Development of an Agent-Based System for Decentralized Control of District Energy Systems

Jakob M. Fritz*, André Xhonneux* and Dirk Müller*†
*IEK-10: Energy Systems Engineering, Forschungszentrum Jülich, Germany, Email: {j.fritz, a.xhonneux}@fz-juelich.de
†EBC, E.ON Energy Research Center, RWTH Aachen University, Germany, Email: dmueller@eonerc.rwth-aachen.de

***Abstract*—This contribution introduces a new system for distributed model predictive control of energy systems. This system uses multiple agents where each agents optimizes a subsystem. A central instance coordinates the individual agents and takes care of feasibility of the combination of the single solutions. The advantages of this approach are increased maintainability and privacy for the individual agents, thus increasing applicability to real-world systems where often multiple parties are involved in a single energy system. Where adequate, the agents perform MPC to control their subsystems. The system and method can be chosen for each agent individually. In order to build this system a framework is developed as existing frameworks lack one or more required features for the overall system. A small example is presented together with first results.**

I. Introduction

According to the Paris agreement the temperature increase should be no more than 2°C. Preferably the increase would be no more than 1.5°C [1]. According to [2] the energy sector is responsible for over a third of the global CO_2 emissions in 2010. This is due to fossil fuels (coal, oil, and natural gas) burned in order to provide heat and electricity. Increasing the share of renewable energy sources in the overall electricity production decreases the CO_2 emissions in that sector. The disadvantage of many renewable energy sources as photovoltaic (PV) and wind power plants is that their energy production is volatile and can only be controlled more limited than conventional power plants. That volatility makes it harder to provide a permanently sufficient power supply as the production is increasingly dependent on wind or solar irradiation. To better cope with the more volatile supply, the demand should be able to follow the changes in production. This is getting more important with a further increasing share of renewable energy sources.

When optimizing the power demand (e.g. reducing costs) of multiple buildings – or consumers in general – different problems can arise. These issues can occur regardless of whether the context is domestic or if the system is planned to control commercial or industrial systems. In all three situations the overall system can exhibit an increasing number of units or machines and – more importantly – often include different participants. As these participants likely want to keep their privacy, creating a mathematical model of the entire system becomes very difficult or even impossible. Furthermore this mathematical model would be difficult and time consuming to solve due to its size and complexity. The last issue is that maintaining or adapting such a big model can be challenging.

One approach to deal with those issues is to split the system in such a way, that each participant is described by a subsystem. Each subsystem is then solved separately and the single results are combined to generate a solution for the total system. This is done in multiple iterations to converge to a feasible solution for the overall system. The digital representation of each participant will be referred to as agent in the rest of the paper.

The second chapter of this publication introduces model predictive control and distributed systems in general and also refers to current systems. Followed by that the new system is shown, explained how it works, and in which aspects it differs from other current implementations. Afterwards, the introduced system is illustrated with an example and some exemplary results are shown. In the end a short summary is given and future development named.

II. State of the art systems

There are multiple ways to control systems. This publication will focus on model predictive control (MPC) and distributed control systems as the developed system uses these strategies for controlling.

Model predictive control is a way to control systems where for a modelled system an objective is minimized, while being subject to given system constraints, that can explicitly be set [3]. The minimization is normally done for a given receding horizon for which values for the surrounding variables are forecasted. First publications about the concept of MPC appeared in the late 1970s [4]–[6]. These publications dealt either generally with the concept, or applied it to petrolchemical processes. In 1989 a survey about MPC was written by García et al. [7]. Ten years later, Morari and Lee gave an overview over past development of MPC and expectations where MPC will further improve [8].

The approach of splitting up larger systems into subsystems to easier model and solve them is not new. Already in 1978 Sandell et al. examined different ways to split up and control systems [9].

Splitting up MPC-systems into multiple parts has already been applied for many years as exemplarily shown in [10]–[12]. One field of application for MPC is the temporal shifting of energy consumption. This is known as demand side management (DSM), altough DSM can also be performed by other means than MPC. The main purpose of DSM is to better match the demand to a given supply, which is particularly

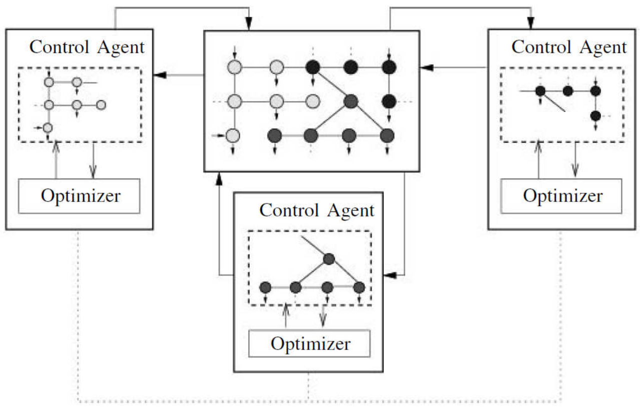

Fig. 1. Example of a distributed optimization of a system by multiple agents with indicated communication. This system does not use a central instance but the communication occurs between the agents directly. Taken from [10].

important for energy systems with a high share of volatile energy sources. This often means reducing peak demands and shifting the demand to times when less energy is consumed. An example of such a system (also in a distributed form) is discussed in [13].This will get more important with increasing share of renewable energy sources.

The behaviour of subsystems often depends not only on the current time step but also on the last time steps. Therefore it is impossible to use a Markov decision process, which would reduce computation time for optimizing [14].

When multiple agents control a larger system, by splitting it up and optimizing each subsystem individually, communication is important, because not necessarily all combinations of solutions for the subsystems are feasible solutions for the full system as well. Different approaches for this communication exist. Agents can communicate with either all or some other agents or with an instance that is solely coordinating. Many publications focus on the first approach, as it is often faster and needs less participants in the overall optimization [10], [15], [16]. Such a way of splitting up a larger system into subsystems, optimizing each part individually by a separate agent, and communicate with the other agents is described in [10] and shown in figure 1, taken from [10]. The second way of communicating only with a central instance is e.g. chosen in [12].

The communication can further be differentiated into iterative and non-iterative approaches [15]–[17]. When communicating with a central instance, the iterative approach is chosen. The advantage of communicating via a central instance is that the agent's privacy can be increased, as the agents do not communicate directly with each other. Depending on the setting this enhanced privacy can outweigh the increased need for communication. One example of a similar system is described in [12], where a central instance is used for communication and to find a convergence in an iterative manner. That system still shares consumption data the other agents with each agents. The system introduced in this contribution, follows a similar approach, but puts a stronger emphasis on privacy in not sharing any consumption data between different agents, but only derived price signals. Further the communication procedure differs here.

III. Structure of the new system

This system aims at representing energy systems from a few units or buildings up to the district level and to coordinate them in real-time in order to be applicable to energy systems for real use and not only as a case study. To better achieve that goal, the system consists of multiple components interacting with each other. This has several advantages. Firstly, it is easier to develop or modify single agents, as they are not coupled with each other and the communication is standardized and only with a central instance. Secondly, adding or removing individual agents to or from the system is easier. Thirdly, due to its modularity, the system can be parallelized easily and distributed across machines. This can enhance scalability as well as redundancy, if needed. Lastly, there is the ability of individual agents to provide a fall-back behaviour if other agents or the overall communication fail. In a strongly interconnected system this would be harder to accomplish. In this scenario, agents can change their behaviour into a safe operation mode. This means the overall efficiency of the system strongly decreases but the components can still operated safely. Therefore, the whole system is less error prone.

The communication follows the MQTT protocol to increase interoperability between different structures and agents developed by separate people or groups. It also means faster development of new agents, as many functions for the communication already exist and are easy to import.

Because the agents optimize their subsystem themselves, there is no need to share the internally used model or approach at all. The communication does not run between agents but between agent and central instance. This central instance serves as a market spectator and market clearer and is expected to be neutral towards individual agents, as it is not competing with them. Communicating only with this instance means that if two agents from the same industry but competing brands are located in the same network they do not share any internal knowledge or information about their processes with each other.

For the use in a general context further features have to be added to the system at later stages. These features include considering the local grid topology for matching supply and demand. Currently, the topology is not a limiting factor but since it could get in the future, it should be considered in general. Additionally, it may become important later to consider several flows of different kinds in the system. These flows can either be electricity of multiple voltages or they can include various energy carriers, such as hydrogen or natural gas.

A. System Components

The individual components of the system can be split up into three categories: agents performing optimizations of subsystems, the central instance having the overview over the agents and the whole system, and some services required for the agents or central instance to operate well.

1) Agents:

Many of the agents perform MPC with their included model of the subsystem they represent and optimize the total financial cost caused by power consumption. This optimization goal is chosen because the system shall be applied to real energy systems and most participants want to reduce their cost. This also allows to include costs for CO_2-emission which leads to situations where the emissions are mitigated effectively. Some other agents with systems with fewer degrees of freedom do not use MPC but simpler approaches. These other approaches within the agent do not change the systems overall approach or structure, since the systems treats the agents as black boxes.

In the system introduced here, most agents represent a building. In the building sector the biggest potential for DSM lies in HVAC systems with an electric heating unit, since most other demands are either lower or cannot be adapted easily without affecting comfort effects [18]–[20]. Most of the modelled buildings are equipped with a heat-pump and therefore allow for DSM and load-shifting to further reduce costs. This is generally possible due to the thermal mass of the buildings. Some buildings additionally include activation of building-components featuring low dynamics to further leverage this option. Examples of such models can be found in e.g. [21], [22].

Since not all buildings feature a heat pump or electric heater, some agents represent buildings with a non-electrical heat-supply so that the electricity is mostly needed for e.g. office equipment – or other plug loads – and therefore cannot be shifted.

Apart from the consumer agents – representing buildings or other consumers – , also producers are represented by agents. This can be simple agents representing the behaviour of a fixed photovoltaic installation (PV). Such agents do not have many degrees of freedom. The latter also holds for agents representing combined-heat-and-power (CHP) plants if those are operated in heat-driven mode. Both characteristics can be combined to model agents representing prosumers. In addition, an agent is needed that represents the connection to the local power grid, as the modelled system is not fully meshed into the local grid but only connected by a few points. This agent e.g. wants to sell electricity to others if there is a surplus of electricity in the local public power grid. Further agents – as for example for electrical storages – are planned but not yet in use.

As these agents are subject to different external constraints and differently complex, the internal structure can be chosen according to the given situation. The same holds true for the used programming language as long as it is able to communicate via MQTT. This offers the possibility to easily deploy more complex MPC-based agents as well as simpler agent with a rule-based strategy.

2) Central instance:

Combining the different objectives of the multiple agents into a feasible solution is the main goal of the central instance. This is the instance that manages the individual agents and works with the solutions returned by the agents. It communicates electricity prices for the receding horizon to all agents and aims to get a feasible combination of loads returned. To achieve this goal the power prices are adjusted at each iteration for the same horizon. Once the returned loads converge for all time-steps in the prediction horizon, the decided control strategy can be applied by the agents. The central instance does not know the control-parameters that are applied but only the electrical loads that are caused by the control-parameters. This central instance further does not communicate the overall load to the agents but only the power-price or in extreme cases the infeasibility of a time step if there is a lack of power-supply for that time step and an increasing power-price will not increase the production, as renewable and conventional power sources combined are not sufficiently available. If there is a large surplus of energy that would cause criticality of the system, the combination of solutions would also be infeasible and agents would have to ramp down production or ramp up consumption to stabilize the system.

The central instance's main task is to determine the power price of the receding horizon. The chosen mode depends on the overall goal that is pursued. The price can be the production price of the single producers. Alternatively, environmental costs, such as plant-specific CO_2-emissions, are added to the production price. This power price is determined similar to a classical market clearing but with the difference that a weighted average of the prices from all accepted offers is calculated. Using this approach, every producer still gets the price for which it offered electricity, but the overall market is more sensitive to smaller changes of demand and supply and therefore the convergence is easier to achieve.

To recognize the convergence, the central instance needs to keep track of the development of the power requests and offers of the single agents. In general, the instance stores only as much information about the agents and their power loads as needed. Storing the power load of the last iteration also helps in the unexpected event that one of the agents fails or shuts down unexpectedly, as then the last power load (forecast) can be used as a hint of the planned power consumption or production.

If new agents are added to the system, they can participate in the communication from the start of the next iteration and will be provided with all needed information. There is no need to restart the central instance or other parts of the system after adding agents. Depending on the behaviour of these agents, no or only a few parameters for convergence of the overall system need to be slightly adjusted.

When starting a new round of iterations, the initial power price is chosen identical to the final power price in the last round of iterations.

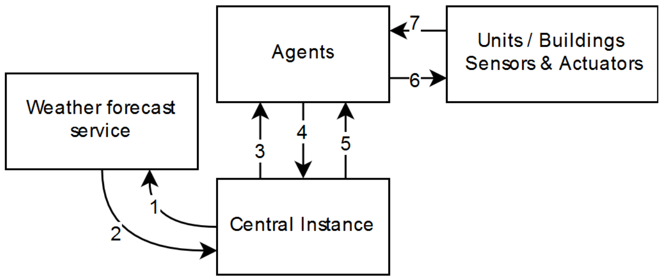

Fig. 2. Overall structure of the control system. First the central instance requests the current weather forecast (1, 2). This forecast, together with a power price is broadcasted to all agents (3). They reply a predicted load for this price curve (4). This iterates until convergence. Then the central instance tell the agents to implement the chosen control in their systems (5,6), where they also regularly get sensor data from (7).

Later deviations from the load forecast will result in penalties. This aims to provide incentives to the agents to offer good predictions and to stick to those predictions where ever possible.

3) Further services:

In order for the agents and the central instance to work properly some additional services are needed. These include a service providing a weather forecast and a MQTT broker. As the system was developed in Germany, the current implementation of weather service aims at German locations.

The service providing the weather forecast uses the open data provided by the "Deutsche Wetterdienst" (DWD), which is Germany's national meteorological service [23]. The service provides weather forecasts for multiple locations in Germany. The closest location is chosen automatically and the data from multiple forecasts are merged and provided for the agents to build predictions of the behaviour of their models on that forecast. This is done repeatedly to always provide a current forecast. Since many agents need a subset of the weather forecast data available, it was decided to provide this service to all agents to avoid multiple implementations and instances of the same functionality.

Another service is a MQTT broker responsible for the transmission of messages between the central instance and the agents. This is no own development but uses available implementations of an MQTT broker and –client. MQTT is one of the main communication protocols for sensors and similar networks, which results in the availability of many implementation in different programming languages. Using this protocol enables an easy communication between agents and central instance, while still allowing for free choice of programming language used inside the agents. Furthermore the communication is not limited to one host machine, which makes it possible to distribute agents on multiple hosts.

B. Interactions of components

The interaction of several components is shown in figure 2 where an overview of the system can be seen. Every few minutes the same procedure starts. First, the weather forecast is requested by the central instance (1) and returned by the weather forecast service (2). After the forecast is received, the central instance provides an initial power price for each time step of the upcoming horizon of e.g. 24 hours (3). With this information the agents each optimize their power consumption or production and return their scheduled power for every time step (4). The central instance collects all planned demands and supplies and checks if total demand and supply match for the receding horizon. If there is a mismatch, the power price for time steps with a mismatch is adapted accordingly and a new price profile for the full horizon is send again to all agents (3). The agents perform an optimization and return their planned profiles (4). This is repeated until the planned profiles converge (or other stopping criteria are met). When supply and demand match an information is sent to the agents that they can implement their controls for the planned load (5). They then sent this information to their physical units they control (6) and from which they continuously receive status information (7).

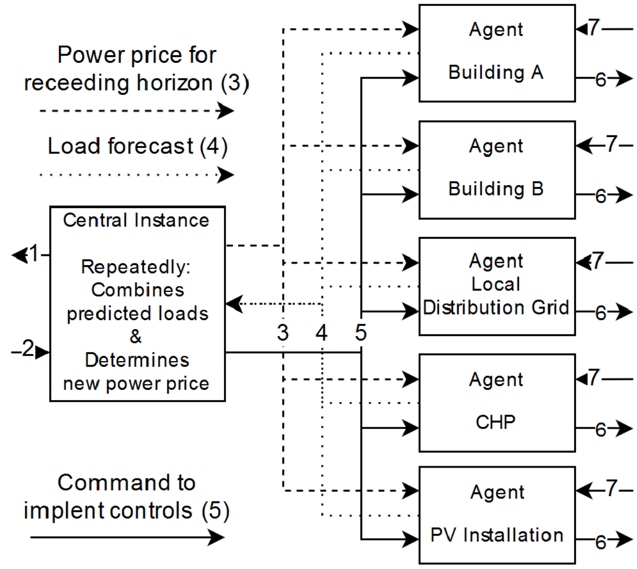

Fig. 3. Detailed structure of the communication between agents and central instance. The broadcasting of the power price by the central instance (3) and the reply with the predicted loads by the agents (4) occurs multiple times, until a convergence is reached. Then the command to implement the controls is send (5).

The main part of the system is the iterative communication between central instance and the agents. This is displayed in figure 3 in more detail. A detailed examination of the arising network traffic will take place at a later stage of research. This also includes the effects of errors occurring during the transmission of messages.

IV. Application and First Results

A. Planned application

This system is developed to work in a real-world implementation. This is a difference to other systems that are developed to study single mechanisms or behaviours theoretically but are not intended to stably control real units over a long period of time. This is intended for the presented system within a broader scope in the “Living Lab Energy Campus” research project combining multiple aspects of the energy sector. This includes the demonstration of multi-modal energy systems with a high share of renewable sources, the re-usage of waste-heat in a low-temperature heat network to heat nearby buildings, as well as the model predictive control of a single office building, which is equipped with many sensors for further research. The presented system shall combine those aspects into a harmonic energy system while still granting a high degree of freedom to the individual subsystems for them to perform well. An Overview of this research project can be found in [24].

B. Showcase and First Results

A simplified example of the introduced system will be discussed in this section. The example consists of five main parts: A central instance, two identical demand agents each representing a building, a supply agent representing connection to the external power grid, and a supply agent representing a photovoltaic installation.

The two identical agents each control a building with concrete core activation and heat supply by a heat pump, but without plug loads. The grid agent supplies large amounts of power for all time steps at a fixed price of 3 (arbitrary units). This agent acts as a dummy-version of an agent representing a connection to the external power grid. The last agent, representing a PV installation, provides power during the day with a power-peak at noon at a price of 1 to give an incentive to use renewable power sources. The PV installation is dimensioned in such a way that around noon it provides enough power to satisfy the maximal demand of the building agents. This means that around noon the power price will drop to 1 and no additional power provided by the grid-agent is needed.

An example can be found in figure 4. There demands and determined prices during the first and last iteration are shown for the whole optimization horizon of 24 h starting at midnight. The single time steps of the horizon are 15 minutes long. As can be seen, in the first iteration a constant power price is set to initialize the system (dark blue dashed line) and no load shifting of the demand agents is done (dark blue solid line). Starting from the 2nd iteration, the PV offers power at a lower price. This leads to a decreasing average price during the day (light blue dashed line). Because of the decreased price the agents representing buildings move their demand towards midday to benefit from the cheaper power (light blue dashed line). Since only one day is optimized, the morning hours do not differ in the scenario with a constant (cf. first iteration) and with a volatile power price (cf. last iteration). The demand in the evening hours after sunset in contrast changes a lot when supplying a volatile power price. The demand after sunset decreases, as the agents used the cheaper power during the day to store some energy in their concrete core. Around midnight this thermal storage seems to be depleted, and the agents need a similar amount of power as with a fixed price.

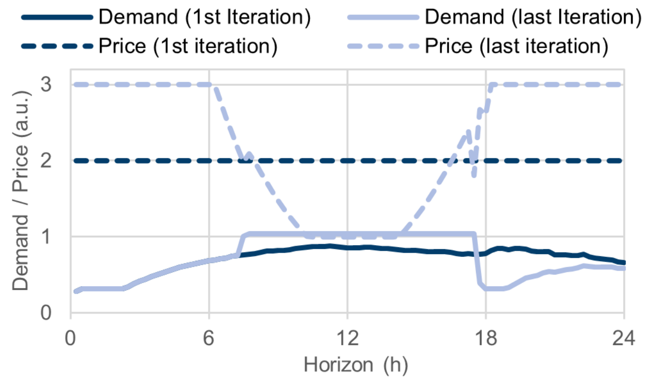

Fig. 4. The determined prices and loads of the agents in the distributed system controlled by a central instance for an exemplary day. The optimization if performed for a horizon of 24 h, starting at midnight. The time steps are each 15 min long, making a total of 96 time steps in the horizon.

This last iteration is already reached after a few iterations, where the exact number depends on multiple aspects and therefore varies. For a new set of iterations dealing with a horizon of 24 h starting later than midnight, the last power price would be used as initial power price. This example is supposed to indicate the cold-start behaviour of the system.

V. Conclusion and Outlook

This contribution introduces a system for district energy systems with a high share of renewable power that provides the possibility for multiple parties to jointly improve the energy-usage and reduce the need for energy storage – which always imply energy losses – without having to share their internally used models and sensor data. This will greatly help implementing such systems in real world situations where more than one party is involved. Agents can use different approaches internally, such as MPC or rule-based approaches. This can be chosen depending on the system to be represented. Furthermore the agents do not share any information with each other but only with a central instance that is expected to be neutral. This instance only receives planned load series for the receding horizon and combines the load of all agents iteratively to a consensus.

To enable that system, a new platform is needed that enables the development of real-time control over such a distributed system while granting individual agents a big degree of freedom to control their specific subsystem in accordance with the bigger context.

Further development will focus on development of more advanced agents (e.g. using machine learning approaches) and evaluate strategies to faster achieve equilibrium of supply and demand for the full horizon at a given time. Examinations of the behaviour of the whole system to unforeseen events, such as losing connection to agents or large deviations from the forecast will be performed in the future. The same holds for benchmarking a complex system controlled with this approach against the same system controlled with a non-distributed approach. Beyond this, the prerequisites that must be met by the agents are studied.

Acknowledgment

The authors would like to thank their colleagues for fruitful discussions.

Part of this work was funded by the German Federal Ministry of Economic Affairs and Energy within the research project "EnOB: LLEC: Living Lab Energy Campus" (03ET1551A).

References

[1] United Nations. (2015) Paris agreement. UNFCC. [Online]. Available: https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement

[2] J. G. J. Olivier, J. A. H. W. Peters, G. Janssens-Maenhout, and M. Muntean, "Trends in global co2 emissions. 2013 report," 2013. [Online]. Available: https://www.osti.gov/etdeweb/biblio/22176292

[3] P. J. Campo and M. Morari, "Robust model predictive control," *1987 American Control Conference*, pp. 1021–1026, 1987.

[4] J. Richalet, A. Rault, J. Testud, and J. Papon, "Model predictive heuristic control," *Automatica*, vol. 14, no. 5, pp. 413–428, sep 1978.

[5] C. R. Cutler and B. Ramaker, "Dmca computer control algorithm," in *AIChE 1979 Houston Meeting, Paper*, vol. 516, 1979.

[6] D. Prett and R. Gilette, "Optimization and constrained multivariable control of a catalytic cracking unit," in *Proceedings of Annual AIChE meeting, Houston, USA*, 1979.

[7] C. E. García, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice - a survey," *Automatica*, vol. 25, no. 3, pp. 335–348, may 1989.

[8] M. Morari and J. H. Lee, "Model predictive control: past, present and future," *Computers & Chemical Engineering*, vol. 23, no. 4-5, pp. 667–682, may 1999.

[9] N. Sandell, P. Varaiya, M. Athans, and M. Safonov, "Survey of decentralized control methods for large scale systems," *IEEE Transactions on Automatic Control*, vol. 23, no. 2, pp. 108–128, apr 1978.

[10] R. R. Negenborn, B. De Schutter, and H. Hellendoorn, "Multi-agent model predictive control of transportation networks," in *2006 IEEE International Conference on Networking, Sensing and Control*. IEEE, 2006.

[11] Y. Ding, L. Wang, Y. Li, and D. Li, "Model predictive control and its application in agriculture: A review," *Computers and Electronics in Agriculture*, vol. 151, pp. 104–117, 2018.

[12] P. Stadler, A. Ashouri, and F. Marechal, "Distributed model predictive control of energy systems in microgrids," in *2016 Annual IEEE Systems Conference (SysCon)*. IEEE, apr 2016.

[13] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings, "Agent-based control for decentralised demand side management in the smart grid," in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, ser. AAMAS '11. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 5–12. [Online]. Available: http://dl.acm.org/citation.cfm?id=2030470.2030472

[14] R. R. Negenborn, B. D. Schutter, M. A. Wiering, and H. Hellendoorn, "Learning-based model predictive control for markov decision processes," *IFAC Proceedings Volumes*, vol. 38, no. 1, pp. 354–359, 2005.

[15] Q. Liu, H. S. Abbas, J. Mohammadpour, S. Wollnack, and H. Werner, "Distributed model predictive control of constrained spatially-invariant interconnected systems and in input-output and form," in *Proceedings of the American Control Conference (ACC)*, 2016.

[16] P. Liu and U. Ozguner, "Distributed model predictive control of spatially interconnected systems using switched cost functions," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 2161–2167, 2018.

[17] M. Baranski, J. Fütterer, and D. Müller, "Distributed exergy-based simulation-assisted control of HVAC supply chains," *Energy and Buildings*, vol. 175, pp. 131–140, 2018.

[18] L. Pérez-Lombard, J. Ortiz, and C. Pout, "A review on buildings energy consumption information," *Energy and Buildings*, vol. 40, no. 3, pp. 394–398, jan 2008.

[19] J. C. Lam, K. K. Wan, and K. Cheung, "An analysis of climatic influences on chiller plant electricity consumption," *Applied Energy*, vol. 86, no. 6, pp. 933–940, jun 2009.

[20] K. Chua, S. Chou, W. Yang, and J. Yan, "Achieving better energy-efficient air conditioning - a review of technologies and strategies," *Applied Energy*, vol. 104, pp. 87–104, apr 2013.

[21] M. Sourbron, "Dynamic thermal behaviour of buildings with concrete core activation," Ph.D. dissertation, 2012. [Online]. Available: https://limo.libis.be/primo-explore/fulldisplay?docid=LIRIAS1672631&context=L&vid=Lirias&lang=en_US

[22] M. Sourbron, S. Antonov, and L. Helsen, "Potential and parameter sensitivity of model based predictive control for concrete core activation and air handling unit," *Tvs*, vol. 2, p. R4, 2013.

[23] Service Open Data - Deutscher Wetterdienst. Open data server. [Online]. Available: https://www.dwd.de/EN/ourservices/opendata/opendata.html

[24] R. Streblow, A. Xhonneux, and D. Hering, "Energiewende im quartier: Living lab energy campus im FZJ," in *Tagungsband 2. Kongress Energiewendebauen*, 2019.