloser, Reutlingen University, and the University of Tuebingen. We present the first result, which is a distributed control architecture for integrating SMEs into a virtual power plant (VPP) without direct control of the SMEs' devices.

A VPP is a distributed power plant comprising the aggregated capacity of distributed energy resources (DERs) at multiple participants. The VPP acts like a single entity towards the energy market. According to the CIGRÉ Microgrid Evolution Roadmap [1], DERs comprise energy generators, controllable loads, and energy storages.

The contribution of this work is an architecture for a loosely coupled VPP with an aggregator acting at the energy market on behalf of the participating enterprises. In contrast to established VPP architectures, the presented approach does not hand over control of SME devices to an aggregator or other external entities. Distributed controllers at the participating enterprises perform local optimizations based on market and environmental information supplied by the aggregator. The aggregator trades at the energy market based on load or production profiles supplied by the participants.

In this work we focus on electrical consumers in industrial production units where control of schedules and information is most relevant. Distributed generators like combined heat and power plants (CHPs) do not suffer from such limitations and, e.g., photovoltaic (PV) generators cannot be scheduled.

This work is structured as follows. Section II discusses related work. Section III presents the design rationale for the distributed architecture. In Section IV we describe the proposed VPP architecture and introduce the components. Section V details the operation of the VPP. In Section VI we show the interactions between an enterprise and the aggregator for the current and subsequent planning periods. Section VII concludes the paper and points out future research directions.

II. RELATED WORK

In this section, we first provide an overview of related work in the area of energy modeling and load scheduling. Then we present selected VPP architectures and VPP research projects.

A. Energy Modeling and Load Scheduling

Pfenninger et. al. [2] provide an overview of energy systems models. They group the models according to the four paradigms (energy systems optimization, energy systems simulation, power systems and market models, and qualitative and mixed-methods scenarios) and identify specific challenges.

Seow and Rahimifard [3] propose a framework for modeling energy consumption in manufacturing systems from a product viewpoint. That means, processes are modeled to determine the so-called embodied product energy (EPE). The EPE comprises theoretical energy and auxiliary energy that are used in the production process and indirect energy needed to maintain the environment. The model can be used for simulations and as a decision tool for implementing energy efficient manufacturing.

Howells et. al. present OSeMOSYS [4], the open source energy modeling system. OseMOSYS is a tool for large-scale and long-run energy planning.

Shrouf et. al. [5] propose a method for schedule optimization of a single machine in manufacturing facilities to minimize total energy costs considering variable energy prices. The authors use a metaheuristic approach based on genetic algorithms to schedule a fixed sequence of jobs and compare the resulting energy costs to the costs caused by scheduling each job as soon as possible.

Bruzzone et. al. [6] present a mathematical model for energy-aware scheduling (EAS) in manufacturing processes that can be used for peak shaving. EAS optimizes an energy-unaware reference schedule by means of mixed integer programming.

Niedermeier et. al. [7] propose a demand side management approach to maximize the use of renewable energy with flexible loads. The authors argue that power planning should be separated from load scheduling. They use a power planner based on quadratic programming to compute optimized load profiles.

Niknam et. al. [8] describe a stochastic model for optimal operation of a microgrid. A multi-objective teaching-learning optimization is used to reduce the solution space.

Paterakis et. al. [9] deal with the prize-optimized scheduling of a smart home using mixed integer linear programming. They use a reactive scheduling approach with dynamic prices and a rolling horizon.

Zhao et. al. [10] investigate the control and bidding problem of VPPs with renewable generation. They propose a strategy to minimize costs in a VPP with inelastic demands at the day-ahead and balancing markets by means of local optimization.

B. VPP Architectures and Projects

Luo et. al. [11] propose a distributed control framework for the coordination of multiple microgrids. The framework is a layered multi-agent system (MAS) implemented using the Java Agent Development Framework (JADE). The authors provide a simulative evaluation to demonstrate the effectiveness of the method.

The PowerMatcher framework [12] aims at providing a communication and trading infrastructure for VPPs at distribution system operator scale. PowerMatcher is implemented as a hierarchical MAS. The agents are connected by the MQTT message bus middleware. A demonstration of PowerMatcher was performed in multiple field tests in the Netherlands [13], [14].

Power Mix Manager (PMM) [15] is an energy management system that is currently under development and will be used in the research microgrids Eco Campus, and Renewable Energy Integration Demonstrator Singapore (REIDS) [16].

The Smart Polygeneration Microgrid (SPM) [17] is a microgrid and VPP testbed at the Savona campus of the University of Genoa in Italy. DERs include CHPs, PV systems, wind turbines, battery and thermal storage as well as absorption chillers and charging stations for electrical vehicles (EV). The DERs are managed using a control architecture based on IEC 61850 [18].

The Intelligent Microgrid Demonstrator [19] is a research microgrid at the Center for Electromechanics at the University of Texas, Austin. The microgrid is used for research on power systems for ships and comprises generators powered by diesel engines and gas turbines with an overall power rating of 2 MW. The control infrastructure for the microgrid is based on NI LabView.

The Nice Grid Demonstrator [20], [21] is a project integrating PV systems, distributed battery storage units and controllable loads in residential buildings in the region of Nice, France. Its goal is the maximization of renewable energy use by means of demand-response and optimized storage utilization. The control infrastructure is based on GE Grid Solutions Network Energy Manager and uses XMPP [22] as communication protocol.

Regenerative Modellregion Harz (RegModHarz) [23] is a VPP project in the Harz region in Germany. RegModHarz combines PV systems, wind turbines, bio gas fueled CHPs and pumped hydroelectric storage to a combined virtual power plant. The control infrastructure is based on web services using IEC 61850 data models [24].

Kombikraftwerk 2 [25] is a VPP project for ancillary services like frequency and voltage control. It uses only renewable energy sources like wind turbines, PV systems, and bio gas engines. Communication and control infrastructure are based on OPC-XML-DA, Modbus TCP [26], and IEC 61850 [18].

The Virtual Power Plant Neckar-Alb [27] is a demonstrator at the Reutligen University Campus. The DER portfolio comprises CHPs, PV systems, a heat pump, battery storage, thermal storages, an adsorption chiller, and EV charging stations [28]. DER controllers communicate to a central management and control system using Modbus TCP [26] or via HTTP using a JSON [29] based data format. Communication is secured using OpenVPN [30].

An overview on distributed control and management techniques in microgrids is provided by Yazdanian and Mehrizi-Sani [31].

III. DESIGN RATIONALE

Electrical energy is traded at different markets. Figure 1 shows the time scales of the energy markets. Long-term delivery contracts can be concluded at the futures market while short-term trading is done at the spot market. The spot market is divided into the day-ahead market, and the intraday market. At the day-ahead market, energy is traded with a granularity of an hour in an auction for the following day. At the intraday market, energy is traded in granularities smaller than one hour, e.g., 15 minutes for the German intraday market at EPEX SPOT, the European Power Exchange [32]. Energy can be traded continuously at the intraday market beginning with the intraday auction three hours after the the day-ahead auction until 30 minutes before fulfillment of the contract. Additionally, there is a separate market for operating reserve.
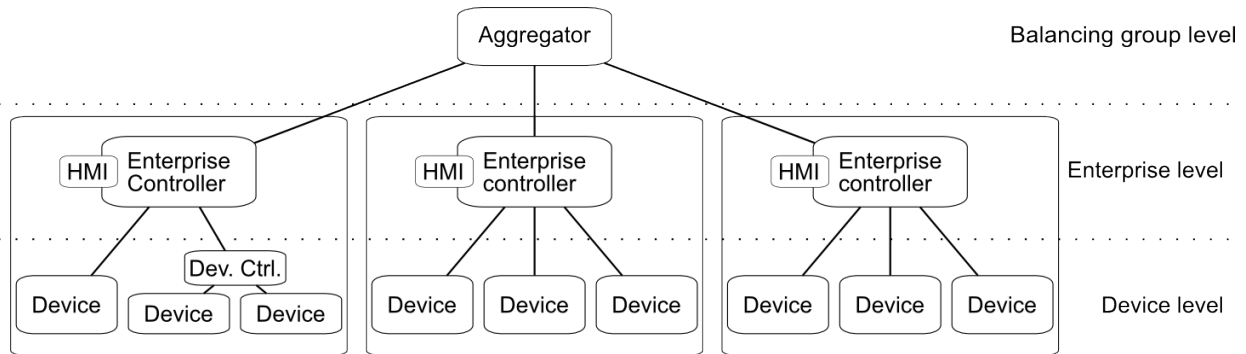
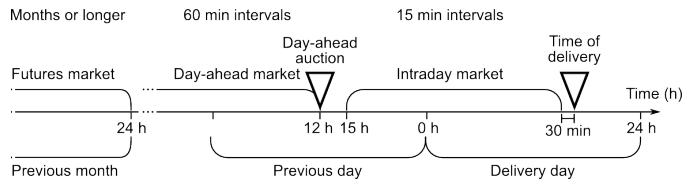Fig. 2: Hierarchy levels of the VPP.

Fig. 1: Time scales of the energy markets.

The formation of prices at the energy markets is a result of energy demand and supply. Weather-dependent generation of renewable energy results in large variability of energy prices at the spot markets. To benefit from variability of prices, industrial units can schedule their production in order to shift energy-intensive processes to times of low prices.

However, to benefit from variability in the day-ahead prices the operation of each device needs to be scheduled in advance. Therefore, reliable price forecasts are required. Additionally, in industrial units each device is part of larger processes. Therefore, dependencies and constraints need to be taken into account. This requires a detailed model of an industrial unit including devices and processes.

While major enterprises could trade energy directly at the energy exchanges, most SMEs lack the capabilities and knowledge to participate in the market or do not meet the minimum requirements. An aggregator can enable SMEs to indirectly participate in the energy market by providing price forecasts and acting at the energy exchange on behalf of the SMEs.

In the electrical power grid, energy demand and energy supply must match. To achieve this goal, the power grid is divided into balancing group. Within each balancing group the energy sold or purchased at the markets must match the total energy consumed or generated. In case of a mismatch, additional costs for balancing energy incur. An aggregator responsible for a balancing group can benefit from load profiles for each member of the group by improving the precision of the match. This is especially important if the enterprises in the balancing group adapt schedules to variable energy prices, as standard load profiles become less useful.

The objective of the proposed architecture is to leverage flexibilities and price differences to reduce energy costs in industrial units without raising the total costs with the help of an aggregator. However, for privacy reasons, control of devices and schedules should not be handed to the aggregator, and internal information should not be disclosed.

IV. ARCHITECTURE

The architecture of the VPP is organized in three levels: the balancing group level, the enterprise level, and the device level.

Figure 2 depicts the logical hierarchy of the VPP architecture. At the balancing group level, the aggregator interfaces to the energy markets and forecasting services. At the enterprise level, an enterprise controller schedules local resources at each VPP participant. At the device level, schedules are executed directly by networked devices or by networked device controllers interfacing to local devices.

In the following, we describe the components of the VPP architecture and their interaction.

A. Aggregator

The aggregator is the central component of the proposed VPP architecture. We assume that the aggregator is responsible for a balancing group in the electrical energy grid. It needs to match the total energy demand or production of the balancing group by purchasing or selling the corresponding amounts at the energy market. To determine the proper energy volume to buy or sell, the aggregator can use standard load profiles or forecasts based on historical data. However, more precise predictions can be achieved by considering local load profiles supplied by the balancing group members.

The aggregator computes price forecasts for the energy markets or acquires price forecasts from external service providers. These price forecasts are distributed to the enterprise level.

B. Enterprise Controller

At each VPP participant, a controller is connected to the aggregator and local devices. Using price forecasts received from the aggregator and a model supplied by the business operator, the controller computes cost-optimized schedules. A load profile derived from the schedule and external factors like the weather is sent to the aggregator.

The Human Machine Interface (HMI) enables the operator of the industrial plant to interact with the controllers. The

operator can enter parameters and specifications at the HMI or confirm proposed schedules.

C. Devices

The device level of the VPP architecture consists of DERs. Examples for generators located at participating enterprises are PV panels or local CHP units. In industrial units, examples for controllable loads are networked manufacturing devices, or devices that can be controlled by external networked device controllers.

D. Communication Infrastructure

Communication between the enterprise controller and the devices or device controllers uses protocols like Modbus-TCP [26], [33], IEC 60870-5-104 [34], or device-specific control protocols.

Communication between the enterprise controller and the aggregator uses TCP communication. The aggregator assumes the role of the server, waiting for connections initiated by the enterprise controllers. To ensure data integrity, source authentication, and confidentiality, the communication between the enterprise controllers are secured using Transport Layer Security (TLS) [35].

Communication of the aggregator to the energy market and information providers uses the interfaces defined by the communication partners. Mostly HTTP-based communication with web service APIs are expected to be used.

V. Operation at Enterprise Level

In the following, we describe the operation of the proposed VPP. As a first step, a model needs to be created for the industrial unit. With the help of this model, production processes are scheduled. In the planning process, additional flexibilities are identified. Based on schedules and external factors, profiles of energy demand or feed-in are computed. Finally, we show by examples how the components of the VPP interact to achieve cost reduction.

A. Enterprise Modeling

Industrial units can benefit from energy price variability, e.g., by shifting operation times of energy-intensive devices to times when energy prices are low. However, in an industrial setting, power consumers are often embedded in complex manufacturing processes and cannot be scheduled independently. Each schedules needs to be applicable to an industrial unit. Besides meeting production targets, schedules must respect dependencies between components and adhere to technical limitations of devices and processes. Additionally, legal restrictions like noise regulations or organizational constraints like workforce availability must be taken into account. To obtain valid schedules, an industrial unit needs to be modeled.

The model is a formal description of an industrial unit and its manufacturing processes. It is represented as a graph model with storage units and devices being the vertices. The edges denote the directions of material and energy flows. The graph model is annotated with time windows, processing rates, and

additional dependencies like, e.g., prerequisite devices, run time limits, warm-up times, cool-down periods, or global rate limits.

The model needs to be specified before integrating an enterprise into the VPP. It needs to be adapted when processes change and parameters like times of operation have been identified for each planning period.
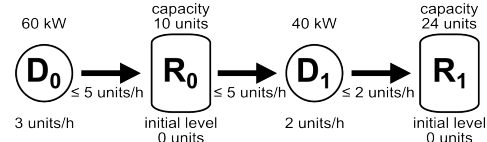

Fig. 3: Graph notation of a minimal example enterprise model.

Figure 3 shows a minimal example of an enterprise model. D_0 is a device capable of producing three units of a product per hour at an electrical load of $60\,\mathrm{kW}$. The output of D_0 is fed into the storage unit R_0. R_0 has a capacity of 10 units and a maximum input and output rate of 5 units per hour. The device D_1 obtains its starting products from R_0. Per hour, D_1 can transform two units of its starting product into two units of final product at an electrical load of $40\,\mathrm{kW}$. The final product is fed into the storage unit R_1. R_1 has a capacity of 24 units and a maximum input rate of 2 units per hour. Both storage units R_0 and R_1 are empty at the beginning of the planning period, i.e., midnight. The minimum run time of all devices is 2 hours. For the sake of simplicity no additional restrictions are imposed. In the following, we use this model to illustrate the operation of the VPP.

B. Planning

Given the dependencies and constraints from the enterprise model and the production targets for the current planning period, schedules can be computed. The planning process starts from the end of a process chain and recursively traverses the graph to find rates and run times for each device to meet the production targets.

Storage units increase temporal flexibilities by decoupling the run times of devices within a process chain. To determine valid rates and run times, a scheduling corridor is computed between the upper and lower bound for the cumulative charge of a storage unit. This method was developed and successfully applied to an optimized operation of CHP units for onsite electricity utilization by Widman et al. [36]. Meanwhile this method has been advanced for heat pump systems forced to run on PV electricity, and in addition corridors for more than one storage unit can be superimposed [37].

Using the recursive planning approach and selecting random values inside the valid ranges, multiple schedules are computed, and each schedule is evaluated based on the resulting energy costs.

We illustrate the planning process using the example model presented in Section V-A. The storage unit R_1 represents the end of the process graph. Therefore, the planning process starts at R_1. We specify the target of production as a storage fully
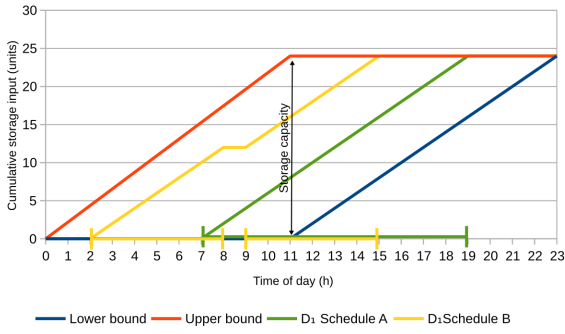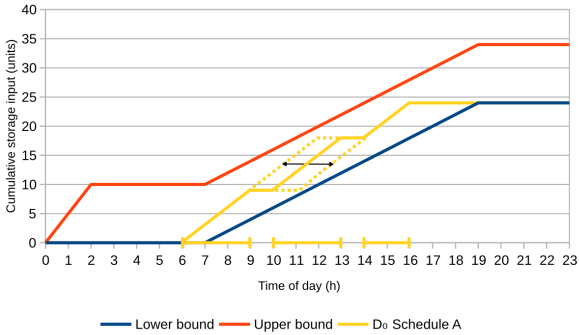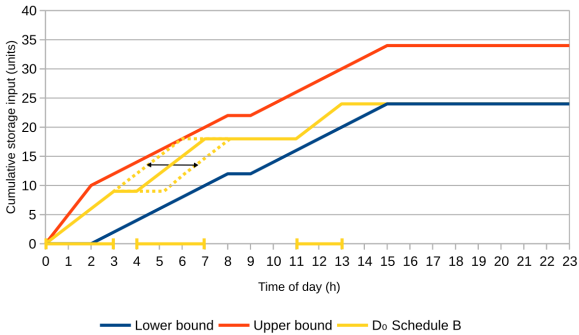
Fig. 4: Scheduling corridor for storage unit R_1.

loaded by exactly 24 units at the end of the planning period. From this target and the maximum input rate we derive the scheduling corridor for R_1 as shown in Figure 4. The area between the upper bound and the lower bound designates the flexibility available for scheduling the operation of the device D_1. The distance between the upper and lower bound is limited by the storage capacity. No discharge of R_1 is scheduled within the planning period, hence the upper bound is limited by the production target of 24 units. The slope of both the upper and the lower bound are given by the maximum input rate of the storage. Figure 4 contains two exemplary schedules for D_1. Schedule A consists of a continuous operation of D_1 for 12 hours, Schedule B uses two separate runs of D_1 of 6 hours each.

(a) Scheduling corridor for R_0 based on Schedule A for D_1

(b) Scheduling corridor for R_0 based on Schedule B for D_1
Fig. 5: Scheduling corridors for storage unit R_0.

These two schedules yield different scheduling corridors for the storage unit R_0. Figure 5(a) shows the scheduling corridor resulting from Schedule A and Figure 5(b) shows the scheduling corridor resulting from Schedule Bi. The lower bound results from the times of operation of D_1 receiving its input from R_0. The upper bound is given by the maximum input rate of R_0 and the storage capacity limiting the offset from the lower bound. Figures 5(a) and 5(b) illustrate schedules for device D_0 within the respective scheduling corridors. Both schedules consist of three runs of D_0.

C. Load Profiles

A load profile is the predicted electrical load over time. A schedule is associated with specific electrical energy consumption or generation. In the following we summarize both energy consumption or generation as load. Thus, schedules generating electrical energy are associated with negative loads.

The electrical load profile associated with a schedule consists of three parts. The first part is constituted by the electrical loads immediately caused by the scheduled devices. This part can be computed from the operating power and the scheduled times of operation. The second part is caused by additional electrical loads that are not part of the model and but depend on the operation times of scheduled devices. Among other things, those loads comprise air conditioning, ventilation, or lighting which are also dependent on external factors like time of day and weather. The third part is a base load that is independent of the schedule, but may depend on external factors like day of week and time of day.

To predict the required electrical energy of an enterprise for the entire planning period, a schedule is fed into a forecasting mechanism to extrapolate the expected total load of scheduled and unscheduled devices based on schedules, historical data, and weather forecasts.

For our illustration we ignore external factors and assume that electrical load is immediately caused by scheduled devices.

D. Optimization

Using the load profile and the price forecast received from the aggregator, energy costs for a schedule can be computed. Additional costs associated with a schedule, e.g., personnel costs, can be computed from annotations in the model. Hence, each schedule can be mapped to a price.

A naive optimization strategy that can be used is the Monte-Carlo-optimization. It generates a large number of schedules and computes the prices for each schedule. The schedule providing the best prices is finally selected. More advanced optimization strategies may include heuristic approaches where given schedules are modified to yield better prices.

Figures 6(a) and 6(b) show energy costs resulting from Schedules A and B with a fictitious price forecast for the day-ahead market ranging between 23 and 42 € per MWh. The bars show the total load of D_0 and D_1 in kW, the yellow line illustrates the price forecast in € per MWh, and the red line shows the cumulative costs resulting from the load profile
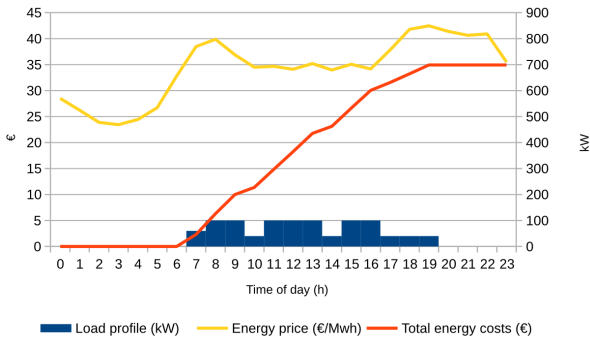
(a) Energy costs for Schedule A
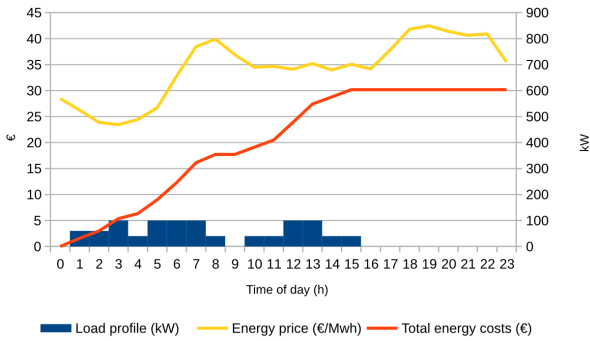
(b) Energy costs for Schedule B

Fig. 6: Energy costs for Schedules A and B with a given price forecast

and the price forecast. In our example, Schedule B yields a reduction in energy cost of approximately 14% compared to Schedule A.

E. Flexibilities

Flexibilities can exist both in the model of an enterprise and in a schedule. Flexibilities in the model exist as temporal flexibilities and rate flexibility. Temporal flexibility denotes that the operation of a device or an entire process can be shifted in time. Temporal flexibility is provided by storage units. Rate flexibility denotes that a device can run at different rates or that a device can be replaced by an alternative device with a different fixed rate.

The flexibilities of the model are exploited during the planning process in order to obtain price-optimized schedules. However, some flexibility remains usable within a schedule. In some cases, a operation of one or multiple devices connected to storage units can be temporally shifted within the limits of storage capacity. If such a temporal shift is possible without affecting other parts of a schedule, we call this kind of flexibility a *schedule elasticity* (SE). SEs are identified during the planning process using the mechanisms described in Section V-B. As an example, Figures 5(a) and 5(b) show SEs for the second run of D_0 which can be temporally shifted by one hour in each direction without violating the scheduling corridor or affecting the other scheduled operations of the same device.

Independent flexibilities (IFs) are a second kind of flexibility that can be exploited after the planning process is completed. IFs are provided by devices that are not scheduled for the current planning period but may run on request without affecting the schedule of the current planning period. Likewise, IFs can be given by devices that are scheduled to run but their operation may be canceled or shifted without affecting the schedule of the current planning period. An IF is only valid within the scope of a schedule and its planning period. The use of an IF in a planning period can prevent its use in subsequent planning periods, e.g., a process that may run once per week when the planning period is 24 hours.

F. Overview of Local Planning and Optimization

Figure 7 shows the process of local planning and optimization. The model of the enterprise is parametrized for the next planning period and fed to the scheduler. The scheduler generates a set of schedules and computes the associated SEs and IFs.
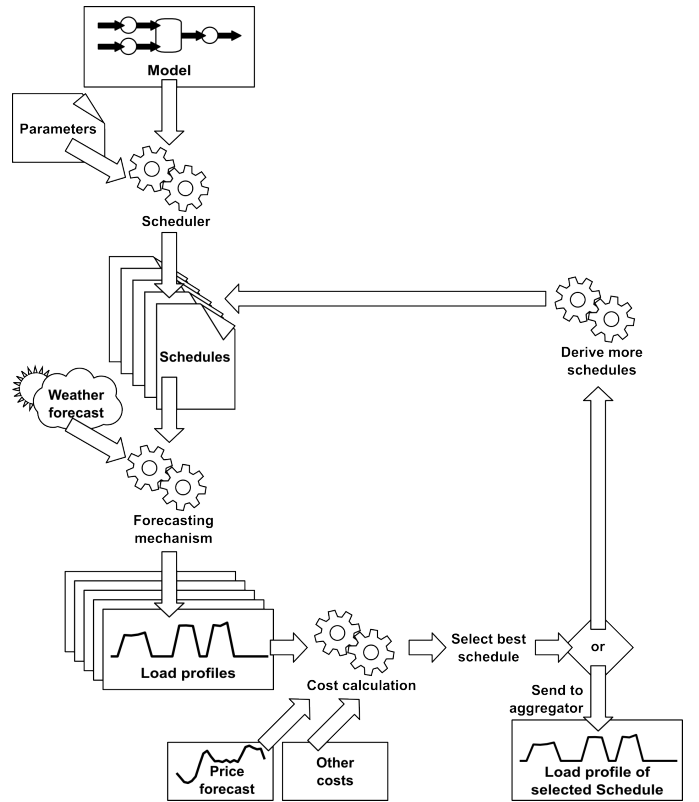

Fig. 7: Local planning and optimization.

Load profiles are derived for the generated schedules using weather forecasts and other relevant information. Based on the load profiles and the energy price forecast the costs of the schedules are computed. The schedule yielding the best costs is either selected for the next planning period and sent to the aggregator or further optimization is performed by modifying the schedule.
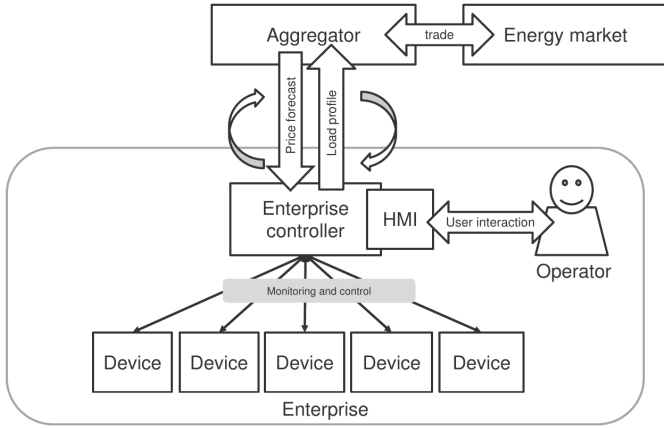
Fig. 8: System overview.

Fig. 9: Day-ahead cycle.

VI. INTERACTIONS OF ENTERPRISE AND AGGREGATOR

Figure 8 provides an overview of the entire system. The aggregator trades at the energy market, supplies price forecasts to the enterprise controllers, and receives load profiles from the enterprise controllers. The enterprise controller schedules processes at the enterprise level, provides a user interface for the operator, and controls devices.

In the following, we describe the day-ahead and intraday interactions between the aggregator and the enterprise controllers. We assume that the planning period is one day.

A. Day-Ahead Cycle

The day-ahead cycle comprises the communication and interactions between the aggregator and the VPP participants to negotiate energy demand and supply for the subsequent planning period, i.e., the next day.

Figure 9 illustrates the day-ahead cycle. The aggregator distributes price forecasts for the day-ahead energy market to the VPP participants (1). The price forecasts reflect the predicted day-ahead prices at the spot market, but they can also be modified by the aggregator to account for already existing energy contracts from the futures market. The VPP participants compute schedules and perform price optimization based on their models and production targets using the price forecast received from the aggregator (2). After the best schedule is selected, the VPP participants send load profiles to the aggregator (3). SEs and IFs can be sent to the aggregator in addition. SE and IF data are annotated with additional organizational and economic information like lead times, and costs. As this discloses additional information about an enterprise to the aggregator, exposing SEs and IFs needs to be confirmed explicitly by the operator.

After receiving load profiles from all VPP participants, the aggregator checks global constraints (4) like cumulative power limits for multiple enterprises. An example for a global constraint is a cumulative power limit for a group of VPP participants imposed by the grid operator. If such a constraint is violated, the aggregator may send modified price forecasts to the affected groups (5). Those modified forecasts could include
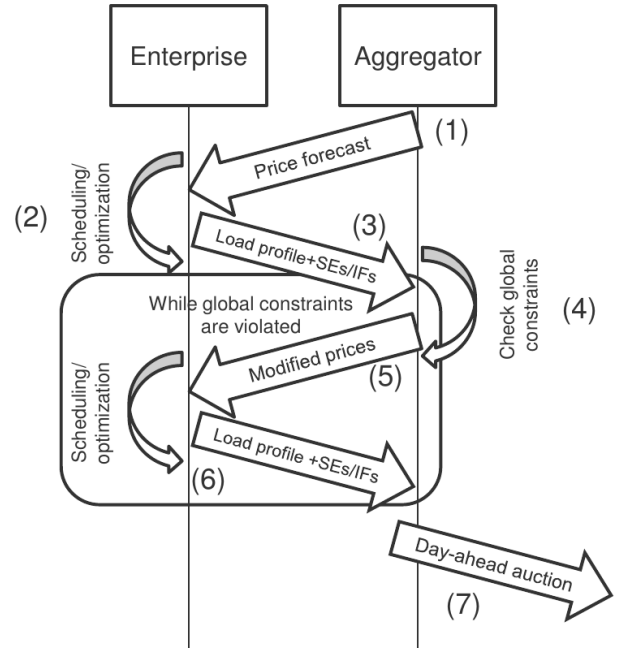
price penalties for times when the constraints are violated or price reductions for other times as incentive to shift loads. The affected VPP participants perform planning and optimization based on the new prices and send new load profiles to the aggregator (6). Steps (4)–(6) are repeated until the global constraints are no longer violated.

The end of this cycle is the deadline for the day-ahead energy auction, which is noon at EPEX SPOT. At the day-ahead-auction, the aggregator orders energy volumes according to aggregated load profiles for the VPP.

B. Intraday-Cycle

The intraday cycle comprises the communication and interactions between the aggregator and the VPP participants to compensate for deviations from the day-ahead forecast.

Figure 10 shows the intraday cycle. If conditions change and deviations from the day-ahead load profile are to be expected, the VPP participants send new load profiles to the aggregator (1). The aggregator assesses the announced SEs and IFs and compares the conditions of SE/IF use to the prices at the intraday market (2). Depending on the outcome of the make-or-buy decision, the aggregator either requests VPP participants to activate SE or IF (3a) or the aggregator buys or sells the remaining energy at the intraday market.

If no deviations are reported, the aggregator can bring SE and IF to the intraday market to generate additional revenues (not depicted).

VII. CONCLUSION

We have presented a distributed control architecture for VPPs. An aggregator acts at the energy markets for the VPP participants based on load profiles supplied by the participants.
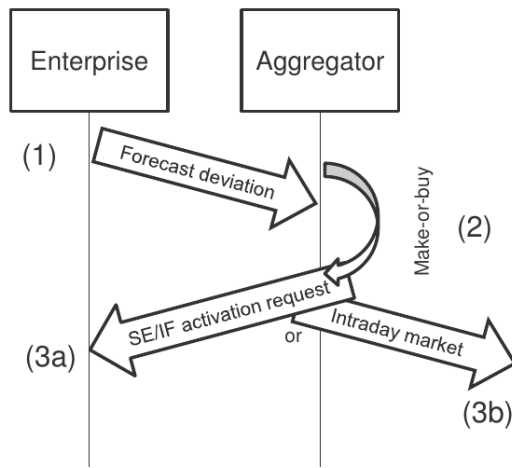
Fig. 10: Intraday cycle.

The participants benefit from lower energy costs by optimizing operation schedules based on price forecasts supplied by the aggregator. The agregator benefits from improved match within his balancing group. The aggregator does not require direct control of devices or internal information about an enterprise. Business operators can influence the scheduling process and may interactively change parameters or reject a proposed schedule.

Besides the control architecture, identification and quantification of flexibilities, evaluation of appropriate forecasting mechanisms, and optimization methods including metaheuristics are in focus of further research. Future extensions of the architecture include the integration of local planning and optimization into enterprise resource planning (ERP) software. In addition to spot market price optimization, providing ancillary services like operating reserve will be investigated as a use case.

REFERENCES

[1] C. Marnay, S. Chatzivasileiadis, C. Abbey, R. Iravani, G. Joos, P. Lombardi, P. Mancarella, and J. von Appen, "Microgrid Evolution Roadmap," in *International Symposium on Smart Electric Distribution Systems and Technologies (EDST)*, 2015.

[2] S. Pfenninger, A. Hawkes, and J. Keirstead, "Energy Systems Modeling for Twenty-First Century Energy Challenges," *Renewable and Sustainable Energy Reviews*, vol. 33, pp. 74–86, 2014.

[3] Y. Seow and S. Rahimifard, "A Framework for Modelling Energy Consumption Within Manufacturing Systems," *CIRP Journal of Manufacturing Science and Technology*, vol. 4, no. 3, pp. 258–264, 2011.

[4] M. Howells, H. Rogner, N. Strachan, C. Heaps, H. Huntington, S. Kypreos, A. Hughes, S. Silveira, J. DeCarolis, M. Bazillian, and A. Roehrl, "OSeMOSYS: The Open Source Energy Modeling System: An Introduction to its Ethos, Structure and Development," *Energy Policy*, vol. 39, no. 10, pp. 5850–5870, 2011.

[5] F. Shrouf, J. Ordieres-Meré, A. García-Sánchez, and M. Ortega-Mier, "Optimizing the Production Scheduling of a Single Machine to Minimize Total Energy Consumption Costs," *Journal of Cleaner Production*, vol. 67, pp. 197–207, 2014.

[6] A. Bruzzone, D. Anghinolfi, M. Paolucci, and F. Tonelli, "Energy-Aware Scheduling for Improving Manufacturing Process Sustainability: A Mathematical Model for Flexible Flow Shops," *CIRP Annals-Manufacturing Technology*, vol. 61, no. 1, pp. 459–462, 2012.

[7] F. Niedermeier, F. Kazhamiaka, and H. de Meer, "Energy Supply Aware Power Planning for Flexible Loads," in *Proceedings of the 5th International Workshop on Energy Efficient Data Centres (E2DC)*, 2016.

[8] T. Niknam, R. Azizipanah-Abarghooee, and M. R. Narimani, "An Efficient Scenario-Based Stochastic Programming Framework for Multi-Objective Optimal Micro-Grid Operation," *Applied Energy*, vol. 99, pp. 455–470, Nov. 2012.

[9] N. G. Paterakis, I. N. Pappi, J. P. S. Catalao, and O. Erdinc, "Optimal Operation of Smart Houses by a Real-Time Rolling Horizon Algorithm," in *IEEE Power and Energy Society General Meeting (PESGM)*, Jul. 2016.

[10] Q. Zhao, Y. Shen, and M. Li, "Control and Bidding Strategy for Virtual Power Plants With Renewable Generation and Inelastic Demand in Electricity Markets," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 2, pp. 562–575, 2016.

[11] F. Luo, Y. Chen, Z. Xu, G. Liang, Y. Zheng, and J. Qiu, "Multiagent-Based Cooperative Control Framework for Microgrids' Energy Imbalance," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 3, pp. 1046–1056, 2017.

[12] K. Kok, C. Warmer, and R. Kamphuis, "PowerMatcher: Multiagent Control in the Electricity Infrastructure," in *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Jul. 2005.

[13] J. Kok, B. Roossien, P. MacDougall, O. Pruissen, G. Venekamp, I. Kamphuis, J. Laarakkers, and C. Warmer, "Dynamic Pricing by Scalable Energy Management Systems – Field Experiences and Simulation Results Using PowerMatcher," in *IEEE Power and Energy Society General Meeting*, Jul. 2012.

[14] B. Roossien, "Field-Test Upscaling of Multi-Agent Coordination in the Electricity Grid," in *CIRED International Conference on Electricity Distribution*, 2009.

[15] T. Jiang, L. Costa, P. Tordjman, S. S. Venkata, N. Siebert, J. Kumar, H. B. Puttgen, A. Venkataraman, S. Dutta, Y. Li, and T. Tang, "A Microgrid Test Bed in Singapore: An Electrification Project for Affordable Access to Electricity With Optimal Asset Management," *IEEE Electrification Magazine*, vol. 5, no. 2, pp. 74–82, Jun. 2017.

[16] A. Venkataraman, S. Dutta, Y. Li, S. Kayal, L. M. Costa, T. Jiang, R. Plana, P. Tordjman, D. Tang, F. H. Choo, and C. F. Foo, "REIDS Microgrid: Development and Demonstration of a Power Mix Management System," in *Asian Conference on Energy, Power and Transportation Electrification (ACEPT)*, Oct. 2016.

[17] S. Bracco, F. Delfino, F. Pampararo, M. Robba, and M. Rossi, "The University of Genoa Smart Polygeneration Microgrid Test-Bed Facility: The Overall System, the Technologies and the Research Challenges," *Renewable and Sustainable Energy Reviews*, vol. 18, pp. 442–459, 2013.

[18] IEC, "IEC 61850: Power Utility Automation," 2003.

[19] R. E. Hebner, J. D. Herbst, and A. L. Gattozzi, "Intelligent Microgrid Demonstrator," in *Electric Machines Technology Symposium 2010 (EMTS)*, May 2010.

[20] A. Michiorri, R. Girard, G. Kariniotakis, C. Lebossé, and S. Albou, "A Local Energy Management System for Solar Integration and Improved Security of Supply: The Nice Grid Project," in *IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe)*, 2012.

[21] A. Michiorri, G. Kariniotakis, and F. Foucault, "An Aggregator for Distributed Energy Storage Units under Multiple Constraints in the Nice Grid Demonstrator," in *CIRED Workshop: Grid Operation and Congestion Management*, Jun. 2014.

[22] P. Saint-Andre, "Extensible Messaging and Presence Protocol (XMPP): Core," RFC 7159 https://www.ietf.org/rfc/rfc6120.txt, Mar. 2011.

[23] P. Hochloff, D. Filzek, G. Heusel, K. Lesch, A. Liebelt, L. Nicklaus, P. Ritter, K. Rohrig, C. Roehrig, F. Schloegl, C. Volkert, M. Wickert, and M. Winter, "Regenerative Modellregion Harz – RegModHarz," in *Kasseler Symposium Energie-Systemtechnik*, Oct. 2011.

[24] IEC, "IEC 61400-25-4: Communications for Monitoring and Control of Wind Power Plants - Mapping to Communication Profile," 2008.

[25] K. Knorr *et al.*, "Kombikraftwerk 2," Fraunhofer-Institut für Windenergie und Energiesystemtechnik (IWES), Final Report, Aug. 2014.

[26] Modbus IDA, "Modbus Messaging on TCP/IP Implementation Guide V1.0b," 2006.

[27] F. Heimgaertner, M. Hoefling, M. Menth, E. Schur, F. Truckenmueller, H. Hagenlocher, J. Zunke, A. Frey, D. Ebinger, S. Jaegers, L. Duerr, T. Roeger, K. Lindner, and C. Kahlert, "The Demonstration Project Virtual Power Plant Neckar-Alb," in *VDE Kongress 2016: Internet of Things*, Nov. 2016.

[28] F. Heimgaertner, E. Schur, F. Truckenmueller, and M. Menth, "A Virtual Power Plant Demonstration Platform for Multiple Optimization and Control Systems," in *International ETG Congress*, Nov. 2017.

[29] T. Bray, "The JavaScript Object Notation (JSON) Data Interchange Format," RFC 7159 https://www.ietf.org/rfc/rfc7159.txt, Mar. 2014.

[30] J. Yonan, "OpenVPN – An Open Source SSL VPN Solution," http://openvpn.net, 2008.

[31] M. Yazdanian and A. Mehrizi-Sani, "Distributed Control Techniques in Microgrids," *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2901–2909, 2014.

[32] EPEX SPOT SE, "EPEX SPOT," https://www.epexspot.com/en/, last visited Feb. 2018.

[33] Modbus IDA, "Modbus Application Protocol Specification v1.1b3," 2012.

[34] IEC, "IEC 60870-5-104: Network Access for IEC 60870-5-101 Using Standard Transport Profiles," 2000.

[35] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2," RFC 5246 https://www.ietf.org/rfc/rfc5246.txt, Aug. 2008.

[36] C. Widmann, D. Loedige, A. Toradmal, and B. Thomas, "Enabling CHP Units for Electricity Production on Demand by Smart Management of the Thermal Energy Storage," *Applied Thermal Engineering*, vol. 114, pp. 1487 – 1497, 2017.

[37] A. Toradmal, T. Kemmler, and B. Thomas, "Boosting the Share of Onsite PV-Electricity Utilization by Optimized Scheduling of a Heat Pump Using Buildings Thermal Inertia," *Applied Thermal Engineering*, vol. 137, pp. 248–258, 2018.