Optimizing annual-coupled industrial energy systems with sequential time dependencies in a two-stage algorithm

Marion Powilleit*. Stefan Kirschbaum. Joram Wasserfall***. Clemens Felsmann******

* *Society for the Advancement of Applied Computer Science (GFaI), Berlin, Germany (+49-30-814-563-527; powilleit@gfai.de).*

** *Society for the Advancement of Applied Computer Science, Berlin, Germany (kirschbaum@gfai.de)*

*** *Society for the Advancement of Applied Computer Science, Berlin, Germany (wasserfall@gfai.de)*

**** *Institute of Power Engineering, Dresden University of Technology, Dresden, Germany (clemens.felsmann@tu-dresden.de)*

Abstract: Our paper presents a two-stage algorithm designed to address year-round-coupled optimization problems encountered in energy system optimization, particularly relevant for scenarios involving seasonal storages or other conditions depending on annual integrals. We apply this algorithm to MILP and MIQCP models. The solution we propose aims to stay feasible with the original problem while getting close to optimal results. It also significantly reduces the computing time compared to solving the original problem alone. This is crucial because the original problem, when coupled, is very complex and sometimes impossible to solve.

Keywords: Energy system optimization, seasonal storage, year-round coupling, MILP, MIQCP

1. INTRODUCTION

For the best design and operation of energy systems, mathematical optimization methods are a well-established tool to increase efficiency, minimize costs, manage capacity restrictions/availabilities, and reduce ecological impact. In the context of energy systems for, e.g., industrial or municipal energy supply or the manufacturing sector, Mixed-Integer Linear Programming (MILP) represents the state of the art for achieving fast and satisfactory results. However, for individual problems or research purposes, nonlinear algorithms are also commonly applied. Especially the inclusion of bilinear terms can help in dealing with problems of temperature-dependence. It augments the problem to Mixed-Integer-Quadratically-Constrained Programming (MIQCP). For optimal design and operation of heat pump and heat storage systems, previous publications (Hering et al., 2021; Hering et al., 2022) and our earlier work (Powilleit et al., 2019; Wasserfall et al., 2019) worked with this subset of nonlinear optimization models.

Energy system optimization is usually applied to quasi-stationary models and therefore may not suffice for certain problems. One example of this is design optimization, which involves the optimal dimensioning of plants. In this scenario, the entire year needs to be considered as one coupled problem. Due to computational constraints, this can quickly become challenging for moderately complex systems. This challenge is addressed by calculating with aggregated time series (e.g., clustering time points or typical days), where the reduced model still yields highly accurate results (Bahl et al., 2017). More details and ideas about optimal time series aggregation give the publications of Hoffmann et al. (2020; 2022). An important aspect is that the sequence of the time steps is not kept in order with clustering aggregation techniques.

In addition to the question of optimal plant size, other circumstances also necessitate the coupling of all time steps into a single optimization problem. These include:

- optimizing the loading of seasonal storages,
- integrating CO_2 emission limitations over the course of a year, and
- considering specific pricing models or regulatory conditions (such as CHP-remuneration).

The integration of seasonal storages was addressed by Kotzur et al. (2018b) through the inclusion of typical charging and discharging days into their aggregation method. Baumgärtner et al. (2020) proposed an approach which decomposes the original problem into smaller subproblems. Kirschbaum et al. (2023) introduced an adapted rolling horizon technique with integer relaxation.

In this work, we propose a two-stage algorithm similar to methods used for structural optimization. The first stage with downsampling and relaxation methods is simplified while maintaining the sequence of the time steps. The second stage is time-resolved in full detail. The crucial aspects include formulating the resulting boundary conditions and selecting the result variables to be transferred to the second stage.

The second stage has to be feasible to the original problem while closely approaching optimality. Our aim is to develop a seamless algorithm capable of handling a broad range of models across both stages, making it suitable for engineering applications. This entails automating the aggregation, transfer, and formulation of the boundaries in the time-resolved stage. The algorithm should solve the problem faster than the original problem.

2. METHODOLOGY

The methods presented in this work apply to energy system optimization models (ESOM) with the characteristics of being quasi-stationary and equipped with hourly resolved annual demand time series and containing 5 to 15 technical components. Each time step is formulated with a mixed-integer objective function J to minimize operating costs while up to k constraints (e.g., energy balances, conversion terms, operational conditions) can either be mixed integer (MILP, if $C_i = 0$) or in addition bilinear (MIQCP). The continuous optimizing variables x are usually converted power or energy flows. Integer variables y are used to describe the minimum part load or piecewise linear part load behavior.

$$\min J(x, y) \;= a^T \cdot x + b^T \cdot y \tag{1}$$

$$\text{s.t. } x^T \cdot C_i \cdot x + \; d_i^T \cdot x + \; e_i^T \cdot y \leq b_i \tag{2}$$

$$\text{with } x \in \mathbb{R}^n, \;\; y \in \mathbb{Z}^m, \;\; i = 1, \ldots k$$

These equations are often formulated for each time step and solved individually (quasi-stationary). However, when time-coupling conditions are necessary, such as optimizing a daily storage, a **rolling horizon** method is employed (attributable to Bellman (2021)). In this approach, multiple future time steps are linked within a control horizon (e.g., 48 h) to determine the optimal plant behavior. From this control horizon, only the first result is retained, and the process is repeated for the next step then. However, for saving computation time, it is also very common to keep a subset of the control horizon in a shorter write-back horizon (e.g., the first 4 h). The horizon continuously shifts forward throughout the year.

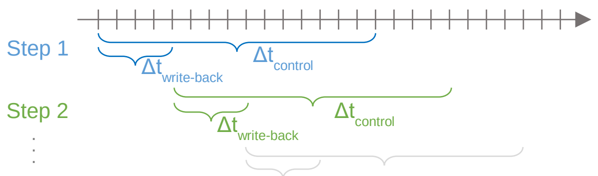

Fig. 1. Illustration of a rolling horizon optimization.

However, some issues, such as the integration of seasonal storage units or regulatory constraints, cannot be addressed by breaking the simulation into smaller parts and prevent solving with conventional receding horizon methods. These require a fully coupled optimization process. This often results in very large and complex optimization problems, which can take a long time to compute or might not even reach a solution. To tackle these cases, we propose our two-stage algorithm.

2.1 Two-stage algorithm: simplified first stage

In the simplified first stage, all time series b_i with the original time step width Δt are downsampled into a coarser grid Δt_j using integral-preserving averaging.

$$b_{i,t_j}^{down} = \frac{1}{\Delta t_j} \sum_{t=\Delta t_j \cdot j}^{\Delta t_j \cdot j + \Delta t_j} b_{i,t} \cdot \Delta t \tag{3}$$

Additional relaxation of the binary variables is possible: The originally binary variables $y \in \mathbb{Z}^n$ are defined continuously as $y \in \mathbb{R}^n$ in the bounds [0,1] to solve only an LP instead of an MILP. These simplifications allow to solve the optimization model of a moderately complex problem in which the whole time frame is coupled in one coupled problem.

2.2 Two-stage algorithm: choosing transfer variables

In an intermediate step of the algorithm, it must be determined which result variables from the first stage are to be retained. These should be the crucial, year-round variables, such as the filling level of a seasonal storage or the amount of CO_2 emitted. These can be brought back to the original time grid through upsampling. The upsampling is optional because the results of the interpolated time steps are not intended to be used further. However, it can still be useful in special cases. The rest of the results are omitted.

From the results, only the transfer variables are selected, which should become boundary conditions for the second stage.

2.3 Two-stage algorithm: Fully-resolved second stage with boundary conditions

In the second stage, the fully-resolved time series and original binary variables are solved with a rolling horizon method. The transfer variables are incorporated as boundary conditions. An intuitive idea is to simply equating the results from each time step from the first stage with the ones from the second, but this leads to infeasibilities and large deviations to optimality. Therefore, we incorporate the boundary conditions into the rolling horizon.

- For integral variables (storage level, amount of CO_2), in each MILP to solve, we only equate the variable x^* to the results from the first stage $x^{*,first\ stage}$ at the last time step of the control horizon:

$$x^*_{t=\Delta t_{control}} = x^{*,first\ stage}_{\Delta t_{control}} \tag{4}$$

- For regular variables (e.g., the electrical power of a CHP), we equate the integral of all variables within the control horizon:

$$\sum_{t=i}^{\Delta t_{control}} x_i^* \cdot \Delta t = \sum_{t=j}^{\Delta t_{control}} x_j^{*,first\ stage} \cdot \Delta t_j \tag{5}$$

- We optionally apply tolerances ε to equation (4) and analogously to (5):

$$x^*_{t=\Delta t_{control}} \geq (1-\varepsilon) \cdot x^{*,first\ stage}_{\Delta t_{control}} \tag{6}$$

$$x^*_{t=\Delta t_{control}} \leq (1+\varepsilon) \cdot x^{*,first\ stage}_{\Delta t_{control}} \tag{7}$$

- We discovered that it is very beneficial to solvability to set the control horizon to a multiple n of the prior downsampling rate Δt_j and write the results back in the same length as the downsampling rate.

$$\Delta t_{control} = \; \Delta t_j \cdot n, \qquad n \in \mathbb{Z} \tag{8}$$

$$\Delta t_{write-back} = \; \Delta t_j \tag{9}$$

This way, the optimality loss accepted in the first stage due to simplification is balanced out to reach approximate optimality. The selection of the values for the downsampling rate, the rolling horizon frame, and the tolerance varied are discussed in Section 3.

2.1 Demonstration of algorithm implementation

Figure 2 shows the results of the algorithm in anticipation of the third case study, the optimal filling level of a seasonal storage. The final optimum filling level of the seasonal storage is marked in green. The grey line is the result of the first downsampled stage, with only a very coarse resolution. The comparison of the two plots shows the progress: In the left one, the red line describes the result of one MILP in the rolling horizon. It starts at t_0. The filling level in the horizon meets the one from the first step (grey line) at $t_{horizon,0}$, which is exactly the defined boundary condition. In the second plot, the consideration starts after the write-back frame at t_4. From there, an MILP with the actualized control horizon is solved and meets the boundary condition again at $t_{horizon\_4}$.

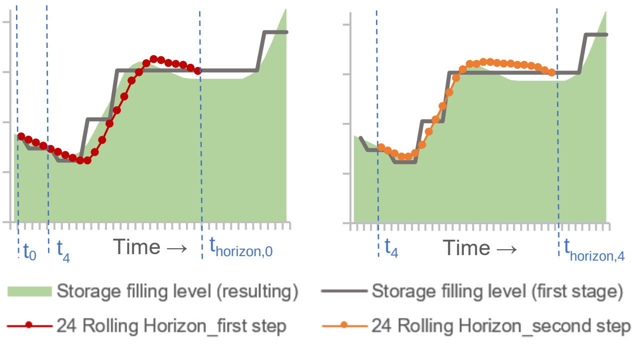

Fig. 2. Exemplary illustration to demonstrate the algorithm: first time step (left), second time step (right).

3. OPTIMIZATION RESULTS

The methodology applies to four different use cases, each of it considering different aspects. For all use cases,

- the operating costs are minimized,
- the demand profiles are extracted and adapted from real use cases or taken from published typical days and have a resolution of 8760 h, and
- energy supply from grid is always depicted with realistic prices.

Because the simulations contain many different combinations of parameters, the nomenclature for the results is standardized as follows.

Table 1: Nomenclature of result diagrams

	The big red dot marks the reference case (if solvable): the complete solution of the annual-coupled fully-resolved problem.
	Colored and shaped group of data points share the same simplified stage model. The legend gives the downsampling rate (“Down 4”) and adds binary relaxation if applied (“Relax”).
48/3 - Tol 0,1	The data labels mark the boundary conditions in the second stage with the control horizon of data to write back in hours (“48/3”) and the applied tolerances ε in % (“Tol 0,1”).
	The red circle marks simulations with an incomplete result set: For some time steps, no solution could be found (within the time limit).

3.1 Software in use

The models are built with the framework *EnergyFrames*, using libraries from the derivate *TOP-Energy*®. Both are proprietary in-house developments by the GFaI. As an optimization solver, Gurobi 11.0.1 is used. While striving for comparability, we calculated with a gap value of 0. However, for some simulations this was not possible in reasonable time. For illustrative purposes, we present computational times in the results, but internally validated them using the dimensionless measurement of Gurobi’s work-units.

3.2 CASE 1: OPERATIONAL OPTIMIZATION DUE TO REGULATORY RESTRICTIONS

Case 1 is a combined heating, cooling, and power system (Fig. 4), as found in small industrial systems. It has certain degrees of freedom: The heating demand of ~1 GWh/a can be met by a CHP (Combined Heat and Power system), a heat pump, or a boiler. The cooling demand of ~0.6 GWh/a can be met by an electric compression chiller or a heat-driven adsorption chiller. The electric grid can either supply power to the plants and other demands of ~1 GWh/a or receive fed-in electricity from the CHP.

The objective function is to minimize the operation costs while complying with one condition: Due to reduction goals of the operator, the whole system should not emit more than 295 t CO_2 per year.

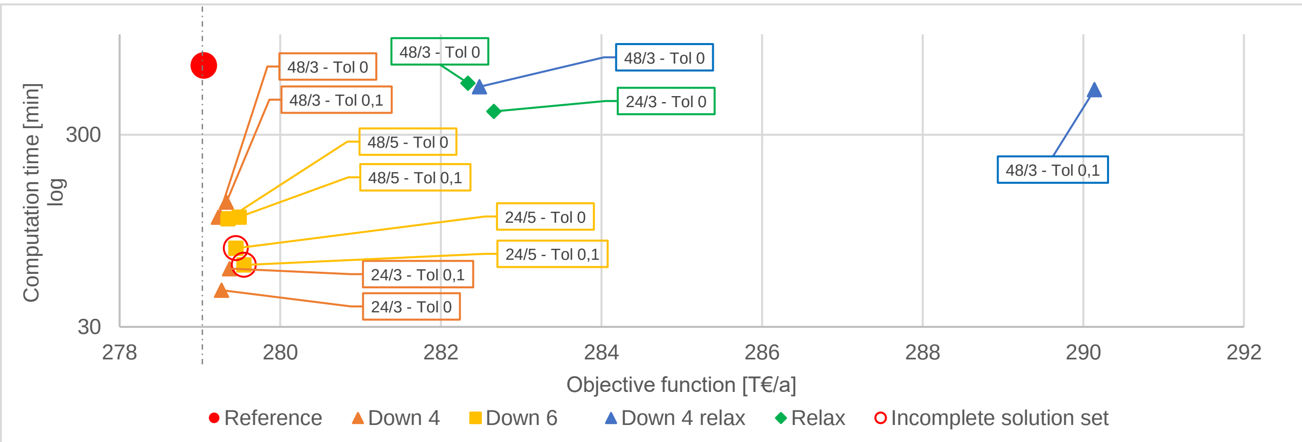

Fig. 3. Computation time vs. objective function results for the two-stage algorithm with varying parameters (Case 1).

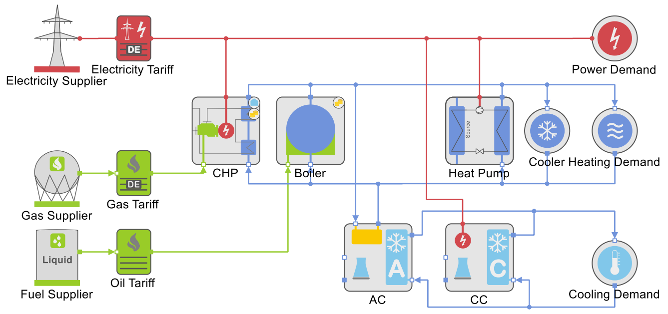

Fig. 4. Scheme of the ESOM with combined heating, cooling and power, with CO_2 emission restriction (Case 1).

Figure 3 shows the results of the two-stage algorithm for different parameters of downsampling rate, tolerances, and horizon frame. The calculation times are compared with each other depending on the objective function values. The aim is to minimize the calculation time with the objective function being as close as possible to the reference solution.

The transfer condition is the integral value of all emissions in each time step. This marks the boundary condition for the second stage.

The main results can be summarized as following.

- Best results can be reached with a downsampling of 4 h and a rolling horizon of 24 h because they require low computation time and still offer a near-optimal objective function.
- The influence of the tolerances in the boundary conditions is unambiguous, but usually, zero tolerance reduces calculation time because no additional degree of freedom is created.
- Binary relaxation of the first step does not show good results (▲ and ◆): Even if the first step solves much faster (~8 s instead of 1670 s), solving the individual steps of the second stage drastically augments the total calculation time and does not even find a feasible solution (○) for each time step.

To understand the reasons for the weak performance of the binary relaxation, we examine the results of the first stage and look at the cumulated CO_2 emissions. The deviation of the relaxed solution to the reference is five times higher than the one from the downsampling case. To disclose that behavior, we take a deeper look into the results of the relaxed binary variables and where they violate the binary condition. We find that the main issue with the violation concerns the operation of the absorption chiller: The minimum part load of 40 % or 160 kW is violated during winter operations with loads around 15 kW (Fig. 6). The resulting boundary condition for the second stage is challenging to meet because unrealistic behavior stems from the outcomes of the first stage.

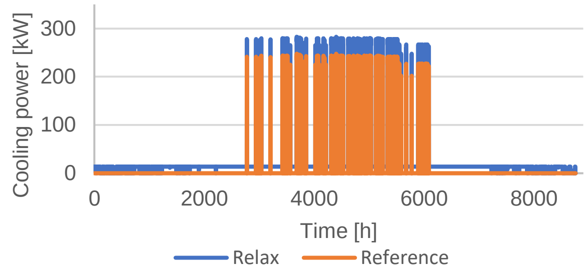

Fig. 6. Comparison of the operation of the AC (Case 1).

3.3 Case 2: Optimal CHP Operation with Annually Restricted Remuneration

In this example (Fig. 7), we delve into a common question concerning the optimal operation of a small-scale CHP system that must effectively meet heat demands of ~1.2 GWh/a while also maximizing electricity sales. Under German law, remuneration is provided for every kilowatt hour of CHP electricity generated, but only for a total of 3,500 full load hours (FLH) per year. Consequently, it is not feasible to optimize each time step individually. Instead, a comprehensive approach spanning the entire year is necessary to determine the most advantageous times for distributing the load of the CHP system.

The CHP is represented with a part load behavior, while a gas boiler and an emergency cooler serve as additional degrees of freedom within the heating grid. Operating costs are minimized as an objective function.

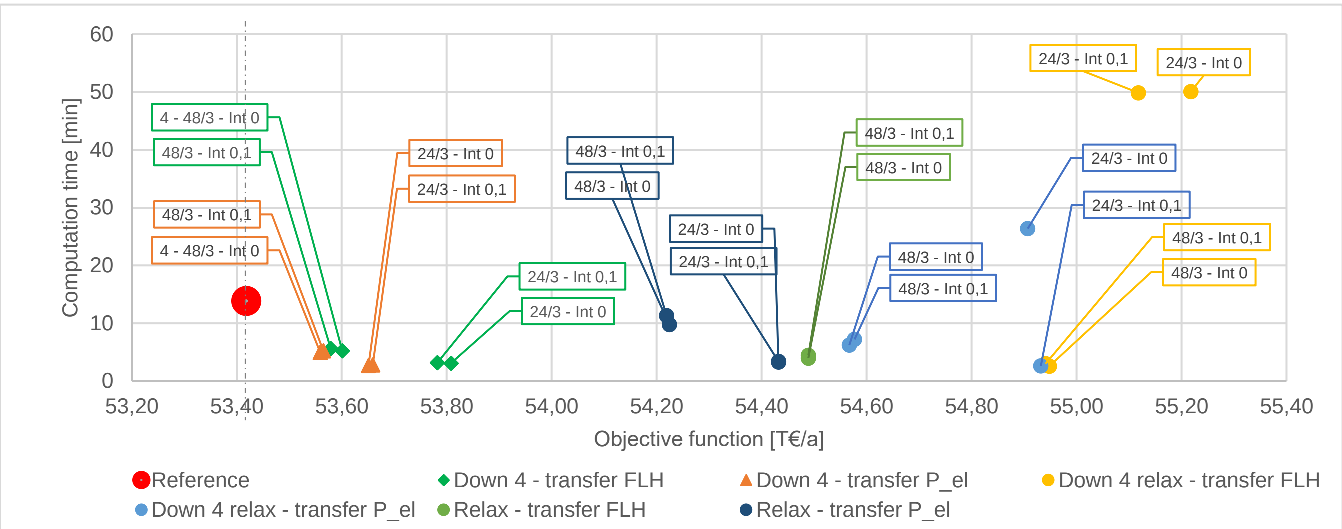

Fig. 5. Computation time vs. objective function results of the two-stage algorithm with varying parameters (Case 2).

The first stage is treated with the downsampling rate of 4 h and optional binary relaxation. In this case, we investigate the quality of two different transfer variables: First, the cumulated FLH (integral variable) is used as a transfer variable for the boundary condition. Unlike the previous case, the time-resolved integral is not necessary because a constraining sum is enough. This allows the remuneration of the FLH to be limited in the first stage and the electricity generation of the CHP to be used as a transfer variable as a second option.

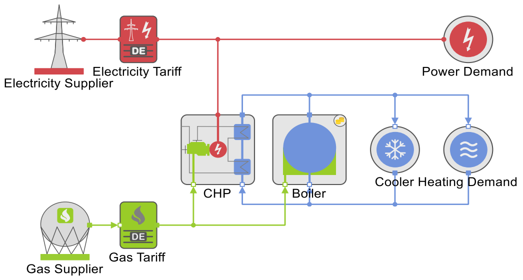

Fig. 7. Scheme of the ESOM with a CHP and annual full load hour restriction of remuneration (Case 2).

The main aspects of the results are:

- As in the previous case, the binary relaxed cases do not improve the total computation time for related reasons as Case 1.
- The other solutions are close to the objective function and reduce the calculation time by around 70 %.
- A larger rolling horizon frame serves a better objective function but increases the calculation time.
- Using the electric power as a transfer variable instead of the cumulated FLH seems to be a better option.
- The tolerance's influence is less clear than in Case 1.

3.4 *CASE 3: OPTIMAL LOADING STRATEGY OF A SEASONAL HYDROGEN STORAGE*

At the core of the fourth use case is a seasonal hydrogen storage system, which serves as a year-round coupling element to increase the use of renewable energy (Fig. 9). During the summer, 1.8 MW_p Photovoltaic (PV) and electrolysis can charge it, while a fuel cell can reconvert the stored energy into electricity to partially cover a demand of ~1.5 MWh/a. A small daily electricity storage system can compensate for short-term fluctuations. In combination with a heat pump, the waste heat from the fuel cell and electrolysis can meet the ~1 GWh/a heat demand, which otherwise would be satisfied by a natural gas boiler.

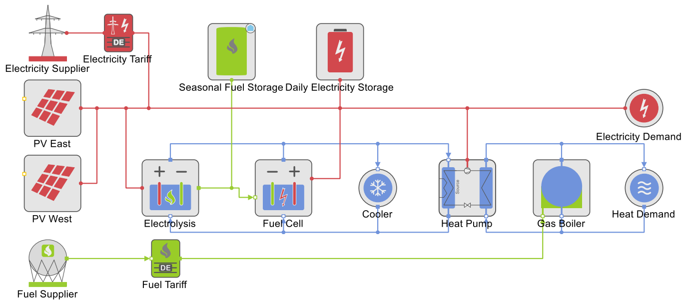

Fig. 9. Scheme of the ESOM with a seasonal storage (Case 3).

As before, the first stage is simplified by downsampling and optional binary relaxation. The transfer condition is initially intuitively the energy level of the seasonal storage (in an integral variable). However, an alternative approach is to define the charging and discharging capacity of the seasonal storage system as a transfer variable. Summarizing the results:

- Good results are achieved by transferring the filling level, employing a downsampling rate of 4 and a rolling horizon of 24 hours. The objective function closely approximates the reference case, while computation times decrease by 20 %.
- Selecting a rolling horizon frame that is too small is problematic because of increased computation times, reduced optimality, and infeasibilities at some time steps.
- Similar to previous cases, the binary relaxation method yields unsatisfying results: With computation time intervals between 200 and 1,800 seconds (not all plotted), the second stage requires too much time for the solution. In addition, it cannot find a solution for

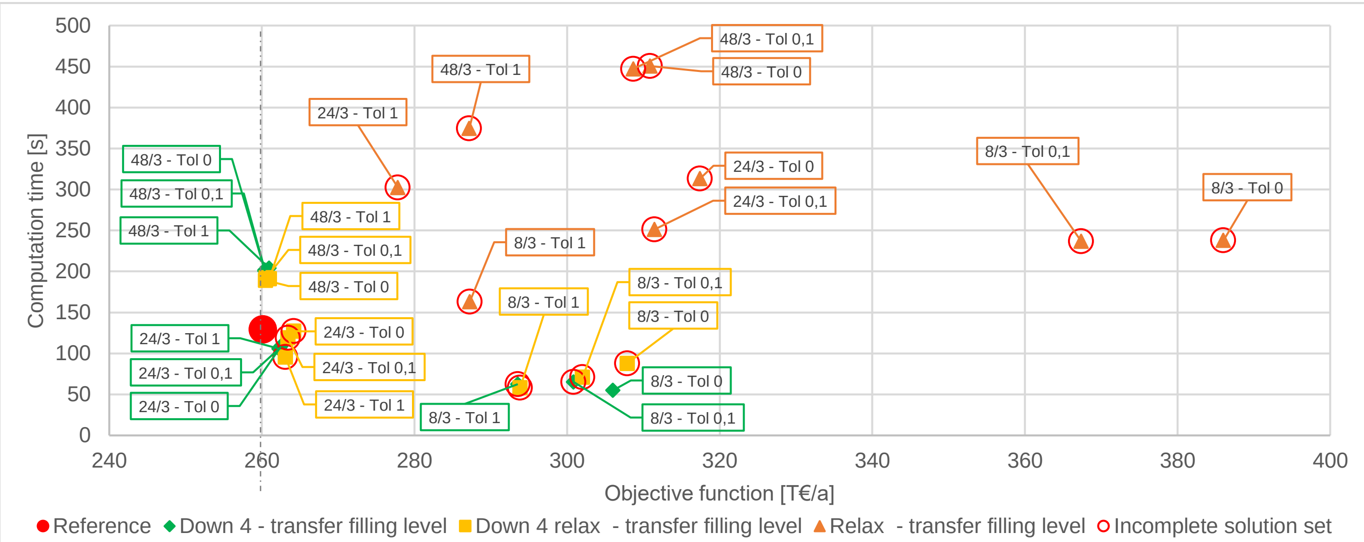

Fig. 8. Computation time vs. objective function results for the two-stage algorithm with varying parameters (Case 3).

all time steps in the maximum time allowed for solving. This leads to an incomplete solution set.

- Coupling the charging and discharging power instead of the integral variable yields highly unfavorable outcomes: Calculation times extend up to 13 hours, consistently resulting in an incomplete solution set. For reasons of clarity, these results are omitted from the plot. It is advisable to refrain from this formulation. Because the integral variable must be computed temporally resolved, no advantage to the previous case study can be expected in the first stage.
- The calculation times of the two-stage algorithm in the downsampling scenario can even exceed those of the reference case (previously observed only when binary relaxation is applied). The solution of individual time steps within a rolling horizon becomes significantly more complex in this scenario. In addition to the coupling condition, the inclusion of a daily storage further enlarges the problem size. Therefore, it is important to carefully choose the horizon frame: To avoid unnecessary complexity, it must not be too large, but must be long enough to effectively operate the daily storage.

3.5 CASE 4: OPTIMAL INTEGRATION OF A SEASONAL HEAT STORAGE TO SUPPLY HEATING OR COOLING COMPRESSION

The fourth model represents a simplified approach to a heating and cooling supply as might be implemented in municipal heat planning (Fig. 11). The core of this model is a 190 MWh seasonal heat storage that stores water at variable temperatures. On the one hand, heat can be extracted as drive heat for an electric heat pump, which raises the temperature level to satisfy the heat demand. As a degree of freedom, a gas-driven boiler can also cover the heat demand. On the other hand, a compression chiller that satisfies a cooling demand can regenerate the storage with its exhaust heat. An additional cooling unit can alternatively fulfill the cooling demand.

The COPs of both plants vary with their inlet temperature, which is reflected as the storage temperature itself. The heat demand is very high in winter, whereas summer season is dominated by cooling demand. The question is how to control the storage temperature throughout the year.

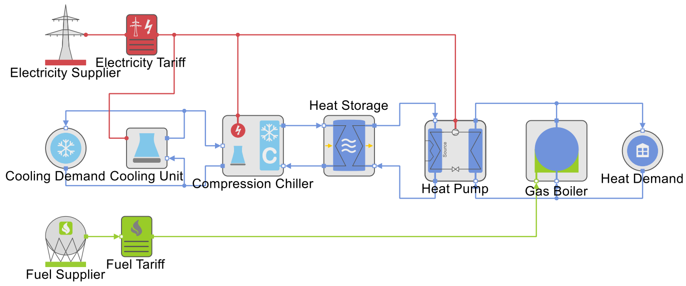

Fig. 11. Scheme of the ESOM with a seasonal heat storage and temperature dependencies (Case 4).

Even if the model is kept simple, the temperature-dependent COP introduces bilinearity into the problem and leads to a non-convex nonlinear optimization problem. That significantly increases its complexity and computational time. The most important bilinear equations are caused by the i enthalpy flows in every energy balance (with c_p being constant) and each of the both COP-dependencies with α being a constant temperature correction factor. Storage mass m and heat capacity c_p are also assumed to be constant.

$$\dot{H}_i = \dot{m}_i \cdot c_p \cdot \Delta T_i \tag{10}$$

$$\Delta U = m_{storage} \cdot c_p \cdot (T_{storage} - T_{ref}) \tag{11}$$

$$\Delta U = \Delta \dot{H}_{charging} + \Delta \dot{H}_{discharging} \tag{12}$$

$$P_{el,j} = COP_j \cdot \Delta \dot{H}_j \tag{13}$$

$$COP_j = COP_{j,nom} + \alpha \cdot (T_{in,j} - T_{in,nom}) \tag{14}$$

In contrary to the other cases, even a coarse downsampling up to 168 h cannot solve the simplified stage to gap zero, but remains at gap values between 5 and 8 %. A quite reasonable gap of around 10 % is reached even in short calculation times, e.g., of 10 minutes.

Even if we could not generate a reference solution from the year-coupled original problem, the algorithm can still compare the different boundary conditions concerning the absolute value of the objective function. Figure 10 gives an overview of the results and the main aspects are summarized:

- With a downsampling of 24 h and a rolling horizon of 48 or 72 h, calculation times of less than 100 min can

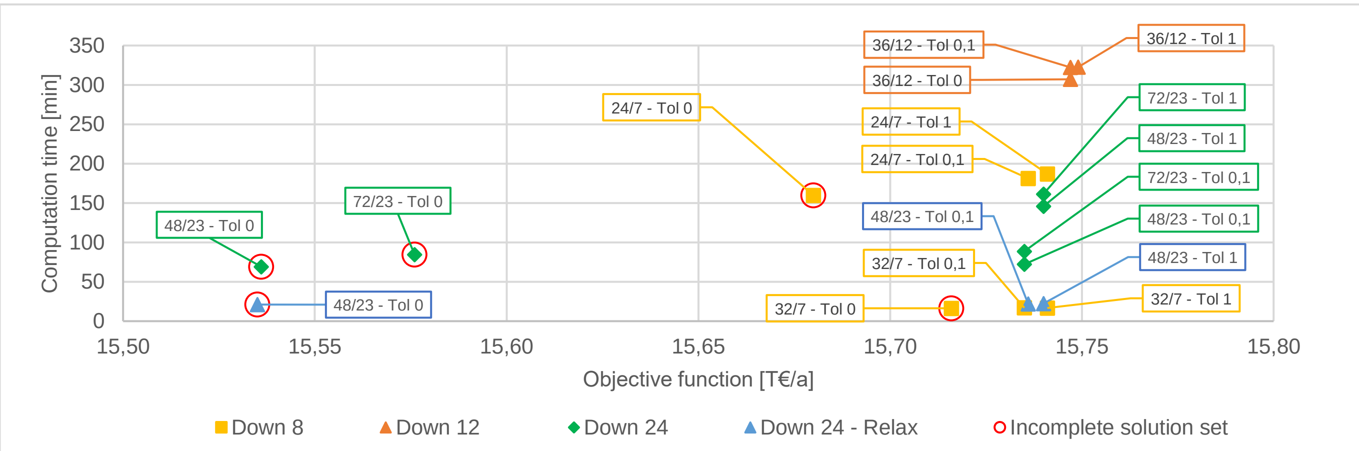

Fig. 10. Computation time vs. objective function results for the two-stage algorithm with varying parameters (Case 4).

be reached to find a feasible solution. The shortest calculation times are reached with a downsampling of 8 h, and a rolling horizon of 32 h. For example, part loads were not modeled in this simple model.

- The objective function values and filling levels do not differ significantly.
- In this case, small tolerances are recommended for the transfer variables in the boundary conditions because most cases with 0 tolerance include time steps for which no solution was found.
- Although in this case the binary relaxation provides advantages, we would not recommend it respecting the experiences of the other use cases.

Even if the model is not very sophisticated (few plants, no heat transfer laws, no part load behavior or runtime conditions), the MIQCP takes much more computation time due to the temperature dependency.

The optimal plant behavior is not only influenced by the energy prices but subsequently also by the temperature-dependent COPs. Three differently efficient heat pumps are compared for illustration (the compression chiller remains unchanged). The algorithm is applied to each heat pump configuration. Figure 12 shows the results of the storage temperatures: The higher the nominal COP, the colder the storage gets in the summer. The highly efficient heat pump can operate economically even with a low driving temperature, whereas the low efficient heat pump requires the storage to keep a higher temperature and hence lowers the full load hours of both compression chiller and heat pump by about 19 % resp. 26 %.

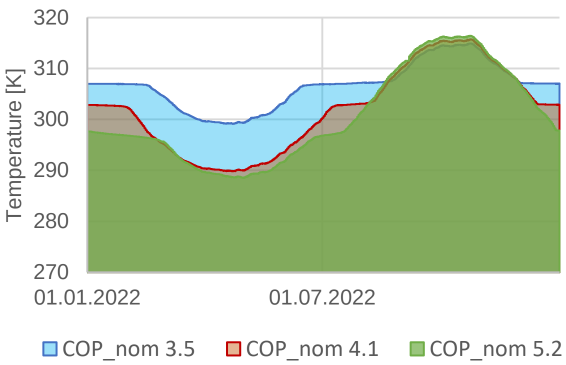

Fig. 12. Influence of nominal COPs on the seasonal storage temperature (Case 4).

4. RESULTS AND DISCUSSION

For all use cases, the two-stage algorithm demonstrates promising results, yielding solutions close to optimal while also achieving significant reductions in computation time (from 15 to 90 %). In instances where the reference solution fails to solve entirely (e.g., due to memory constraints with Gurobi), a feasible solution may still be attained, although without certainty regarding its proximity to the optimum.

Generally, the selection of downsampling rates and horizon windows requires careful consideration: while coarse downsampling accelerates the solution of the first stage, ensuring compliance with boundary conditions in the second stage may require more time. An overview about the achieved computation time saving vs. objective function losses shows Figure 13, where two of the most appropriate parameter settings where chosen. With a too high resolution, the computation time for Case 3 with the seasonal storage in combination with a daily storage may exceed the reference time. Given the absence of a reference for the MIQCP seasonal storage in Case 4, a direct comparison is not possible. However, obtaining a plausible and feasible solution in approximately 16 minutes is a promising outcome. The consistency of the similar storage charging results across different algorithm settings is also promising.

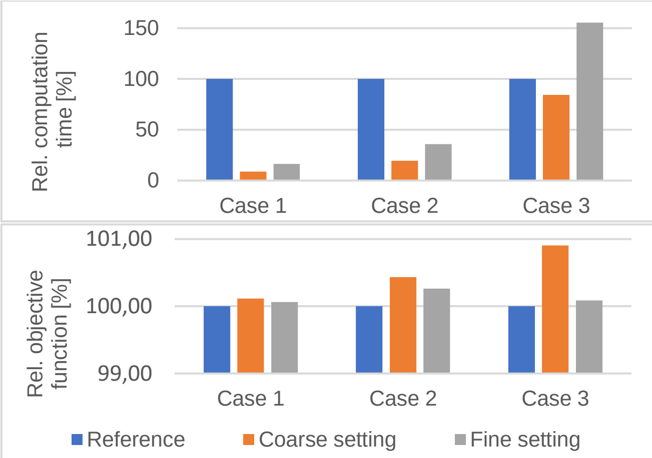

Fig. 13. Comparison of effectiveness of the two-stage algorithm.

Binary relaxation yields unsatisfactory results in this study. While the first stage solves rapidly even without downsampling, the second stage entails long computation times. It is important to note that we did not explore context-specific relaxation methods. There are likely opportunities for improvement by treating specific descriptions for which binary variables are intended differently (e.g., piecewise linear characteristics, minimal part loads or to flow direction decisions), for more details see Özbeg (2022). In cases where downsampling is impossible (e.g., due to critical importance of peak power capabilities or peak power prices), revisiting this method may be warranted.

In this work, a highly capable commercial solver was utilized, allowing even the reference solution to be solved in a reasonable amount of time. However, if using open-source solvers, the parameters found here would need to be adjusted. The pure solution times were taken into account in this study. When utilized within a comprehensive program, other factors such as data handling and LP creation will inevitably come into play. Moreover, it is worth noting the advantage of managing smaller MILP files, particularly in scenarios where memory resources may be limited.

The downsampling process itself can be reconsidered: Currently, equidistant integral averaging is applied, but a more intelligent segmentation of the time series could be implemented, focusing on a higher resolution at important data points while preserving the chronological order, as discussed by Kotzur et al. (2018a). This approach could provide a means to address issues related to peak loads or power pricing better.

5. CONCLUSIONS

Though the models considered here are not extensively complex (to allow a comparison with the reference solution), they nevertheless quickly escalate in complexity with the addition of a few more binary variables or constraints. At this point, year-round coupling can become unfeasible, necessitating the application of multi-stage methods.

Our aim was to develop a highly generic method capable of solving a variety of year-round coupled models within acceptable computation times, yielding feasible solutions. While it is possible to fine-tune each individual model with parameters, we believe that these results allow us to offer a general solution for MILP models in this domain. A coarser and faster solution could be achieved with a downsampling rate of 7 h and a rolling horizon of 28 h, whereas a better solution could be achieved with a downsampling rate of 4 h and a rolling horizon of 24 h. However, expressing this generality for the MIQCP case is more challenging: Here, it is advisable to examine the results of the first stage and assess their plausibility with engineering insight. Nonetheless, by doing so, very good results were achieved in this case as well.

In this work, we only addressed operational optimization problems. The method can easily be adapted to questions combined with design optimization as well when considering the aforementioned issue of peak levelling.

The formulation of MILP models is well-established in research and application, but there is room for further refinement in formulating MIQCPs to enhance their efficiency.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support and funding provided by BMWK, which made the research project "*BiliOpt - Optimierung von Energiesystemen unter Verwendung bilinearer Nebenbedingungen*" (03EI1066B) possible.

Supported by:
Federal Ministry for Economic Affairs and Climate Action
on the basis of a decision by the German Bundestag

REFERENCES

Bahl, B., Kümpel, A., Seele, H., Lampe, M. and Bardow, A. (2017) Time-series aggregation for synthesis problems by bounding error in the objective function. *Energy Energy*, 135, 900–912. doi: 10.1016/j.energy.2017.06.082.

Baumgärtner, N., Shu, D., Bahl, B., Hennen, M., Hollermann, D.E. and Bardow, A. (2020) DeLoop: Decomposition-based Long-term operational optimization of energy systems with time-coupling constraints. *Energy Energy*, 198, 117272. doi: 10.1016/j.energy.2020.117272.

Bellman, R.E. (2021) *Dynamic Programming*. Princeton University Press: Princeton, NJ.

Hering, D., Faller, M.R., Xhonneux, A. and Müller, D. (2022) Operational optimization of a 4th generation district heating network with mixed integer quadratically constrained programming. *Energy Energy*, 250, 123766. doi: 10.1016/j.energy.2022.123766.

Hering, D., Xhonneux, A. and Müller, D. (2021) Design optimization of a heating network with multiple heat pumps using mixed integer quadratically constrained programming. *Energy Energy*, 226, 120384. doi: 10.1016/ j.energy.2021.120384.

Hoffmann, M., Kotzur, L. and Stolten, D. (2022) The Pareto-optimal temporal aggregation of energy system models. *Applied Energy*, 315, 119029. doi: 10.1016/ j.apenergy.2022.119029.

Hoffmann, M., Kotzur, L., Stolten, D. and Robinius, M. (2020) A Review on Time Series Aggregation Methods for Energy System Models. *Energies*, 13(3), 641. doi: 10.3390/en13030641.

Kirschbaum, S., Powilleit, M., Schotte, M. and Özbeg, F. (2023) Efficient Solving of Time-Coupled Energy System MILP Models Using a Problem Specific LP Relaxation. In: ECOS 2023 (Ed.) *36th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems,* 25-30 June, Las Palmas De Gran Canaria, Spain. ECOS 2023, pp. 2774–2785.

Kotzur, L., Markewitz, P., Robinius, M. and Stolten, D. (2018a) Impact of different time series aggregation methods on optimal energy system design. *Renewable Energy*, 117, 474–487. doi: 10.1016/j.renene.2017.10.017.

Kotzur, L., Markewitz, P., Robinius, M. and Stolten, D. (2018b) Time series aggregation for energy system design: Modeling seasonal storage. *Applied Energy*, 213, 123–135. doi: 10.1016/j.apenergy.2018.01.023.

Özbeg, F. (2022) *Entwicklung einer problembezogenen Relaxierung zur effizienten Beschreibung von MILP-Modellen für Energiesysteme*. Bachelorarbeit, Berlin, TU Berlin.

Powilleit, M., Wasserfall, J., Kirschbaum, S., Tretau, A. and Steinmann, D. (2019) Energy system optimization with short term storages respecting nonlinear temperature dependencies. In: Stanek, W. (Ed.) *32nd International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2019),* 23-28 June, Wrocław, Poland. Silesian University of Technology, Department of Thermal Engineering.

Wasserfall, J., Powilleit, M., Kirschbaum, S. and Wrobel, G. (2019) Efficient formulation for optimizing temperature dependent energy systems with mixed-integer quadratically constrained programming. In: Stanek, W. (Ed.) *32nd International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2019),* 23-28 June, Wrocław, Poland. Silesian University of Technology, Department of Thermal Engineering, pp. 953–965.