

Kallabis, Thomas

Working Paper

Rolling-horizon optimization as a speed-up method - assessment using the electricity system model JMM

HEMF Working Paper, No. 06/2020

Provided in Cooperation with:

University of Duisburg-Essen, Chair for Management Science and Energy Economics

Suggested Citation: Kallabis, Thomas (2020) : Rolling-horizon optimization as a speed-up method - assessment using the electricity system model JMM, HEMF Working Paper, No. 06/2020, University of Duisburg-Essen, House of Energy Markets & Finance (HEMF), Essen

This Version is available at:

<https://hdl.handle.net/10419/234071>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



House of
Energy Markets
& Finance



Rolling-horizon optimization as a speed-up method - assessment using the electricity system model JMM

HEMF Working Paper No. 06/2020

by

Thomas Kallabis

August 29, 2020

UNIVERSITÄT
DUISBURG
ESSEN

Open-Minded

Abstract

Energy system models are limited in their scope and level of disaggregation by the availability of fast computing hardware. While improvements in hardware and solver developments have led to an increasing size of solvable models, problems with high temporal and geographical resolution remain difficult to solve in one loop. In this paper, we evaluate the use of rolling planning as a speed-up method for energy system models. In a stylized model, we highlight potential issues that occur at the boundary of optimization horizons, especially regarding time-linking constraints such as energy storage balances. In multiple configurations of the energy system model WILMAR-JMM, we investigate the tradeoff between solution quality and problem size / solution time that characterize the use of rolling planning.

Keywords : Energy system model, rolling-horizon optimization, linear programming

Thomas Kallabis
House of Energy Markets and Finance
University of Duisburg-Essen, Germany
Universitätsstr. 12, 45117 Essen
thomas.kallabis@uni-due.de

The authors are solely responsible for the contents which do not necessarily represent the opinion of the House of Energy Markets and Finance.

1 Introduction

The global energy sector is undergoing a paradigm change as the efforts to mitigate climate change have required regulatory changes and initiated technological innovations, all with the goal to reduce greenhouse gas emissions and lead to a more sustainable system. Energy systems are complex and highly disaggregated systems, in which techno-economic rationales affect the decision-making of market actors, regulators, and consumers. To allow analyses of new technologies, policy measures and individual investment decisions, mathematical energy system models have been developed and are continuously applied in various studies and analyses (Pfenninger et al., 2018). Depending on the model and the application at hand, the level of temporal, geographical and technological resolution varies widely. Historically, large-scale electricity system models have focused on the transmission grid and power plants above a certain size threshold, which limits the number of variables and results in acceptable problem sizes. With the energy transition leading to increasingly smaller assets often connected at the distribution grid level, and further challenges coming from modeling the corresponding investment decisions, there is an increasing need to allow for more detailed system modeling.

A common classification of modeling approaches distinguishes the following types: equilibrium models (applicable to electricity grid studies and equilibria in markets with strategic actors), simulation models (often so-called agent-based models), and optimization models (using the objective function for single-actor revenue or system welfare maximization) - see Ventosa et al. (2005). In the following, we focus on optimization models where the solution of an integrated optimization problem is the key characteristic. Optimizing energy system models are limited in their scope and level of disaggregation by the availability of fast computing hardware. While improvements in hardware and solver developments have continuously increased the size of solvable models, problems with high temporal and geographical resolution remain challenging. Multiple aggregation and decomposition techniques have been proposed to overcome this issue (see Section 2.1). The use of a rolling planning horizon, which shortens the time periods included in each individual optimization problem, is one such approach. Without time-linking constraints or variables, each time period could be considered as an individual optimization problem and be solved in parallel without affecting the solution. However, energy system models commonly include time-linking elements. Technical constraints such as ramp speeds or uptime and downtime constraints affect the short-term operation of thermal generators. Additionally, storage balance constraints affect the use of storages. Therefore, there exists a tradeoff between solution quality and model size / computation time when implementing a rolling planning horizon. Focusing on the electricity system model JMM (see Section 3), we therefore investigate the following research questions.

- How much reduction in problem size and runtime can be achieved using a rolling planning horizon (using a *Week* and a *Day* horizon configuration)?

- Analyzing the solution using quantitative and qualitative measures, how large is the reduction in solution quality from applying the heuristic?
- As a secondary question, how do three different computer classes (high-end desktop, server, supercomputer node) compare in terms of runtime?

The remainder of the paper is structured as follows. Section 2 gives an overview of the existing literature and provides a motivating example using a stylized model. In Section 3, we describe the model setup, data, and problem instances. Section 4 contains the results of the analysis, while Section 5 provides a summary and concludes.

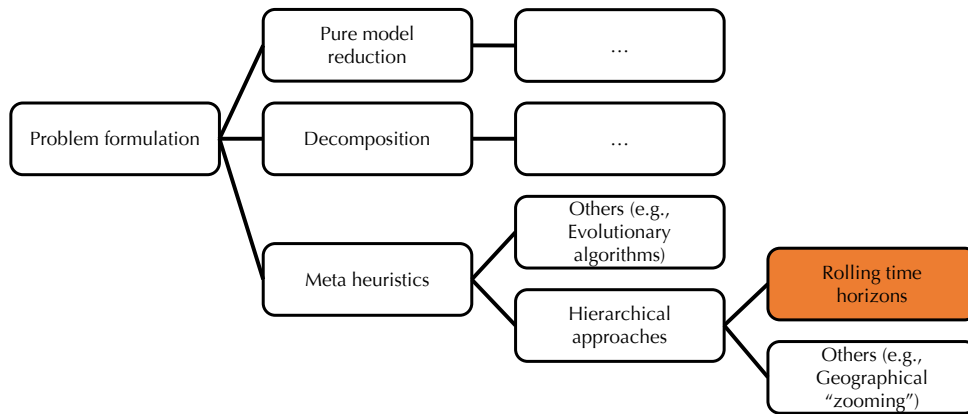
2 Methodology

2.1 Literature

In a perfect world without any computational limits, an ideal energy system model would include high resolution in every dimension: hourly time periods, endogenous investment decisions over multiple decades, nodal market clearing and DC or even AC load flow modeling. However, while not particularly challenging to implement (each of the aforementioned characteristics exists in individual models), the combination of all leads to computationally unsolvable problems given the capabilities of current computers. Therefore, methods to increase solution speed, reduce problem sizes or use decomposition approaches to allow solving large problem instances have been prevalent in energy system models for years. Cebulla et al. (2017) and Schreck (2016) provide a structure to categorize the various approaches, the first stage distinguishing at pure model reduction techniques, model decomposition, and meta heuristics (cf. Figure 1).

To illustrate this point, we subsequently give examples of the tradeoffs / methods applied by several current-generation energy system models. Leuthold, Weigt, and von Hirschhausen (2012) provides an overview of the energy system model ELMOD, which is also applied in various other publications. This model, which exists in multiple configurations (among others, MIP and NLP including a DC load flow approximation), is solved for 24 hours at a time and uses heuristics to achieve realistic thermal startup and storage operation behavior. Due to the level of detail, the optimization period is one day "for the time being". Hirth (2017) is the documentation to the European Electricity Market Model (EMMA), in which a pragmatic approach is used to overcome computational constraints. Due to the focus on market operation and renewable dispatch rather than conventional unit commitment, neither hydro reservoirs and thermal constraints on uptime or minimum load are modeled. This allows a high temporal resolution and long modeling periods, but comes at the cost of lower technological detail. Wiese et al. (2018) give an overview of the July 2017 implementation of BALMOREL. This model is quite flexible, allowing the implementation of

Figure 1: Reformulation approaches to reduce problem size (Schreck, 2016)



regional and temporal scopes depending on the application at hand. As such, both long-term investment equilibria and sub-hourly dispatch can be investigated, however not in the same model instance. For investment decisions, a rolling-planning implementation exists in which multiple years are solved simultaneously and then shifted forward.

The only systematic investigation of a rolling planning horizon in a large-scale energy system is to our knowledge Schreck (2016). In this student research paper, the system model REMIX is extended with rolling planning. The author finds computation time reductions of 35 to 50%, at the cost of objective function increases below 1.5%. In a somewhat related setup, Marquant, Evins, and Carmeliet (2015) investigate a type-day approach and a rolling planning horizon for use in a neighborhood system model called “energy hub framework”. Their analysis finds that given perfect parametrization, the rolling-horizon approach can result in a 0% gap compared to the full run, but this result does not generalize. Computation time reductions are between 0% and 30%, with larger problem instances benefitting more. JMM, the model used in this study, is an hourly-resolution one-year energy system model with a focus on optimal dispatch and (in its MIP formulation) on unit commitment. It was developed during the WILMAR project (Barth et al., 2006; Meibom et al., 2006) and has been applied in studies of market splitting / bidding zone configuration, RES integration and interconnector investment valuation. One defining characteristic of all studies is the use of a rolling planning horizon, which has not previously been compared to the results of a fully integrated optimization run. The model is described in Section 3.

A separate stream of literature provides a treatment of rolling-horizon optimization from a mathematical standpoint and with applications differing from the large-scale system models described above. Birge and Louveaux (2011) give an overview of multi-stage problems which allow recourse action. Recourse action, or the ability to change previously made decision given new information, is a key element of (stochastic) rolling-horizon methods. This is a realistic representation of applications such as hydro planning. Consequently, Guigues and Sagastizábal (2012) find a rolling-horizon configuration to be significantly faster than an SDDP formulation of a hydro reservoir scheduling problem. However, the authors add the closing remark that while they find

Table 1: Stylized model nomenclature

Symbol	Type	Unit	Description
$t \in T$	Set	N/A	Time periods
$i \in I$	Set	N/A	Generators
P_i	Parameter	EUR/MWh	Production costs
C_i	Parameter	MW	Generation capacity
D_t	Parameter	MW	Demand levels
$x_{i,t}$	Variable	MW	Production per generator
y_t	Variable	MW	Storage operation
f_t	Variable	MWh	Storage filling level
λ_t	Dual variable	EUR/MWh	Electricity price

rolling-horizon to be a good heuristic, their results are not systematic but clearly problem and parameter dependent. Another application of rolling-horizon optimization as a speedup method can be found in Kallabis, Gabriel, and Weber (2019), where the authors apply this method to MPECs and investment models.

2.2 Stylized model

A stylized energy system model serves as a motivating example for considering time-linking constraints when implementing a rolling horizon into an optimization model. Consider a simple four-period ($T = \{t1, ..., t4\}$) dispatch problem, whose nomenclature is stated in Table 1. Assume two demand levels, peak and offpeak, where $D_{peak} > D_{offpeak}$ and $t1, t2$ are offpeak, while $t3, t4$ are peak. There are two generators, gas and coal, with $P_{gas} > P_{coal}$. Also, let $D_{peak} > C_{coal}$, so that during peak hours, the capacity of the cheaper technology is insufficient to satisfy demand. Finally, consider an ideal storage with no production or storage volume limits. Its operation is given by y_t and its filling level by f_t . The objective function, stated in Equation 1, is the minimization of production costs. Equation 2 enforces load serving, with the dual variable λ_t giving the price of electricity. Additional constraints apply to maximum production capacity (Equation 3) and storage balance across time periods (Equation 4).

$$\min z = \sum_T \sum_I x_{i,t} P_i \quad (1)$$

$$y_t + \sum_I x_{i,t} \geq D_t \quad : \lambda_t \quad \forall t \in T \quad (2)$$

$$x_{i,t} \leq C_i \quad \forall i \in I, \forall t \in T \quad (3)$$

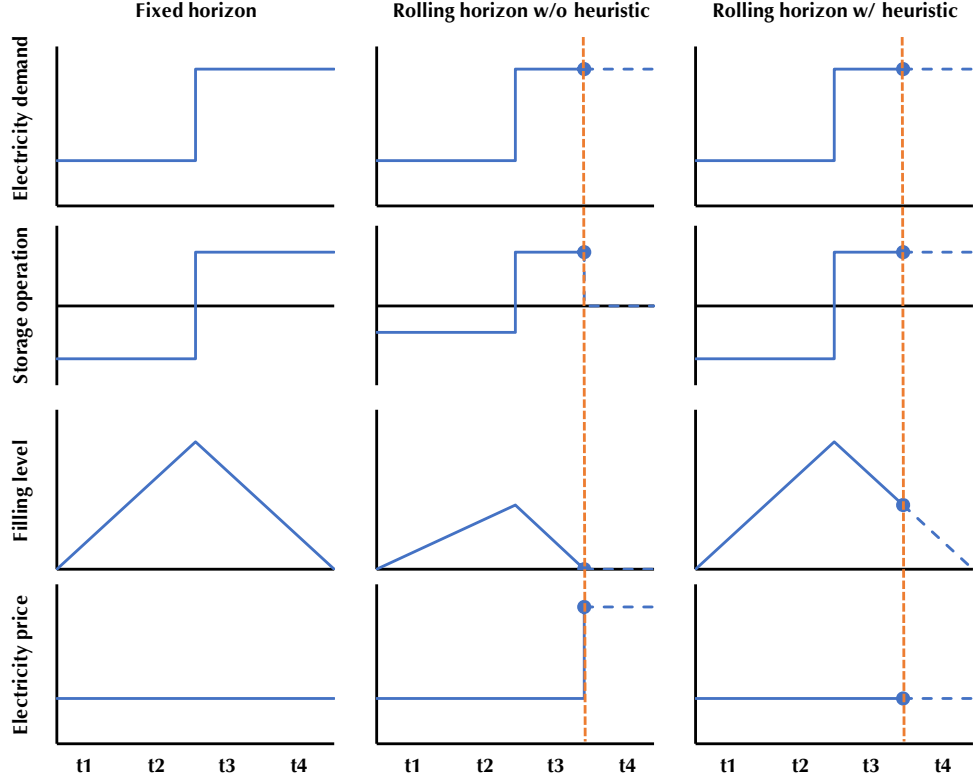


Figure 2: Results of the stylized model for three configurations: integrated optimization (left), rolling horizon without heuristic (center), rolling horizon with heuristic (right)

$$f_t = f_{t-1} - y_t \quad \forall t \in T \quad (4)$$

In the following, we consider three configurations. The first is an integrated, fixed-horizon optimization of all four time periods t_1, \dots, t_4 . The other two both use a shortened optimization horizon in which at first, only time periods t_1, \dots, t_3 are included in the optimization, while t_4 is optimized subsequently. The (simplified) results are shown in Figure 2. For the integrated optimization over the entire time horizon (shown on the left), the result is trivial. During the initial offpeak periods t_1 and t_2 , the cheaper coal generator produces at maximum capacity to serve demand and charge the storage. During the subsequent peak demand periods, the storage discharge produces enough electricity so that the more expensive gas generator is not required to produce electricity to satisfy demand. The storage charges during offpeak and discharges during peak, leading to a uniform electricity price at the cost level of the cheaper coal generator. This result is also cost-optimal.

A rolling planning horizon is characterized by including only a subset of time periods into each individual optimization problem. While overlapping time periods are not strictly required, it is a common trait in rolling-horizon energy system models. For the stylized model, we assume two disjoint planning horizons: the first including time periods t_1, \dots, t_3 , the second including only t_4 . However, this issue of end-of-horizon effects as described in the following can also apply

to overlapping horizons. Without any additional constraints, the rolling-horizon optimization of t_1, \dots, t_3 leads to the results shown in the middle of Figure 2. The storage charges and discharges as before, however it is empty at the end of time period t_3 since the model is no longer aware of the additional peak load time period t_4 . The individual optimization of said time period then requires the use of more expensive gas-fired generation, thus leading to higher electricity prices in t_4 and an overall higher system cost. This result is clearly suboptimal and results from the lack of awareness of the time period t_4 in the initial planning loop, a concept known as an end-of-horizon effect.

There are multiple ways to overcome this issue. The key problem is that the storage content has no value at the end of the planning horizon, therefore giving the model an incentive to empty it completely. Two main approaches are common. Value-based approaches give the storage content an economic value by introducing an additional objective function term reflecting the storage value, i.e. the storage volume multiplied by a shadow price term. The second approach is volume-based, i.e. fixing the final storage filling level to a predetermined value (see Ripper, 2000). Hybrid approaches also exist, where a penalty term introduced to the objective function penalizes deviations from a predetermined storage filling level path. All approaches share the difficulty of parameter calibration, as choosing optimal parameters requires knowledge of the optimal solution over the entire planning horizon. This solution is not available - if it were, the rolling-horizon configuration would be pointless to begin with. With ideally chosen parameters, the results of the rolling-horizon would be identical to the entire horizon optimization. This result is shown in Figure 2 on the right.

Given optimal parameter choices, both value-based and volume based approaches can lead to identical results. However, one issue of volume-based approaches is their inflexibility. When using time series data of historical years as input for simulation of future years, the difference in the overall system usually renders the historical filling level time series obsolete. Fixing filling levels at the end of the horizon also strips the model of flexibility. Value-based approaches on the other hand allow temporary deviations, since a value heuristic gives incentives based on the overall system state. Ripper (2000) finds value functions generally more efficient in modeling hydro dispatch, at least when not systematic forecast error is present. Therefore, this is the heuristic we use in the following.

3 Case study setup

In the following, we describe the model framework and the data used in our case study to evaluate rolling-horizon optimization in a large-scale electricity market model.

3.1 Model JMM and data

The Joint Market Model JMM is a large-scale electricity system model. It was originally developed as part of the EU-funded research project WILMAR and is a linear optimization model of the European power system. JMM is implemented in GAMS and commonly CPLEX is used as the solver for both its LP and its MIP configuration. One defining model aspect is the dynamic geographical scope, allowing to include Europe in its entirety (Meibom et al., 2006; Weber et al., 2009) or individual countries. There is also a high level of detail in modeling the various constraints in the thermal generation fleet (upper and lower bounds on electricity output, dynamic uptime and downtime constraints, combined optimization of electricity and heat output, integrated optimization of long-term and short term storages, among others). The model can be configured to model individual power plant blocks, which leads to a mixed-integer optimization problem (Trepper, Bucksteeg, and Weber, 2015), but is commonly used in a linear approximation, in which similar generation assets are grouped into "vintage classes" with common technological and operational characteristics.

The high level of detail corresponds with high computational requirements. During its inception, the model was designed to run in a rolling-horizon configuration of 24 to 36 hours. This significantly reduces the size and solution time of individual instances, and in fact was required since at the time no integrated optimization of an entire year appeared feasible. During the time since the model's inception in 2006, there have been significant advancements in the overall capability of home desktop and server computers, and the machines available today facilitate another attempt for an integrated year-round optimization. The availability of supercomputing hardware at the Jülich Supercomputing Centre (through the BEAM-ME project) allows another layer of research, comparing three hardware configurations in the following.

We configure the model using a three-country geographical scope: Germany, Austria and Switzerland (abbreviated by the acronym "D-A-CH"). While a full-fledged run of the entire European electricity system was also initially considered, the resulting problem size proved prohibitive (see below). However, the scope of our case study provides the relevant variety: the Alpine countries, which are largely driven by dispatchable seasonal hydro storages, provide flexibility to the Germany market, which is heavily reliant on less flexible thermal and inexpensive renewable generation capacity. Particularly Switzerland, a country in which hydropower is essential for security of supply, makes for a good test case.

We parametrize the model using 2016 historical data for fuel prices, electricity demand, and renewable production as well as hydro inflow. Hydro storage filling levels are fixed to their historical values at the beginning (all configurations) and the end of the year (only enforceable for the integrated Year configuration). Since the three D-A-CH countries are heavily interconnected with their European neighbors, we fix the historical hourly interconnector flows to third countries as parameters. The used data is outlined in Table 2. Thermal unavailability data was provided by Beran, Vogler, and Weber (2017). For comparison, we also provide metrics for a problem instance

Table 2: Data used for the three-country D-A-CH case study

Parameters	Source	Specification
Load (time series)	ENTSO-E transparency	Total Load - Day Ahead / Actual
Demand (annual)	IEA electricity statistics	Final electricity consumption (incl. losses)
Renewable production	ENTSO-E transparency	Hourly production per production type
3rd country exchange	Dito	Scheduled commercial exchanges
Net transfer capacity	Dito	Day-ahead forecasted capacities, aggregated into monthly
Thermal power plant fleet	HEMF-internal database	Capacities, technology and age-specific parameters
Unavailability	EEX transparency	Technology-specific data for DE/AT, manually adjusted for CH
Hydro reservoir levels	ENTSO-E transparency	Water reservoirs and hydro storage plants
Hydro reservoir inflow	Dito	Calculated from filling levels and production, data cleaned
Coal price	Energate-messenger.de	API#2 (CIF ARA)
Gas Price	Dito	NCG Day-ahead
Oil Price	Dito	OPEC basket
CO2 Price	Dito	EU allowances (EUA)

whose geographical scope includes Europe in its entirety rather than the D-A-CH subset (based on the work given in Felling et al. (2019) and Felten et al. (2019)). Since it is only used for comparisons regarding problem size and solver time (but not solution quality, incl. objective function value), we do not state the parametrization details here.

3.2 Optimization horizon configurations

The optimization horizon configurations we use for analysis are described in the following. We compare three configurations: a closed loop optimization over an entire year (called *Year* configuration), which encompasses 8760 hours and is used as both a reference for problem size and a benchmark regarding solution quality. We compare its results to two different rolling-horizon implementations, one using weekly, one using daily planning horizons. The *Week* configuration consists of 180 hours, of which 12 hours are overlap to the following optimization loop. The *Day* configuration consists of 36 hours of optimization, of which 12 hours are overlap from the previous loop. However, the first 12 hours of the horizon (from noon the previous day until midnight) are only included as a "ramp-up" period for which all variables are fixed. Therefore, the optimization horizon for the solver is only 24 hours long.¹ The structure of both rolling-horizon configurations is given in Figures 3 and 4.

One noticeable difference between the two rolling-horizon implementations is the treatment of end-of-horizon effects. The *Week* configuration has a 12-hour overlap, which provides an incentive to the solver not to shut down thermal plants at the end of the horizon.² Contrarily to this, the

¹This setup is derived from a separation in day-ahead and intraday loops, which is commonly used in JMM applications. The implementation in this paper can be interpreted as a sequence of day-ahead dispatch problems with no in-between intraday market.

²End-of-horizon effects for storages are treated separately using a heuristic, see Section 3.3.




Figure 3: Day horizon looping structure




Figure 4: Week horizon looping structure

Day configuration does not include any overlap: each optimization loop ends at midnight, and the optimization of the following one picks up with the first hour of the next day. As outlined before, the tradeoff in rolling-horizon optimization achieves problem sizes which can be generated and / or solved in reasonable time given available computation infrastructure, but it comes at the cost of reduced solution quality. We therefore provide metrics for both.

3.3 Hydro reservoir shadow price heuristic

The optimal use of seasonal hydro storages is one key determinant of system costs in heuristics such as rolling-horizon optimization. The hydro storage modeling in JMM includes upper and lower constraints on reservoir filling levels, limits on hourly reservoir production, and a dynamic storage filling level constraint. Since our analysis is focused on the economically incentivized (i.e. value-based) reservoir usage, we do not show all hydro equations here. A detailed description

of the hydro modeling of JMM is given in Meibom et al. (2006) and Jahns, Podewski, and Weber (2019).

As outlined in Section 2.2, giving stored water a value at the end of the optimization horizon is key in rolling-horizon configurations. In integrated optimization runs, the value of water is given implicitly as the dual variable of the dynamic storage equation, JMM's version of which is given in Equation 5. The dual variable, also known as shadow price, mathematically states the marginal change to the objective function value given a small relaxation to the constraint. We define this dual variable $WV_{t,a}^{dual}$, giving the water value per time period $t \in T$ and area $a \in A$ in the model.

$$FillLevel_{t,a} = FillLevel_{t-1,a} - Prod_{t,a} + Inflow_{t,a} : WV_{t,a}^{dual} \quad \forall t \in T, a \in A \quad (5)$$

In rolling-horizon configurations, water remaining in the reservoir at the end of a horizon has no intrinsic value to the model. This can be modified by making two adjustments in the model code. In a first step, the objective function is amended to include a value term for the reservoir content filling level at the end of the planning horizon. We denote the time period that defines the end of the horizon as $p \in P \subseteq T$, where p is the final time period in the current optimization loop while P is the set of all final time periods (in any optimization loop). The objective function term is given in Equation 6. The filling level is a primal variable in all optimization configurations. The shadow price parameter $WV_{p,a}$ would ideally be set to the value of the dual variable $WV_{p,a}^{dual}$ from the Year run as deccribed in Equation 5, but this value is not known in the rolling-horizon configurations. Therefore, it calculated using the heuristic described in the following.

$$+ \sum_A \sum_P FillLevel_{p,a} * WV_{p,a} \quad (6)$$

For the first loop, $WV_{p,a}$ is initialized with a pre-determined start value. Beginning with the second loop, the value of $WV_{p,a}$ is updated as described in Equation 7. In each loop, the previous loop's $WV_{p,a}$ is modified by calculating the difference between a predefined reference time series $FillLevel_{p-1,a}^{ref}$ and the actual filling level $FillLevel_{p-1,a}$. This difference (in percent) is scaled by parameter α and added to the previous $WV_{p-1,a}$. If the reservoir level is lower than in the reference profile, the shadow price increases, incentivizing less water use in the following optimization loop. Vice versa in the opposite case, this approach is beneficial in that it creates an economic incentive to keep the reservoir filling level relatively close to the reference time series. At the same time, it allows deviations when the objective function can be improved significantly by deviating from the reference path.

$$WV_{p,a} = \max(1, WV_{p-1,a} + \alpha \times (FillLevel_{p-1,a}^{ref} - FillLevel_{p-1,a})) \quad \forall a \in A, p \in P \subseteq T \quad (7)$$

Table 3: Computer classes used in this case study

Computer class	High-end desktop	Dell R620 Server	JUWELS HPC node
Build year	2019	2014	2018
CPU	Core i9 9900k	2 x Xeon E5-2637 v2	2 x Xeon Platinum 8168
Cores (total)	8	8	48
Core speed [GHz]	3.60	3.50	2.70
Boost speed [GHz]	5.00	3.80	3.70
RAM [GB]	96	192	192

This heuristic is parametrized with the following input: the initial $WV_{p,a}$, the reference profile $FillLevel_{p,a}^{ref}$, and the adjustment scalar α . The latter we choose as 1 or 100%, a value which has proven robust during years of applying JMM in various studies and is not varied in the following. For the other two parameters, there is an *Optimized* and a *Historical* parametrization. The optimized parametrization uses results of the *Year* configuration as input: $WV_{p1,a}^{dual}$ is assigned as the initial value of $WV_{p,a}$, while $FillLevel_{p,a}^{ref}$ uses the realized $FillLevel_{p,a}$. The advantage of these optimized parameters is that they allow to evaluate the hydro heuristic under idealized conditions. The disadvantage is that these parameter values are not commonly known in rolling-horizon applications: if yearly results were available, the rolling horizon configuration would be obsolete. Therefore, we also test a configuration using historical inputs which are known independently of the *Year* run: the initial $WV_{p,a}$ is set at the average historical day-ahead electricity price (per country), while $FillLevel_{p,a}^{ref}$ uses the historical filling level time series. While this input is not necessarily optimal, it allows to test the heuristic given realistic input parameters.

3.4 Computing hardware and problem sizes

We compare three different computer classes: a custom-built, high-end desktop; a Dell R620 server workstation; and a compute node of the JUWELS supercomputer at Jülich Supercomputing Center (JSC), computation time on which was made available through the BEAM-ME project. Descriptions of the three classes are listed in Table 3. The main differences between the classes lies in their computing architecture. While all three use Intel processors, their clock speeds differ significantly. RAM size on the desktop is chosen at 96 GB so that the largest individual problem instance remains solvable.

The computational effort required to solve a problem with CPLEX scales with problem size. Table 4 gives an overview of the problem size for different model configuration. Regarding geographical scope, it is noteworthy yet expected that the problems including the entire European electricity system are roughly four times the size (referring to the number of non-zero elements) of those only modeling the D-A-CH region. Of course, this is directly related to the number of nodes / countries

Table 4: Problem sizes under different rolling horizon configurations and parametrizations

Horizon configuration	Geographical scope	Stats (typical instance, before pre-solve)			Relative delta	Peak RAM usage [GB]	Relative delta
		Rows	Columns	Nonzeroes			
Year	D-A-CH	61,863,127	35,378,920	228,229,549	0.00%	73.19	0.00%
Week	D-A-CH	1,271,161	732,880	4,694,809	-97.94%	1.49	-97.96%
Day	D-A-CH	236,593	129,421	843,513	-99.63%	0.25	-99.66%
Year*	Europe	275,076,776	200,145,832	914,629,816	0.00%	393.60	0.00%
Week	Europe	6,236,281	4,144,095	20,788,416	-97.73%	7.71	-98.04%
Day	Europe	943,321	732,507	3,025,716	-99.67%	1.40	-99.65%

***Extrapolation based on Day and Week configuration results, no actual run data available**

and power plants, which is smaller in the latter case. Please note that computational restrictions (in particular, the available amount of RAM) made it impossible to compile or solve a yearly planning configuration for the European geoscope. Therefore, the statistics given in Table 4 are extrapolated from the *Day* and *Week* configurations using the ratios from the D-A-CH runs. This gives an estimate of the computational efforts required to generate and solve this configuration. Since current commercial servers can be configured with 512 GB of memory (or more), sufficient budget and time likely allows to solve these problems.

The size differences between the different horizon configurations also yields relevant insights. The problem size (again, referring to non-zeros) scales with the horizon size in a roughly linear fashion. The optimization horizon of the *Week* configuration includes 5 times the number of the 12hour configuration (180 vs 36 hours), while the problem size ratio is 5.6. Comparing *Year* to *Week*, the it is an almost perfect match: 48.6 times the hours in the optimization horizon yield 48.7 the number of non-zeros. RAM usage follows a similar pattern, as one would expect: the amount of memory required to store a problem instance is directly driven by its row, equation, and non-zero element count. This first assessment of the model configuration clearly shows the lower computational requirements of the rolling-horizon configurations as measured by problem size and required memory.

4 Results

4.1 Computational benefits

The main computational burden comes from the time it takes to solve a problem, commonly known as solution time. While sufficient memory is a necessary condition, the difference in solution time may be a deciding factor when determining the value of a shortened optimization horizon. The results for the three different computer classes are outlined in Table 5. Please note

Table 5: Solution times under different rolling horizon configurations and parametrizations

Horizon configuration	Geographical scope	High-end desktop		Dell R620 Server		JUWELS HPC node	
		Total Solver time [hours]	Relative delta	Total Solver time [hours]	Relative delta	Total Solver time [hours]	Relative delta
Year	D-A-CH	10.50	0.00%	20.10	0.00%	13.53	0.00%
Week	D-A-CH	0.28	-97.33%	0.60	-97.03%	0.33	-97.54%
Day	D-A-CH	0.14	-98.71%	0.29	-98.57%	0.14	-98.93%
Year*	Europe	239.76	0.00%	452.73	0.00%	302.99	0.00%
Week	Europe	7.94	-96.69%	17.42	-96.15%	9.42	-96.89%
Day	Europe	2.35	-99.02%	4.57	-98.99%	2.38	-99.22%

*Extrapolation based on Day and Week configuration results, no actual run data available

that all figures given in the table refer to the overall solution time, i.e. one loop for the year, 52 loops for the week, and 365 for the *Day* configuration. Obviously, computation times for the D-A-CH region are much lower than for entire Europe as the geographical scope of Europe is significantly larger than D-A-CH (again, the results for the EU-year configuration are extrapolations). Within each computer class, longer horizons also lead to higher solutions times. However solution time increases less than linear with problem size: for the smaller D-A-CH geoscope, the *Week* configuration takes "only" twice as long to solve as the *Day* configuration, while for the larger European model, this ratio is closer to 4. However, all these statistics refer to overall solution time (i.e., the sum over all problem loops) - the solution times of an individual loop of each horizon configuration are significantly shorter, indicating a superlinear relationship between problem size and solution time.

The results for the different computer classes are somewhat inconclusive. It is clear that the R620 server lacks behind the other two classes, leading to a 50% higher solution time. It appears that this is not due to pure core speed though, since both the JUWELS node and the R620 server operate with similar boost clocks. Potentially the compute architecture plays a role, since the Core i9 of the high-end desktop and the Xeon Platinum of the JUWELS node are several years younger than the Xeon E5 of the R620 server. It should however be noted that given their higher core count, the server and especially JUWELS would significantly benefit from parallelization. While CPLEX, the solver used in this study, allows parallelization, studies involving the JMM do not commonly use it.³ The underlying Simplex algorithm is sequential at its core⁴. In any case, the solution time of the Europe-Year configuration (which we extrapolate) would be close to two weeks, which may or may not be acceptable depending on the task at hand.

³Parallelization is very limited in CPLEX when requiring deterministic results, that is to say identical results for a given model instance. An opportunistic implementation using multiple CPU threads (e.g. barrier and concurrent optimizers) is available, but are not commonly used in JMM to allow reproduction of run results. www.gams.com/latest/docs/S_CPLEX.html#CPLEXparallelmode

⁴All runs use the CPLEX parameters `lpmethod 0` and `threads 1`, leading the solver to default to its dual simplex configuration. www.gams.com/latest/docs/S_CPLEX.html#CPLEXlpmethod

Table 6: Objective function values under different rolling horizon configurations and parametrizations

Horizon configuration	Parameter selection	Uncorrected Vobj [mEUR]	End-of-year Res. filling level [TWh]	Filling level correction [mEUR]	Corrected Vobj [mEUR]	Absolute delta [mEUR]	Relative delta
Year	N/A	8,687.52	5.76	0.00	8,687.52	0.00	0.00%
Week	Optimized	8,670.02	4.48	58.03	8,728.04	40.52	0.47%
Week	Historical	8,669.59	4.40	61.90	8,731.49	43.97	0.51%
Day	Optimized	8,670.59	4.42	60.90	8,731.48	43.96	0.51%
Day	Historical	8,692.79	4.31	66.17	8,758.96	71.44	0.82%

4.2 Objective function impact

The previous paragraphs have outlined the computational benefits of the rolling-horizon configurations. However, these benefits come at a cost: the solutions found are different from that of the integrated yearly optimization run. This leads to "less optimal" results, ultimately caused by less efficient use of resources in the electricity system. In a first stage, we aggregate this solution difference to the objective function value. These results are given in Table 6. It appears that the rolling-horizon configurations lead to lower system costs than the integrated run. However, these initial objective function values do not show a complete picture. This is related to the use of hydro storages. When comparing the end-of-year filling levels of the *Day* and *Week* configurations to the integrated yearly run, it is apparent that they are lower than in the *Year* run. While these are likely not systematic and could be higher in different problem instances, they do indicate that the proposed shadow value heuristic in the rolling-horizon configurations does not yield perfect incentives to use the available water as efficiently as in the yearly run. *Ceteris paribus*, using up more of the stored water reduces the required fuel used in thermal generation assets, therefore lowering the objective function value. To make the different configurations comparable, we propose an adjustment of the objective function by the "filling level adjustment term" Adj_{FL} as given in Equation 8. By this logic, we value the remaining storage content using the final hydro shadow price from the *Year* run. This corrects the objective function value for the difference in end-of-year reservoir filling level by using the hydro shadow price from the *Year* run as common valuation basis.⁵

$$Adj_{FL} = (FillLevel_{p_{final},a,RH} - ResLevel_{p_{final},a,Year}) * WV_{p_{final},a}^{dual} \quad (8)$$

The end-of-year filling levels, adjustment terms, and the adjusted objective function values are given in Table 6 (for brevity, all figures are aggregated among the D-A-CH countries). With the adjustment, the order of objective function values is as one would expect: the integrated optimization leads to the lowest system costs, while the *Week* rolling-horizon configuration comes

⁵It is important to use the same shadow price for all adjustments in order to avoid inconsistent valuations. The use of the shadow price from the *Year* run is the logical choice as this is the reference configuration.

out ahead of the *Day* configuration. Given the fact that shorter optimization horizons lead to more frequent use of the imperfect hydro heuristic (as well as potential end-of-horizon effects), this is a plausible outcome. Also, the optimized parametrizations of the heuristic lead to lower objective function values than the ones using historical parameters, again, an expected outcome. When looking at the absolute and relative delta between the configurations, it appears that the impact on solution quality is quite low in cost terms. However, there are two points which suggest additional analysis: first, the absolute differences are in excess of 70 million EUR for one year, which may be a relevant welfare change for some applications. Secondly, solution quality cannot be measured by objective function values alone. Therefore, we devote the remainder of this chapter to providing a more qualitative analysis of the solution variations between the different configurations centered around the efficient use of fuels and water.

4.3 Qualitative solution assessment

One key difference between the configurations is the use of hydro storages during the year. It is clear that the integrated yearly optimization must yield the optimal⁶ result. Figures 5 and 6 show the closely related shadow price of water and the filling level of the Swiss hydro storages⁷. The two are linked, since the value of reservoir content strongly affects its use. In the *Year* configuration, there are three discrete shadow price levels: 33.5 EUR / MWh from January to late May, 30.8 EUR / MWh until mid October, and 45.4 EUR / MWh until the end of the year. The changes in water value occur in May and October at those points in time when the reservoir reaches its minimum and maximum filling level, as is expected in a deterministic setting (Böcker and Weber, 2019). The shadow price development also corresponds with the filling level, as the lowest shadow price level coincides with the largest inflows due to snow melting during early summer. The large jump in water value in the fourth quarter is also related to an unplanned outage of the Leibstadt nuclear power plant, the largest of Switzerland's five nuclear stations. The plant went offline for its main revision in August.⁸ Initially planned for only four weeks, then extended to eight weeks, and finally to six months, this caused a shortage in relatively cheap nuclear power supply for the remainder of 2016, and increased Swiss electricity prices and thereby, the value of stored water.

The reservoir water usage in the rolling-horizon configurations is mainly driven by the water value. The water value in turn is affected by the filling level compared to the reference profile. During the first and second quarters, both the *Day* and the *Week* configuration use slightly more water than the *Year* run, leading to lower filling levels than the reference profiles (Figure 6) and higher shadow prices than the *Year* run (Figure 5). It is visible that the water value adjustment heuristic (as described in Section 3.3) leads to shadow prices and filling levels which are relatively close to the

⁶"Optimal" refers to optimality in a mathematical sense, i.e., that the objective function value cannot be improved further. Evaluating the model fit against historical data may lead to a different result, since the model premise of perfect foresight is not present in reality. This is briefly discussed in Section 5.

⁷For brevity, we limit this evaluation to Switzerland. It is the country in which hydropower has the largest generation share within our case study, and the results for Germany and Austria do not differ significantly.

⁸<https://www.kkl.ch/unternehmen/medien/archiv/2016.html>

Figure 5: Swiss reservoir shadow price under different horizon configurations


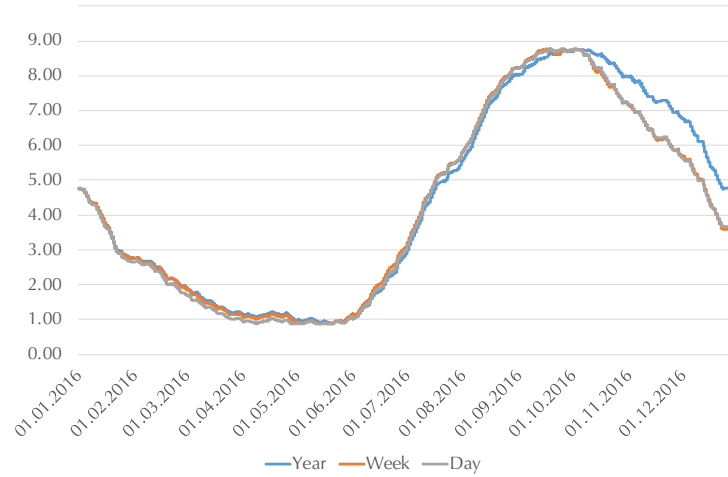
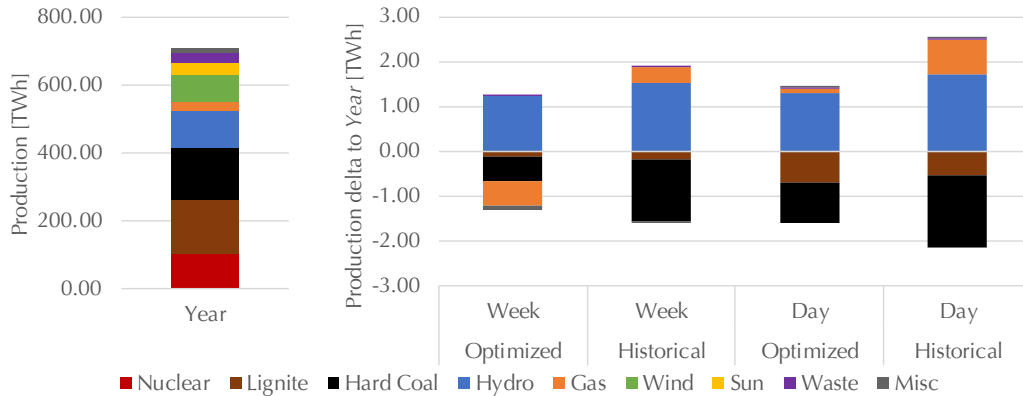


Figure 6: Swiss reservoir filling levels under different horizon configurations



Year run. However, they both deviate significantly in the fourth quarter when the shadow price jumps in the *Year* configuration. Here it is obvious that the heuristic requires several weeks to reach the higher optimal value, in the process leading to an above-optimal use of hydropower in the two rolling-horizon configurations. This appears to be the main cause of the lower end-of-year reservoir filling levels in all rolling-horizon runs. This deviation could potentially be diminished by setting the adjustment speed parameter α to a higher value. On the other hand, a higher value of α would likely lead to more volatile water values. Overall, the heuristic performs as expected: imperfect compared to the *Year* run, but leading to reasonable results within its limitations. Note that we only discuss configurations using the optimized parameters for reference profile and starting shadow price and not the ones using historical parameters. For brevity, these were omitted here. The qualitative results do not differ significantly from the ones shown, but the objective

Figure 7: Electricity production volumes from different fuels in the *Year* configuration (left) and deltas for the remaining configurations (right)



function values given in Table 6 show the less optimal outcome using the historical parameter values.

Another assessment of solution quality are differences in production volume per technology / fuel type. Figure 7 shows the production volumes for the different configurations. The total electricity production for the *Year* configuration is 710 TWh, of which 22% each are hard coal and lignite, and 15% each are nuclear and hydro (both run-of-river and reservoirs), while wind and solar make up for a combined total of 16%. This result mirrors the generation system of the D-A-CH case study region, being strongly influenced by the German fuel mix. On the right, we show the deltas between the *Year* run and the other configurations. Here, it is noteworthy that the rolling-horizon runs all use more hydro reservoir content than the integrated run (mirroring the data shown in Table 6) because of the sub-optimal use of hydropower due to the water value heuristic. This is made up for by using less mid- and baseload fuels, mainly coal and lignite, a pattern visible in all configurations. However, the production deltas of gas-fired generation assets (as a more expensive midload to peak technology) deviates from this pattern. In the *Week Optimized* configuration, the model uses less natural gas, while the other three configurations show additional use of this fuel. The difference shows the impact of the horizon configuration on the technology replaced by the additional hydropower. In *Week Historical* and the two *Day* configurations, the solver only manages to replace relatively cheap coal and lignite production, and actually requires additional gas-fired generation to satisfy demand. In the *Week Optimized* configuration however, almost 50% of the additional hydropower usage replaces natural gas generation, explaining the lower system costs as expressed by the lower objective function value.

Finally, we review the average day-ahead market prices (see Table 7). We also show the historical day-ahead prices for comparison, but note that this study has not aimed to provide a model back-testing (which would require additional parametrization). It is obvious that prices increase with shorter horizon length and less-than-optimal parametrization of the water value heuristic. Unlike

Table 7: Average electricity prices in the D-A-CH countries and a weighted average under different model configurations and historical market results

Configuration	Parameter selection	Average Price (uncorrected)		
		AT	CH	DE
Historic Day-ahead prices		28.96	37.83	28.96
Year	N/A	30.24	34.70	30.21
Week	Optimized	30.18	34.30	30.16
Week	Historical	30.30	35.32	30.28
Day	Optimized	30.81	35.19	30.79
Day	Historical	30.90	36.22	30.88

the objective function values shown in the previous subsection, these prices are not and cannot be corrected for different end-of-year reservoir filling levels. Three of the rolling-horizon configurations (*Week Historical* and both *Day* configurations) show higher electricity prices, in line with higher (corrected) objective function values. This is likely another sign of the less-than-optimal dispatch resulting from the shortened optimization horizon. Electricity prices, in system models defined as the shadow price of the market-clearing condition, are set by the marginal plant. Since the merit-order is by definition a convex function, small changes in production volumes can lead to large price increases in some hours and small price decreases in others. The net effect of this appears to be positive, an argument supported by the fact that we observe the largest individual increase between *Year* and *Day Historical* (1.52 EUR / MWh) in Switzerland, the country with the largest influence of hydro reservoir production. The exception in this case is the *Week Optimized* configuration, which leads to lower electricity prices than the *Year* run. This is in line with the results on fuel usage, which in this configuration showed less use of (expensive) natural gas than the *Year* run.

5 Conclusion

The aim of this paper is to investigate the tradeoff between runtime and solution quality that comes from implementing a rolling planning approach. Referring to the research questions raised in Section 1, the answers are as given below. Please note that all statements are made in relation to the tested problem instances and generalized where they can be. However, many of the answers are problem-specific, and different problem instances likely require to revisit each topic.

- The reduction in problem size and computation time are significant. Depending on the horizon configuration and the geographical scope, the problem size can be reduced by 97.9% to 99.7%. Clearly this is an improvement that allows running the model on much smaller machines or implementing much larger, more detailed problem instances. The reduction

in runtime is in a similar range, decreasing by 97.0% to 99.0%. Again, this creates flexibility when configuring problem instances encompassing e.g. the entire European electricity system.

- Regarding solution quality, results are not straightforward. On the one hand, the relative increase in the objective function value or system cost is small, remaining below 1% in all instances. On the other hand, the absolute deltas are around 40 million EUR and exceed 70 million EUR in one configuration. When comparing the welfare impact between two model runs, e.g. to analyze the impact of a new interconnector (see Trepper, 2015), this difference may be meaningful. On the other hand, it is apparent that the *Week* horizon configuration performs better than the shorter *Day* one, suggesting that some solution quality can be preserved by this somewhat longer planning horizon. The implementation of the water reservoir heuristic plays a role as well, and even better results are likely to be achieved by optimal parametrization.
- Regarding the different computer classes, there appears to be a significant benefit from using the most recent architecture. While the two classes using 2018 and 2019 hardware (the high-end desktop and the supercomputer node) vary in (boost) clock speed, both outperform the 2014 server by up to 100% (cutting runtime in half). This suggests improvements in the computer architecture or lower-level adjustments of the solver to more recent systems. In any case, the results highlight the relevance of improvements in computing hardware to the application of energy system models.

As mentioned in Section 3, the JMM model is used in various applications, often investigating the impact of individual investments / assets on overall system welfare. It would clearly be problematic if the system cost increase introduced by the rolling horizon approach were to distort these results. However, these analyses commonly involve comparing two relatively similar runs with each other: in this case, any systematic deviation due to the rolling horizon configuration would be present in both, and only secondary effects (likely smaller) would affect the quantification of the deltas. Therefore, analyses derived from comparing two similar model instances using the same rolling-horizon configuration are likely still valid. However, please note that this approach still requires correcting for any changes in the end-of-year hydro reservoir filling level.

When focusing on production volumes per technologies, the heuristic also appears to perform reasonably well. As this metric is commonly used to assess the amount of carbon dioxide emissions, it is clearly relevant. For any focus on prices (also used in model backtesting to improve parametrization), the differences are larger and likely meaningful. Here, additional work would be useful to understand which horizon configuration comes closest to modeling actual market behavior - after all, reality has no perfect foresight, so an integrated year-long optimization loop may actually not be the best proxy.

References

- Barth, R., Brand, H., Meibom, P., & Weber, C. (2006). A stochastic unit-commitment model for the evaluation of the impacts of integration of large amounts of intermittent wind power. *2006 9th International Conference on Probabilistic Methods Applied to Power Systems, PMAPS*.
- Beran, P., Vogler, A., & Weber, C. (2017). Kurz- und mittelfristige Preisprognosen: Auswahl optimaler Modellierungsansätze unter Berücksichtigung des Prognosehorizonts. *VDI-Berichte*, 2303.
- Birge, J. R., & Louveaux, F. (2011). *Introduction to Stochastic Programming* (2nd). Springer International Publishing.
- Böcker, B., & Weber, C. (2019). How will storage affect future energy systems? Fundamentals of storage operation and investment. *HEMF Working Paper*.
- Cebulla, F., Cao, K.-K., Wetzel, M., Fiand, F., & Gils, H. C. (2017). Reducing computing times in optimising energy system models - challenges and opportunities. *Presentation at OR2017, Berlin*.
- Felling, T., Felten, B., Osinski, P., & Weber, C. (2019). Flow-Based Market Coupling Revised - Part II: Assessing Improved Price Zones in Central Western Europe. *HEMF Working Paper*.
- Felten, B., Felling, T., Osinski, P., & Weber, C. (2019). Flow-Based Market Coupling Revised - Part I: Analyses of Small- and Large-Scale Systems. *HEMF Working Paper*.
- Guigues, V., & Sagastizábal, C. (2012). The value of rolling-horizon policies for risk-averse hydro-thermal planning. *European Journal of Operational Research*, 217, 129–140.
- Hirth, L. (2017). The European Electricity Market Model EMMA Model documentation. *Neon Neue Energieökonomik*.
- Jahns, C., Podewski, C., & Weber, C. (2019). Supply Curves for Hydro Reservoirs - Estimation and Usage in Large-Scale Electricity Market Models. *HEMF Working Paper*.
- Kallabis, T., Gabriel, S. A., & Weber, C. (2019). Strategic generation investment using a stochastic rolling-horizon MPEC approach. *HEMF Working Paper*.
- Leuthold, F. U., Weigt, H., & von Hirschhausen, C. (2012). A Large-Scale Spatial Optimization Model of the European Electricity Market. *Networks and Spatial Economics*, 12, 75–107.
- Marquant, J. F., Evins, R., & Carmeliet, J. (2015). Reducing computation time with a rolling horizon approach applied to a MILP formulation of multiple urban energy hub system. *Procedia Computer Science*, 51, 2137–2146.
- Meibom, P., Larsen, H., Barth, R., Brand, H., Weber, C., & Woll, O. (2006). Wilmar Joint Market Model Documentation. *Risø National Laboratory*.
- Pfenninger, S., Hirth, L., Schlecht, I., Schmid, E., Wiese, F., Brown, T., Davis, C., Gidden, M., Heinrichs, H., Heuberger, C., Hilpert, S., Krien, U., Matke, C., Nebel, A., Morrison, R., Müller, B., Pleßmann, G., Reeg, M., Richstein, J. C., ... Wingenbach, C. (2018). Opening the black box of energy modelling: Strategies and lessons learned. *Energy Strategy Reviews*, 19, 63–71.

- Ripper, G. (2000). *Vergleich von Steuermechanismen für den kurzfristigen Einsatz hydraulischer Jahresspeicher* (Ph.D. dissertation). RWTH Aachen.
- Schreck, S. (2016). *Implementation and Analysis of a Rolling Horizon Approach for the Energy System Model REMix* (tech. rep.). Deutsches Zentrum für Luft- und Raumfahrt. Stuttgart.
- Trepper, K. (2015). On the operational and economic effects of ramping restrictions on interconnectors exemplified by the Nord.Link investment. In *K. trepper: Modelling and evaluation of selected market design options for electricity systems with high renewable penetration* (pp. 123–159). Ph.D. dissertation, University of Duisburg-Essen.
- Trepper, K., Bucksteeg, M., & Weber, C. (2015). Market splitting in Germany - New evidence from a three-stage numerical model of Europe. *Energy Policy*, 87, 199–215.
- Ventosa, M., Baíllo, Á., Ramos, A., & Rivier, M. (2005). Electricity market modeling trends. *Energy Policy*, 33, 897–913.
- Weber, C., Meibom, P., Barth, R., & Brand, H. (2009). WILMAR - a stochastic programming tool to analyse the large scale integration of wind energy. In J. Kallrath, P. Pardalos, S. Rebennack, & M. Scheidt (Eds.), *Optimization in the energy industry*. New York, NY, Springer.
- Wiese, F., Bramstoft, R., Koduvere, H., Pizarro Alonso, A., Balyk, O., Kirkerud, J. G., Tveten, Å. G., Bolkesjø, T. F., Münster, M., & Ravn, H. (2018). Balmorel open source energy system model. *Energy Strategy Reviews*, 20, 26–34.

Correspondence

Thomas Kallabis

Research Associate

House of Energy Markets and Finance

University of Duisburg-Essen

Berliner Platz 6-8, 45127 Essen

E-Mail thomas.kallabis@uni-due.de