



LUKAS MEHL
NIELS LINDNER¹
KAROLINA BARTOSZUK
JANINA ZITTEL²

Prototypical warm-starts for demand-robust LP-based energy system optimization

¹  0000-0002-8337-4387

²  0000-0002-0731-0314

Zuse Institute Berlin
Takustr. 7
14195 Berlin
Germany

Telephone: +49 30 84185-0
Telefax: +49 30 84185-125

E-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Prototypical warm-starts for demand-robust LP-based energy system optimization

Lukas Mehl,
Niels Lindner,
Karolina Bartoszuk,
Janina Zittel

Abstract The expressiveness of energy system optimization models (ESOMs) depends on a multitude of exogenous parameters. For example, sound estimates of the future energy demand are essential to enable qualified decisions on long-term investments. However, the enormous demand fluctuations even on a fine-grained scale diminish the computational performance of large-scale ESOMs. We therefore propose a clustering-and-decomposition method for linear programming based ESOMs that first identifies and solves prototypical demand scenarios with the dual simplex algorithm, and then composes dual optimal prototype bases to a warm-start basis for the full model. We evaluate the feasibility and computational efficiency our approach on a real-world case study, using a sector-coupled ESOM with hourly resolution for the Berlin-Brandenburg area in Germany, based on the oemof framework.

1 Introduction

The aim of an energy system optimization model (ESOM) is to find an optimal design of an energy system determined by investment and operational decisions. As, for example, the availability of renewable energy sources, the content of storages, and the energy demand heavily fluctuate over time [8], it is common to consider time-dependent ESOMs [2]. The longer the planning horizon, and the finer the resolution, the larger are the resulting models, and the higher is the computational effort for solving. Not only does this apply to deterministic perfect-foresight scenarios, but also to models that capture uncertainty by investigating a catalog of demand scenarios, where all scenarios can in

Lukas Mehl, Karolina Bartoszuk, Janina Zittel
Zuse Institute Berlin, Takustr. 7, 14195 Berlin, Germany,
e-mail: mehl@zib.de, e-mail: bartoszuk@zib.de, e-mail: zittel@zib.de

Niels Lindner
Freie Universität Berlin, Institut für Mathematik, Arnimallee 3, 14195 Berlin, Germany,
e-mail: lindner@zib.de

principle be combined to a single, but extremely large ESOM in a stochastic or robust setting (see, e.g., [9, 7]). However, for such models, solving merely a single time step typically requires negligible effort. This motivates a decomposition approach to solve each individual timestep, or longer time slices independently, and then to combine the solutions to a warm-start, i.e., an initial solution for the full ESOM. In this paper, we will explore this idea in the case when the ESOM is based on linear programming.

2 Method

We first set the formal preliminaries for our method in Section 2.1, before explaining the key details in Section 2.2 and Section 2.3.

2.1 Formalization and Preliminaries

We will consider linear programs (LP) of the form

$$\text{Minimize} \quad c^\top x + \sum_{t \in T} c_t^\top y_t \quad (1)$$

$$\text{subject to} \quad Ax + \sum_{t \in T} A_t y_t \leq a, \quad (2)$$

$$Bx \leq b, \quad (3)$$

$$B_t y_t \leq b_t, \quad t \in T. \quad (4)$$

for matrices A, B, A_t, B_t and vectors a, b, c, b_t, c_t , where t is an element of a set $T = \{1, \dots, T_{\max}\}$ of time steps. The LP (1)-(4) has *time-independent* variables x and *time-dependent* variables y_t , restricted by time-independent (3) and time-dependent constraints (4), respectively. The variables x and y_t for $t \in T$ are linked by *coupling* constraints (2).

In the context of energy system optimization, the time-independent variables x are *investment* variables, e.g., the capacity that should be installed for a specific technology in a given area. These variables represent decisions that are valid for the whole time horizon. The energy flows between different components of an energy system, e.g., from a primary source to a storage component, and the contents of storages, e.g., batteries, make up the time-dependent variables y_t . The right-hand sides b_t – assuming perfect foresight – stem from time series data for the demand of the different fuels. The other coefficients in the LP are governed by the possible network structure, and assumptions on the efficiencies of technologies, as well as on investment and operational costs. We want to illustrate two examples of coupling constraints (2): First, the installed capacity, e.g., of a primary energy source, may be modeled as the maximum outflow over all time steps. Second, the content of a storage at time $t + 1$ depends on the content at time t .

A linear program can be solved in numerous ways [4]. For large-scale energy systems, interior point methods do usually perform best, but do not admit a powerful warm-starting method that can exploit knowledge about a feasible good-quality solution. The situation is different for the dual simplex method: The core of this algorithm is to maintain a *basis*, i.e., an invertible submatrix of maximal rank of the coefficient matrix of the LP. In more detail, the constraints are transformed into equations by introducing non-negative *slack variables*, with the exception of those constraints that specify lower and upper bounds for the variables. In this terminology, a basis is given by stating which variables and constraints (by means of slack variables) are part of the basis (*basic*), and which are not (*non-basic* at lower or upper bound), provided that the corresponding submatrix is regular. The dual simplex procedure then iteratively explores basis exchanges to finally terminate with an optimal dual basis. The advantage is that a dual basis remains feasible when right-hand sides b and b_l are altered, and can therefore be used again as initial basis for a modified LP. If the modifications are minor, the warm-start basis is likely to be close to an optimal solution, so that the dual simplex will terminate within few iterations.

Speaking in terms of ESOMs, the dual simplex method hence allows for efficient warm-starts, when energy demands are only slightly changed.

2.2 Clustering and Prototypes

To find out when warm-starts could be reasonably applied, we propose a clustering procedure, whose goal is to obtain few small sub-models which are representative for the full model. The first step is to divide the time series data into slices of consecutive time steps. These slices all have the same length l which is specified by an exogenous parameter. Then, we partition these slices into groups using a clustering method. We opted for the standard k -means clustering, where the number of clusters k is also specified by an exogenous parameter. Now, for each cluster, we set the representative slice of this cluster to be the medoid minimizing the Euclidean distance among all slices from the cluster. In the case where the full time horizon of the model is not a multiple of the slice length l , some of the last time steps are not part of any slice. Those remaining time steps are then simply assigned to the cluster which contains the last slice, i.e., the slice closest to the end of the full time horizon.

To sum up, each time step $t \in T$ is now assigned to a cluster and to the representative slice of this cluster. Our idea is to solve the model only for the representative slices and then take advantage of the information obtained from these solutions by making use of the assignment from time steps to representatives. To this end, we create a *prototype* for each cluster, a model that is equivalent to the original model, but covers only the time horizon of a representative slice. Then, we solve each prototype model independently and store the optimal prototype basis computed by the dual simplex algorithm.

2.3 Composed Basis

The basis which shall be used for the warm-start is a combination of the optimal bases of the prototype models. In the following, we describe how we create a *composed basis* from these prototype bases.

For a time-dependent variable y_t , and a time-dependent or coupling constraint G_t of (1)-(4), consider the cluster to which the time step t belongs and its representative. We want to derive the basis status of y_t or G_t from a corresponding timestep in the optimal basis of this representative. To this end, we calculate the relative time index of t w.r.t. the slice length l as $t' := t \bmod l$. Then, we set the basis status of y_t or G_t in the composed basis to the basis status of $y_{t'}$ or $G_{t'}$ in B , respectively.

For an investment variable x , we want a warm-start basis in which the maximum value of this variable among all optimal solutions of prototypes can be attained. Therefore, if the variable is non-basic at its upper bound in at least one of the prototype bases, we set it to non-basic at its upper bound in the composed basis. Else, if the variable is basic in at least one of the prototype bases, we set it to basic in the composed basis. An investment variable is only set to non-basic at its lower bound in the composed basis, if the variable is non-basic at its lower bound in all of the prototype bases.

For the time-independent constraints (3), e.g., emission limits, we apply a majority voting procedure. That is, we take the basis status which occurs most often among the prototype bases. For some constraints, this voting turns out to be unanimous, i.e., these constraints have the same basis status in all prototype bases.

For ESOMs, we also consider a strategy where we ignore the storages in the prototype models, i.e., they are being removed from these models. In this setting, no hints for the basis statuses of storage-related variables and constraints are available and thus general rules are needed. We choose to set all variables which correspond to a storage content and all variables which correspond to an investment in storages to basic. The motivation here is to allow for flexible usage of the storages. All equality-constraints related to storages are set to non-basic as their slacks should be zero in any feasible solution. On the other hand, all inequality-constraints related to storages are set to basic as their slacks can potentially be non-zero.

A basis composed along these instructions might not have the required size, i.e., the number of basic variables (including slack variables) might not be equal to the number of constraints of the full model. To address this issue we adopt a naive approach by adding or removing arbitrary variables to or from the basis, until the required size is obtained: When storages are ignored in the prototype models, we repair the basis by removing arbitrary (non-slack) variables from the basis, as it appears that the size of the basis is always too large for this storage strategy. When we include storages, we always observe that we have too few basic variables. In this case, we add arbitrary slack variables to the basis.

3 Experiments and Results

We evaluate our clustering and decomposition strategy on the *oemof-B3* model ([6, 5]) based on the open energy modeling framework *oemof* ([1]). The model follows a greenfield approach to optimize the sector-coupled energy system of the Berlin-Brandenburg area in Germany, hence considering both a metropolitan area, its suburban surroundings, and rural parts. On the technical side, the model is a linear program, covering a planning horizon of one year with a resolution of one hour.

We partition these 8760 time steps into different slices of constant length from 1 day onward, and then use k -means clustering to assign these slices to a given number k of clusters. As a preprocessing step, we compute an optimal basis for the cluster representatives. When doing so, we can choose between two storage strategies: Either we include the storage variables in the LP or remove them from the model, i.e., completely ignoring storages. Then, we assemble a basis for the original model as described in Section 2.3, and use the dual simplex implementation of Gurobi 10.0.3 [3], warm-starting from this composed basis. All computations have been carried out on a high-throughput computing cluster with Intel Xeon Gold 6342 CPUs.

The results are summarized in Table 1. The preprocessing times are negligible in comparison to the time required to solve the warm-started model. While our approach is viable as such, it is inferior to solving the full linear program, both in terms of runtime and simplex iterations. In general, including storages performs better than ignoring the storages, but still leading to an average increase of a factor of roughly 2.5 in dual simplex iterations compared to solving the original model from scratch. Concerning different slice lengths and numbers of clusters, the best results being obtained with slice lengths of 48 hours. Especially when storage constraints are ignored, there is no clear picture how to choose these two parameters.

4 Conclusion

We have presented a method to solve large-scale energy system optimization models formulated as linear programs by a technique relying on clustering, decomposition, and composing a warm-start basis. Evaluating on the *oemof-B3* model, we confirm that the method works, but is clearly outperformed by solving these linear programs without any decomposition from scratch. We identify two reasons: First, the deviations in demand seem to be too numerous to allow for meaningful warm-starts. Second, given that storages unfold their full potential only over longer time horizons, estimating their behavior remains a sweet spot. As seen in our results, including storage constraints has a clear and positive impact on computation times. In future work, the challenge is therefore to find a good handle on predicting storage usage. If this succeeds, then our method is in principle able to be applied to larger models that accumulate several demand scenarios in one, e.g., in the context of robust or stochastic programming [7].

storage strategy	slice length [h]	# clusters	preprocessing time [s]	runtime [s]	iterations
ignore	24	365	3.70	40220	2379290
ignore	24	182	1.84	62648	3910671
ignore	24	91	0.92	34181	2478864
ignore	24	12	0.14	27255	1440323
ignore	48	182	3.66	62204	1933474
ignore	48	91	2.01	77906	2808932
ignore	48	12	0.24	38029	1492415
ignore	168	52	4.55	78232	2076934
ignore	168	13	1.19	52892	2006363
ignore	720	12	10.04	60267	2641596
ignore	720	4	3.14	75313	2856891
ignore	8760	1	52.30	47063	2460672
include	24	365	26.72	23734	1190337
include	24	182	13.55	26783	1157102
include	24	91	6.88	152558	2628719
include	24	12	0.87	27060	1321364
include	48	182	23.87	17349	1249808
include	48	91	12.33	16874	1161215
include	48	12	1.57	16657	1141430
include	168	52	103.22	24367	1202173
include	168	13	27.59	22113	1037014
include	720	12	469.74	22637	1234910
include	720	4	157.53	19491	1234269
include	8760	1	*—	*5289	*520998

Table 1 Computational results for our method. *The last line is the baseline approach with the full undecomposed original model, requires no preprocessing, and is not warm-started.

Acknowledgements The work for this article has been conducted in the Research Campus MODAL funded by the German Federal Ministry of Research, Technology and Space (BMFTR) (fund number 05M20ZBM), and in the project “Stadt-Land-Energie - Robustheit und Übertragbarkeit von interkommunalen Energiewendeszenarien im Stadt-Land-Nexus” (Stadt-Land-Energie) funded by Forschungszentrum Jülich GmbH with funds of the Ministry for Economic Affairs and Energy (BMWE) (fund number 03EI1051A-B).

References

1. oemof community: Open energy modelling framework – a modular open source framework to model energy supply systems. URL <https://github.com/oemof>
2. DeCarolis, J., Daly, H., Dodds, P., Keppo, I., Li, F., McDowall, W., Pye, S., Strachan, N., Trutnevyte, E., Usher, W., Winning, M., Yeh, S., Zeyringer, M.: Formalizing best practice for energy system optimization modelling. *Applied Energy* **194**, 184–198 (2017). DOI 10.1016/j.apenergy.2017.03.001. URL <https://www.sciencedirect.com/science/article/pii/S0306261917302192>
3. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2025)
4. Klotz, E., Newman, A.M.: Practical guidelines for solving difficult linear programs. *Surveys in Operations Research and Management Science* **18**(1), 1–17 (2013). DOI 10.1016/j.sorms.2012.11.001. URL <https://www.sciencedirect.com/science/article/pii/S1876735412000189>
5. Launer, J., et al.: oemof-B3 (2023). URL <https://github.com/r1-institut/oemof-B3>
6. Launer, J., Haas, S., Gering, M.C.: oemof-B3 Raw Data (2023). DOI 10.5281/zenodo.8364204

7. Lindner, N., Mehl, L., Bartoszek, K., Berendes, S., Zittel, J.: Demand uncertainty in energy systems: Scenario catalogs vs. integrated robust optimization. In: Proceedings of 38th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (2025)
8. Löffler, K., Burandt, T., Hainsch, K., Oei, P.Y., Seehaus, F., Wejda, F.: Chances and barriers for germany's low carbon transition - quantifying uncertainties in key influential factors. *Energy* **239**, 121901 (2022). DOI <https://doi.org/10.1016/j.energy.2021.121901>. URL <https://www.sciencedirect.com/science/article/pii/S0360544221021496>
9. Yue, X., Pye, S., DeCarolis, J., Li, F.G., Rogan, F., Gallachóir, B.O.: A review of approaches to uncertainty assessment in energy system optimization models. *Energy Strategy Reviews* **21**, 204–217 (2018). DOI <https://doi.org/10.1016/j.esr.2018.06.003>. URL <https://www.sciencedirect.com/science/article/pii/S2211467X18300543>