The value of space-time load-shifting flexibility for 24/7 carbon-free electricity procurement

Iegor Riepin, Tom Brown
Department of Digital Transformation in Energy Systems, TU Berlin
26 July 2023

Preface

- This study is done in a spirit of open and reproducible research: GitHub.
- **Funding:** This study was supported by a grant from Google, LLC.
- **Acknowledgements:** The authors thank members of the Google energy and climate teams for their feedback and inputs on earlier drafts of this study. We also thank the **PyPSA team** and many contributors to the open-source energy system modelling ecosystem used for this study (see: **github.com/PyPSA**). Warm thanks to Fabian Hofmann for making complex optimization simpler with **linopy**.
-
- The content of this study, including any errors or omissions, are the responsibility of the authors alone.

Executive summary: context and scope of the study

- This work is a follow-up to an earlier study written by the authors in October 2022 on the **"System-level impacts of 24/7 carbon-free electricity procurement in Europe"**.

- In this study, we explore *how and why* space-time load-shifting flexibility can be used to meet high **24/7 carbon-free energy targets**, as well as what potential benefits it may offer to 24/7 participants and to the rest of the energy system.

- To answer these questions, we expand the mathematical model developed in the previous work by incorporating **spatial and temporal demand flexibility** provided by electricity consumers that follow 24/7 carbon-free energy goals. The space-time flexibility is based on the example of *data centers*; however, the findings of this study are generally applicable to a wide range of companies with flexible demand.

- We model the European power system (**ENTSO-E area**) clustered to **37 zones**. The model **co-optimizes** investment and dispatch decisions of **locally procured** generation & storage assets to meet electricity demand of data centers (the 24/7 CFE participants), as well as investment and dispatch decisions of assets in the rest of the European electricity system to meet the demand of other consumers. Furthermore, depending on a level of flexiblity available, data centers could benefit from co-optimizing load shifting (across space and/or time) and procurement strategies to match every kWh of electiricty consumption with carbon-free energy around-the-clock more efficiently. We place data centers in a selection of European countries: Ireland, Denmark, Germany, Finland, and Portugal. All model runs are done for **2025** with **hourly resolution**.

Executive summary: key findings

1. Demand flexibility enables **better access to clean electricity** and creates **more options** for consumers to match demand with carbon-free electricity around-the-clock.
2. Some flexible electricity consumers, such as data centers, can shift computing jobs and associated power loads in both time and location. These mechanisms facilitate the **efficiency and affordability** of 24/7 CFE procurement. The co-optimized space-time load-shifting can reduce the costs of 24/7 CFE by up to 34%, depending on the level of flexibility and technologies available.
3. Demand flexibility is **especially helpful for resource-constrained locations** where hourly matching with 24/7 CFE is difficult.
4. Space-time load-shifting facilitates **economically efficient redistribution of loads** to locations with good carbon-free resources. When paired with long-duration energy storage, the efficiency gains of this effect are even larger.
5. In the European energy system, the hourly profiles of wind power generation have a low correlation over long distances due to different weather conditions. Spatial load flexibility enables the system to move load to locations when and where there is high wind generation, thus **saving costs of energy storage** and **reducing curtailment** of excess generation.

Table of Contents

Introduction

Introduction

- Climate change is driving a global effort to **rapidly decarbonise** electricity systems across the globe. Many public and private energy buyers join this effort. For example, more than 400 members of the **RE100 group** have committed to procure enough renewable energy to match 100% of their electricity consumption on an annual basis.

- Fully decarbonizing electricity grids, however, requires covering demand with carbon-free energy at all times, not just during periods of abundant sunshine or wind. This challenge requires embracing **innovative strategies** for decarbonization. There is growing interest from leaders in voluntary clean electricity procurement to cover their consumption with carbon-free energy supply on a **truly 24/7 basis**. Achieving 24/7 Carbon-Free Energy (CFE) means that every kilowatt-hour of electricity consumption is met with local carbon-free electricity sources around-the-clock.

- The **24/7 Carbon-Free Energy Compact**, coordinated by the United Nations now includes more than 120 signatories on a mission to realize a 24/7 Carbon-Free Energy future.

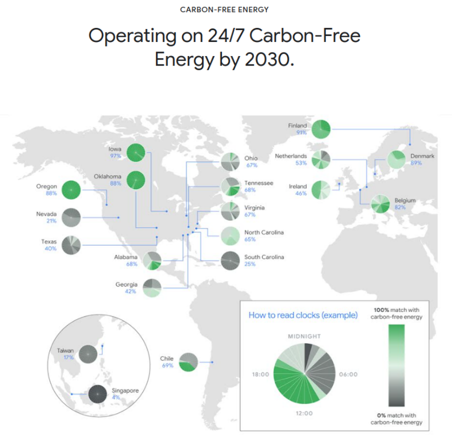

Image: sustainability.google/progress/energy/

Introduction

- In October 2022, we published a study on the
  "System-level impacts of 24/7 carbon-free electricity procurement in Europe"
  Code behind the study.

- In the study, we investigated the **means and costs** of pursuing different clean electricity procurement strategies for companies in a selection of European countries. We also explored how the 24/7 CFE commitments **affect the European electricity system** as a whole.

- The study concluded with the following take-aways:
  (i) 24/7 CFE commitments lead to lower emissions for both the participants and the system;
  (ii) 24/7 CFE also reduces the needs for flexibility in the rest of the system;
  (iii) Reaching CFE for 90-95% of the time can be done with only a small cost premium. Reaching 100% CFE target is possible but costly with existing renewable and storage technologies, with costs increasing rapidly above 95%. 100% CFE target could have a much smaller cost premium if long duration storage or clean firm generation technologies are available.
  (iv) 24/7 CFE procurement stimulates innovation and learning, and creates an early market for the advanced technologies.

- These European study results align with the results in studies done by **Princeton ZERO lab (2021)** for regions in the United States and by **IEA (2022)** for India and Indonesia.

Motivations

- In the previous study, we focused on a large range of European companies from the commercial and industry (C&I) sectors that join 24/7 CFE efforts in aggregate. The implicit assumption we made was that all 24/7 CFE participants have **inflexible demand**.
- In reality, many participants of the 24/7 CFE movement have some degree of flexibility in their electricity consumption. This flexibility takes the form of various mechamisms for temporal demand management available for **a wide range** of C&I consumers.
- A **large potential for demand side flexibility** is available in the information and communications technology (ICT) sector. Big companies such as Amazon, Google, IBM, and Microsoft are centralizing data centers to achieve economies of scale and form a computing infrastructure that is managed collectively via network operation centers. Thus, data center operators have the ability to **shift computing jobs and associated power loads** in time (via scheduling of flexible compute jobs) and in space (via migration of flexible compute jobs across locations).

Why is this important? 1/2

Technische Universität Berlin

- Demand for digital services is rapidly growing. Since 2010, internet traffic has expanded **25-fold**. Global data center energy use represents now nearly **1% of final electricity demand** worldwide. Data centres and data transmission networks are responsible for **0.9% of energy-related GHG emissions** (around 300 Mt CO_2-eq in 2020).
- Despite rapidly growing demand for digital services, the growth of associated emissions was modest due to energy efficiency improvements, decarbonisation of electricity grids and renewable energy purchases by ICT companies above and beyond the policy obligations. Based on **IEA (2022)** estimates, Amazon, Microsoft, Meta and Google have become the four largest purchasers of corporate renewable energy, having contracted over 38 GW to date with power purchase agreements (PPAs).
- Moreover, some of the ICT companies have become the front runners of the 24/7 CFE movement. Google has committed to the goal of **24/7 Carbon-Free Energy by 2030**. Similarly, Microsoft has announced its own **100/100/0 by 2030** commitment.

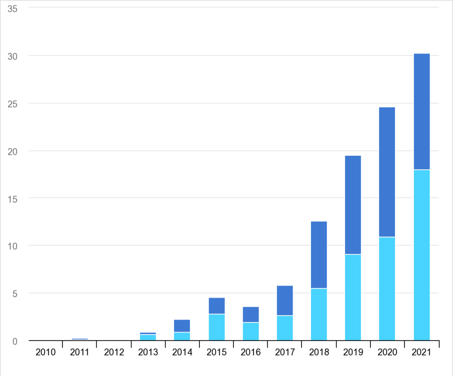

Renewable energy capacity procured with power purchase agreements globally [GW].

ICT sector (dark blue), all other sectors (light blue)

Image: IEA 2022

Why is this important? 2/2

Data centre operators (Pact Associations) that signed the Climate Neutral Data Centre Pact

- The initiatives to measure and reduce the environmental impacts of digital infrastructure is spanning far beyond big companies like Google and Microsoft.
- In 2021, over 100 data data centre operators and industry associations in Europe signed [the Climate Neutral Data Centre Pact]() aiming to make data centres climate neutral by 2030. The pledged targets include measures to increase power usage effectiveness and carbon-free energy supply. The CFE target is declared to be "[..] 75% of renewable energy or hourly carbon-free energy by December 31, 2025 and 100% by December 31, 2030."
- Considering (i) a constant growth of global internet traffic, (ii) a large electricity consumption of data centers distributed in power grids worldwide, and (iii) the need to rapidly decarbonise electricity systems across the globe, it is **important to understand the possible efficiency benefits that space-time load shifting flexibility can provide for the 24/7 carbon-free energy paradigm**.

Image: climateneutraldatacentre.net/signatories/

A growing body of research

- The unique characteristics of data centers as electricity consumers and the active interest of ICT sector companies in sustainable energy drive a growing interest in the research community. Among many other, **Wang et al. (2015)**, **Toosi et al. (2017)**, **Grange et al. (2018)**, **Velasco et al. (2018)**, and **He & Shen (2021)** investigated selected aspects of spatial or temporal demand management strategies in the context of supplying data centers power demand with intermittent renewable energy supply.

- **Zhang & Zavala (2022)** elaborated a mathematical problem that captures both spatial & temporal load-shifting flexibility provided by data centers. The authors suggest market clearing formulation treats data centers as prosumers that simultaneously request load and provide a load-shifting flexibility service to the grid. The illustrated clearing formulation satisfies fundamental economic properties of the competitive markets, such as revenue adequacy and cost recovery.

- The Google research team published a paper on **Carbon-Aware Computing for Datacenters** (Radovanović et al. (2023)). The paper introduced methodology and principles behind a carbon-intelligent compute management system, which minimizes electricity-based carbon footprint and power infrastructure costs by shifting temporally flexible workloads for all datacenter clusters across Google's fleet.

Focus of the study

Technische Universität Berlin

- In this study, we explore *how and why* space-time load-shifting flexibility can be used to meet high 24/7 carbon-free energy targets, as well as what potential benefits it may offer to 24/7 participants and to the rest of the energy system. We aim to answer the following questions:
  – How can demand flexibility reduce the **resources** and **costs** for 24/7 CFE matching?
  – How can spatial and temporal demand flexibilities be utilized to achieve high 24/7 CFE goals?
  – What are the **individual effects** of spatial and temporal demand flexibility, as well what are the synergies from their co-optimization?
  – How would advanced technologies, such as long duration storage, affect **the value of demand flexibility**?

- For this purpose, we elaborate the mathematical model developed in the previous study, by including spatial and temporal demand flexibility provided by electricity consumers following 24/7 CFE goals. Thus, a flexible 24/7 participant could benefit from co-optimizing utilization of available demand flexibility (across space and/or time) and procurement strategies to match every kWh of electricity consumption with carbon-free energy around-the-clock **more resource-efficient**.

- The modelling exercise in this study is based on the example of *data centers*, i.e., facilities used to house networked computer servers that store, process and distribute large amounts of data. Nevertheless, the findings of this study are generally applicable to a wide range of companies and organisations with flexible demand and an interest in 24/7 carbon-free energy procurement, as well as to energy industry experts and stakeholders with an empirical interest in the European energy system.

Study design

A quick overview

- This study is done in a spirit of open and reproducible research. The whole scientific workflow from the publicly available raw input data to optimized electricity system, visualizations and compilation of this study is available at **github.com/PyPSA/247-cfe**.

- In this study, we build upon the mathematical model of 24/7 CFE procurement developed in the former work of authors: **System-level impacts of 24/7 carbon-free electricity procurement in Europe** (October 2022)

- We encode a set of new equations and routines, which allow for modelling **spatial** (computing jobs migration) and **temporal** (computing jobs scheduling) load flexibility provided by data centers. The mathematical model of temporal flexibility generalizes a broad range of flexible C&I consumers.

- We place data centers (i.e., electricity consumers committed to 24/7 CFE goals) in a selection of European countries: Ireland, Denmark, Germany, Finland, and Portugal. These countries have different weather patterns, renewable potentials, national energy and climate policies, legacy fleets of generation capacities, degree of interconnectons, etc. Apart from that, we consider several scenarios for CFE procurement targets, degrees of data center flexibility, and technologies available for 24/7 consumers. These differences help to **understand and generalize** the interplay of demand flexibility and 24/7 CFE procurement.

Study design: European power system

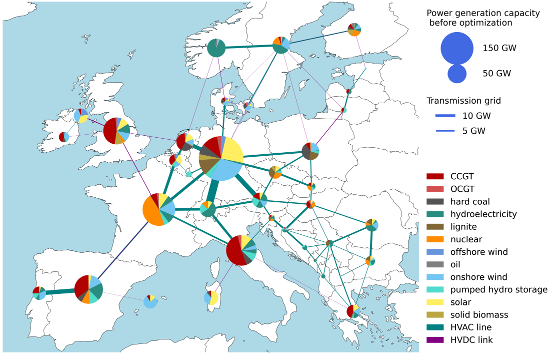

European electricity system clustered to 37 zones
NB power generation capacity fleet before optimization

- In each scenario, we model the full European power system (**ENTSO-E area**) clustered to **37 zones**. Each zone represents an individual country. Some countries that straddle different synchronous areas are split to individual bidding zones, such as DK1 (West) and DK2 (East).

- The model **co-optimizes** investment and dispatch decisions of generation & storage assets to meet electricity demand of data centers (flexible 24/7 CFE consumers), as well as investment and dispatch decisions of assets in the rest of the European electricity system to meet the demand of other consumers.

- The modelling is done for **2025**. Input data such as technology cost assumptions, national renewable policies, decommissioning of legacy power plant fleet, and system-wide assumptions (e.g., price for EU ETS) are parametrised accordingly.

- All model runs are done with **hourly resolution**, i.e., no time sampling.

Study design: data centers

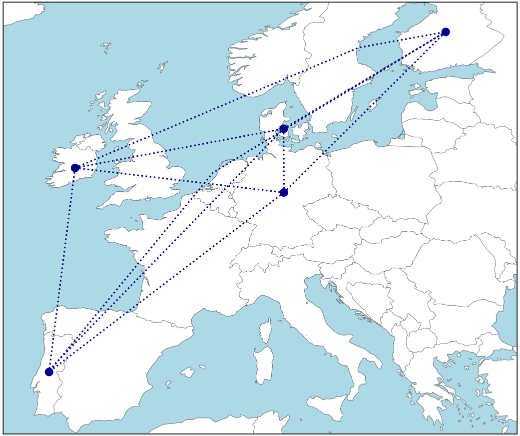

Five data centers interconnected by virtual links, forming a complete graph

- We consider **five data centers** that are located in Ireland, Denmark (West/DK1), Germany, Finland, and Portugal. These locations (i) include zones where data centers have an important share in national electricity demand [1,2,3], and (ii) include zones that have electricity systems with unique characteristics, such as local generation mix, renewable potentials, national energy and climate policies, degree of interconnections, etc.
- Data centers have a nominal load of **100 MW** (baseload profile). The data center operator aims to achieve a given 24/7 CFE matching score **at all locations**.
- Load shifts take place via "virtual links". Virtual links form a **complete graph**, i.e., every pair of data centers is connected by a unique link.
- Data centers have the **same share of flexible workloads**.

Study design: data center load flexibility

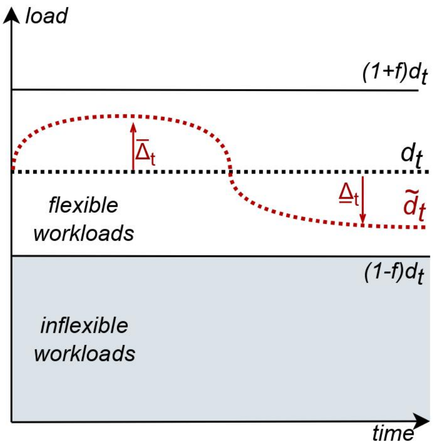

- The premise of data center flexibility is that a known number of computing jobs, and associated power usage is "flexible", i.e., electricity loads can potentially be shifted in space (across datacenter locations), or to other times (by delaying jobs' execution).[a]

- Thus, the **dispatched load** $\widetilde{d}_t$ of a data center can deviate from the nominal requested load d_t. The dispatched load $\widetilde{d}_t$ is constrained by the data center capacity (an upper limit) and the inflexible loads (a lower limit). The range of possible deviations of the dispatched and nominal loads is assumed to lie within f [%] of the nominal load, such as:

$$[1 - f] \cdot d_t \leq \widetilde{d}_t \leq [1 + f] \cdot d_t \quad \forall t \in T \qquad (1a)$$

$$\widetilde{d}_t = d_t + (\overline{\Delta}_t - \underline{\Delta}_t) \quad \forall t \in T \qquad (1b)$$

where $\overline{\Delta}_t, \underline{\Delta}_t \in \mathbb{R}_+$ stand for positive/negative deviation of $\widetilde{d}_t$ and d_t in hour t.

[a] A change in cluster-level CPU usage can be accurately mapped into a change in its power usage, see Radovanovic et al. (2021)

Study design: space-time load shifting

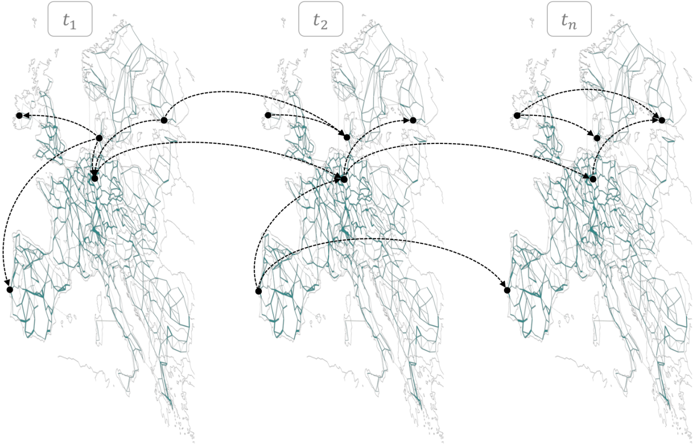

- To capture spatial flexibility, we model a spatial load management system that can shift load across data center locations via virtual links.
- To capture temporal flexibility, we model a load scheduling system that can shift load of a data center over time.
- The spatial and temporal load shifting are subject to a shared set of computing capacity constraints; the temporal shifting is further constrained by the daily compute usage conservation rule.
- For the mathematics and detail on the optimization model, see **Annex A: Methodology**.

Study design: scenario space

- We model various scenarios for data center demand flexibility, which include

  Three modes of operation:

  – Co-optimized **spatial and temporal** load management

  – Isolated **spatial** load management (shifting flexible loads across locations)

  – Isolated **temporal** load management (shifting flexible loads in time)

  and four scenarios for flexible loads range: $f = \{\mathbf{0\%, 10\%, 20\%, 40\%}\}$.

- Two scenarios for 24/7 CFE hourly matching targets: $CFEscores = \{\mathbf{98\%\ and\ 100\%}\}$.

- Further, we assume two palettes of carbon-free technologies available for procurement for data center operators participating in 24/7-CFE:

  – **Palette 1** includes technologies available on the European market now: onshore wind, utility scale solar PV, battery storage.

  – **Palette 2** includes all above plus Long Duration Energy Storage (LDES) system.

  Technology assumptions are provided in **Annex B: Tools and data sources**.

- For interested parties, we publish an [online Annex] with a full pack of modelling results alongside this study. The materials include modelling results for all scenario combinations in a form of plots and summary CSV files.

Study design: limitations

Technische Universität Berlin

- This study is done in a spirit of **modelling for insight** rather than **modelling for numbers**. The design of this study does not aim at quantifying the real-life benefits of demand side flexibility for data centers. It is rather a model experiment to explore *how and why* flexibility of demand can be beneficial for achieving 24/7 carbon-free energy goals. The results we present should thus be viewed with a fair degree of caution, i.e., as a modelling-based insight rather than quantitative projection.

- Quantifying the actual costs and benefits for the ICT industry of utilizing demand flexibility requires **additional empirical research**. Further studies could usefully explore the costs and technical potentials of achieving a certain share of flexible workloads, which are **not considered** in this study. Thus, a range of flexible loads is fixed per scenario and flexibility utilization (i.e., shifting of loads in space and time) is modeled as a "zero-cost" variable. Including information on implicit flexibility costs would help to quantify the flexibility benefits with a greater degree of accuracy. Another empirical improvement could address technical aspects and properties of flexible workloads, such as physical constraints associated with quick ramping of power usage up/down, reliability & performance constraints, etc. Further research is needed to capture the promising role of clean firm generation technologies in achieving 24/7 CFE goals with some degree of spatial and/or temporal load flexibility. This case is particularly relevant, since the authors' **previous modeling work** demonstrated that 24/7 CFE procurement could create an early market and drive deployment of advanced technologies, such as LDES and clean firm generation.

Modelling results and analysis

An overview

The results section is organized as follows:

1. Procurement and 24/7 CFE costs as a function of load flexibility.
2. Economic efficiency of co-optimized space-time load shifts.
3. Isolating values of spatial and temporal load shifting.
4. Insights from time-series data for optimized space-time load shifts.
5. Economically efficient redistribution of data center average loads.
6. Further remarks.

Procurement as a function of load flexibility (100% CFE)

Technische Universität Berlin

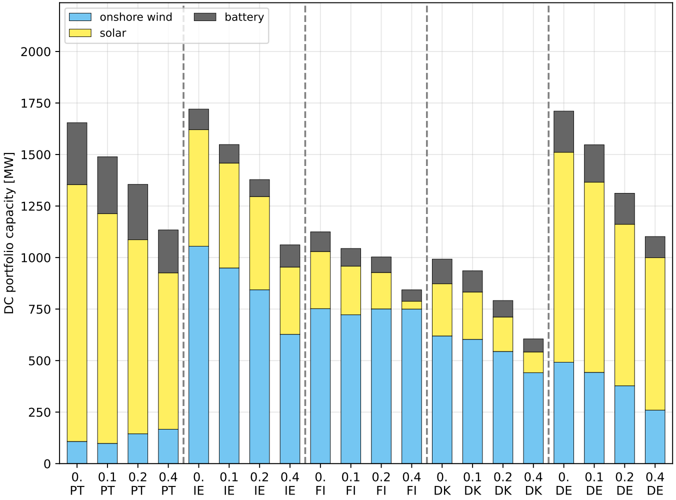

Let us first consider the following scenario: (i) 100% CFE target, (ii) technology palette 1 (without LDES), and (iii) co-optimized spatial & temporal load flexibility.

A plot on the left shows the cost-optimal **portfolio capacity** required to match demand with carbon-free electricity around-the-clock. Results are displayed per each location and share of flexible loads f.

As shown in the previous study, 100% hourly matching with renewable generators and battery storage requires a large portfolio for inflexible demand case. The cost-optimal mix of solar PV, wind and battery storage depends on the local resources.

The required portfolio capacity **is significantly reduced when load shifting becomes possible**. This effect takes place in all locations.

24/7 CFE costs as a function of load flexibility (100% CFE)

Technische Universität Berlin

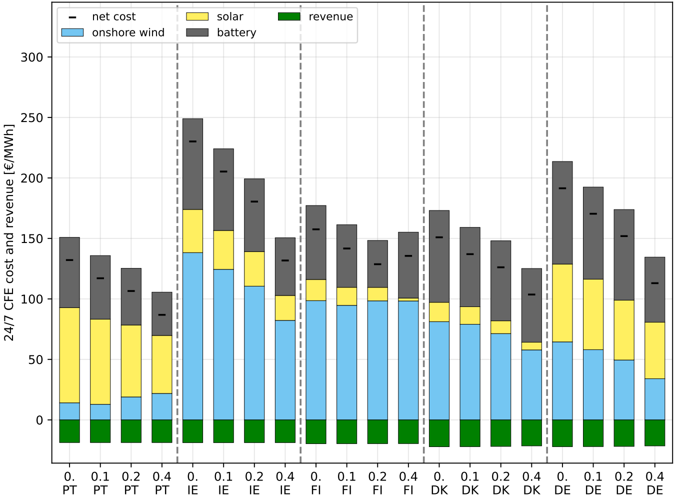

The **cost breakdown** on the left shows the average costs (per MWh of consumption) of meeting demand with the 24/7 CFE policy netted by revenue sold to the regional grid. (NB for the 100% CFE target, 24/7 consumers have nearly no grid imports in the consumption mix.) With inflexible demand, only selected regions benefit from good resources for solar (PT) or wind (DK) and achieve hourly CFE matching with lower costs. Overall, the 100% CFE hourly matching target remains costly with palette 1 technologies.

Load shifting reduces the costs for 24/7 procurement in all locations, and especially in locations where hourly matching with CFE is expensive (IE, DE). Thus, demand flexibility enables achieving 24/7 CFE in a more cost-effective way, and this effect is particularly notable for **the resource-constrained places**.

Procurement as a function of load flexibility (98% CFE)

Technische Universität Berlin

If we now turn to a scenario with **lower CFE target of 98 %** (keeping technology palette 1 and both spatial & temporal load flexibility enabled), we see that this procurement policy can be met by procuring much less onshore wind and solar capacity.

This observation is in line with the previous study, which showed that the last 2% of hourly CFE matching nearly doubles the required resources and costs (without LDES or clean firm generation).

With a CFE target of 98%, the total procured capacity still reduces in all locations with increasing potential for demand flexibility, i.e., increasing share of flexible workloads. The effects include (i) a reduction of battery storage, (ii) a small reduction of overall renewable capacity, and (iii) a swap of solar capacity with wind due to higher capacity factor of the latter.

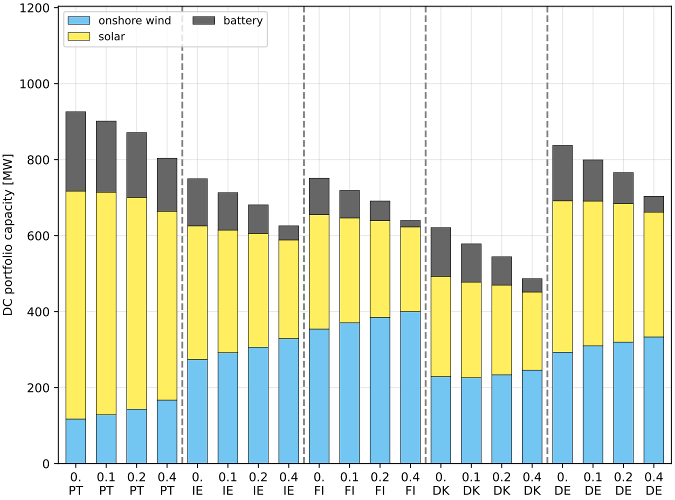

24/7 CFE costs as a function of load flexibility (98% CFE)

There are two distinct observations in the breakdown of costs for meeting the CFE 98% policy: (i) the cost component associated with imports of electricity from the regional grid enters the mix of options to meet the CFE policy; (ii) as mentioned above, the net average cost of CFE procurement is much lower than for CFE 100%.

Shifting of load across space and time **enables access to clean electricity** at times of day when certain locations have high renewable penetration and **creates more options to match demand with CFE** for times and locations where renewable potential is scarce (for details, we provide a time-series analysis below).

As a result, CFE 98% policy is more affordable in all locations with increasing demand flexibility.

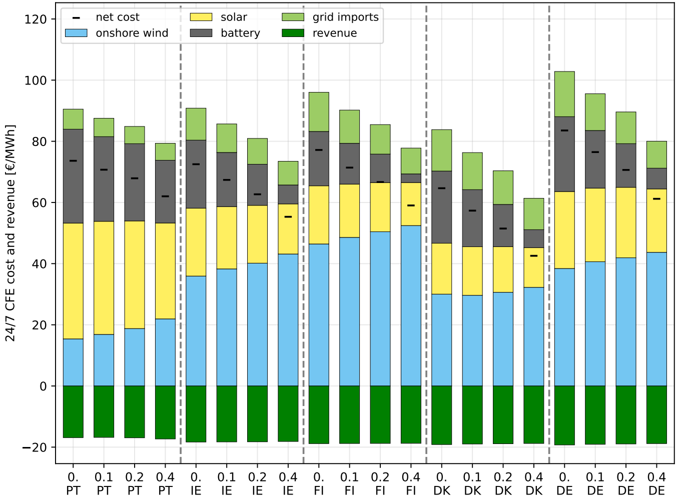

Procurement as a function of load flexibility (100% CFE w/ LDES)

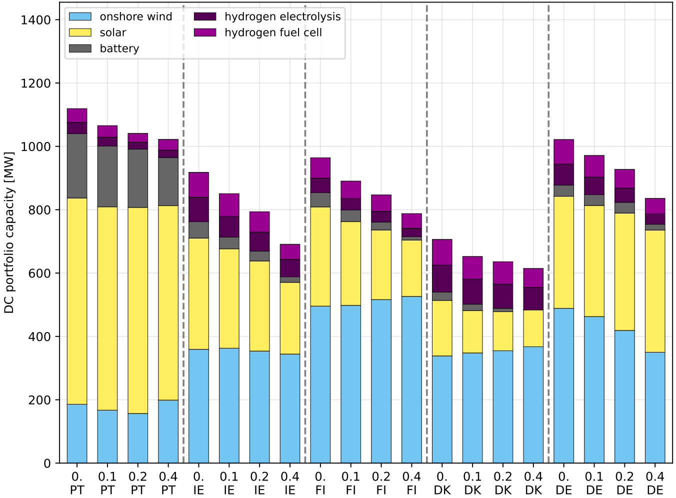

Let's now look at the results for the **technological palette 2 (with LDES)**, keeping the CFE 100% target and both spatial & temporal load flexibility enabled.

When 24/7 consumers have access to a LDES system, the required portfolio of renewable capacity for the 100% CFE target is significantly reduced. The LDES system helps to align the load with the generation of procured variable renewable resources.

Co-optimization of demand flexibility with LDES **promotes further efficiency gains**. LDES paired with load shifting makes it possible to smooths out variations of renewable generation and achieve 24/7 hourly matching with even fewer resources. Though, the absolute values of capacity reduction with higher flexibility are expectedly lower than w/o LDES.

24/7 CFE costs as a function of load flexibility (100% CFE w/ LDES)

Technische Universität Berlin

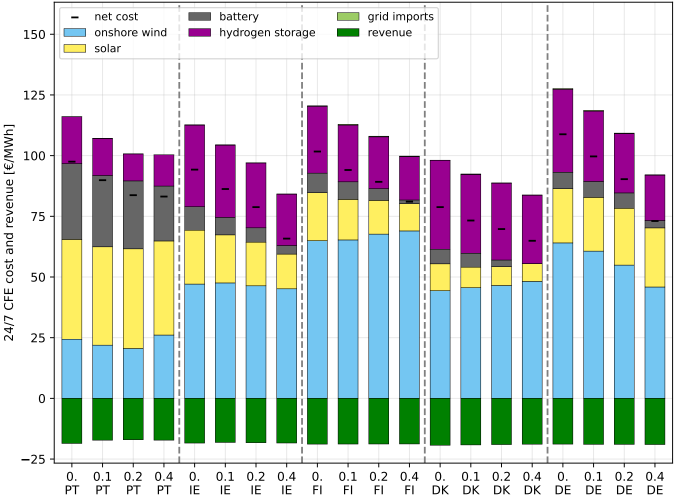

These efficiency gains translate to further 24/7 costs reduction.

For the inflexible demand case, a LDES system (with costs of 2.5 €/kWh) helps to bridge hours with no renewable feed-in and reduce the costs compared to the technology palette 1 case.

Once load shifting is possible, co-optimization of spatial shifting, temporal shifting, and LDES helps a 24/7 consumer to match its demand with carbon-free electricity around-the-clock at lower costs.

Co-optimization of load shifting and LDES eliminates the need for battery storage for the 100% CFE target in wind-dominant locations (DK); however, for solar-dominant locations (PT), some share of batteries is still cost-optimal, even with high degree of flexibility.

Economic efficiency of co-optimized space-time load shifts (100% CFE)

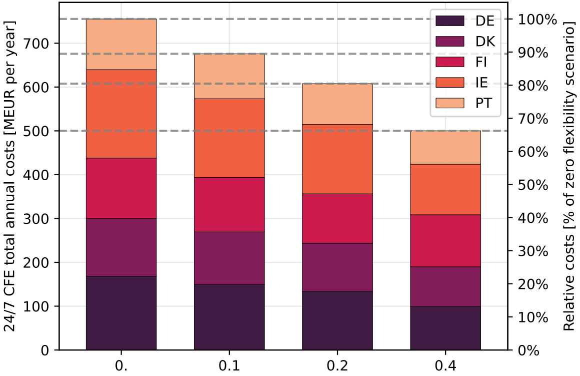

Share of flexible workloads.
Costs reduction in max flexibility scenario: 34% (255.3 MEUR/a)

A plot on the left shows the **total annual costs [€/a]** for achieving 24/7 CFE policy in all locations (left y-axis) and their relative representation as a percentage of the zero flexibility scenario's costs (right y-axis). The costs are plotted as a function of the load flexibility potential. The values represent the total procurement costs, not a 24/7 "premium" (the additional costs to the price of electricity in a local market).

The plot gives a summary perspective on the observations above: increasing potential of demand flexibility facilitates the **efficiency and affordability** of 24/7 CFE procurement. If 10% of loads are flexible, the total costs of achieving 100% CFE decreased by 11%. The costs decrease even further as flexibility increases.

The cost reduction is proportionally higher in the resource-constrained locations (IE, DE) where hourly matching with CFE is more expensive. This suggests that demand flexibility is especially helpful for locations **where 24/7 approach is difficult**.

Isolating values of spatial and temporal load shifting

The results above are for the case when data centers **co-optimize** shifting of the flexible loads across locations and over time. It is also interesting to look at the economic efficiency when flexibility usage is isolated, i.e., data centers implement either a spatial or a temporal load management system. The plots below show total annual costs for achieving 24/7 policy with isolated **spatial (left)** and **temporal (right)** load shifting.

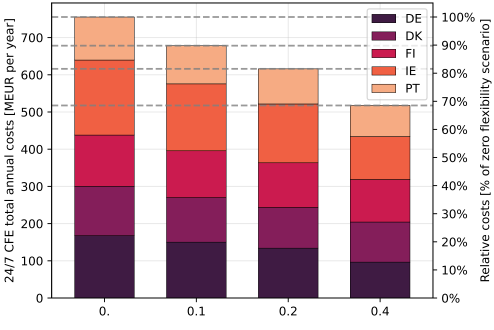

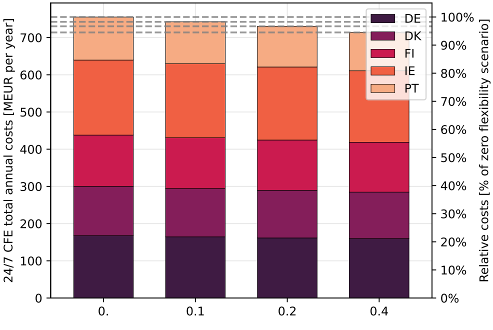

Isolating values of spatial and temporal load shifts (100% CFE)

The modelling results for isolated spatial and temporal load shifting shown above reveal the following:

- The estimated value of spatial load management (for this scenario) is nearly **six times bigger** (31% and 6% of cost reductions, accordingly).
- When implemented together, the space-time load shifting can yield higher overall economic efficiency gains (34%); however, **the effects do not add up** because the spatial and temporal load shifts are subject to a shared set of computing capacity constraints (eq. 13a-13c).
- When analysing time-series data below, we show four individual channels for cost savings attributed to both spatial and temporal load management (two channels for each). The relatively low value of temporal flexibility can be explained by the fact that the two cost saving channels of temporal load shifting are limited in this scenario, while the two cost saving channels of spatial load shifting are actively utilized.[1]
- **Supplementary graphics** in the Annex reveal that spatial load management enables reduction of both locally procured generation and battery storage capacity, while temporal load management mainly reduces the needs for battery storage.

[1] For temporal shifts: (i) the variability of grid emission intensity does not play a role because data centers have to rely on locally procured resources at 100% CFE score; (ii) the daily compute usage conservation rule (eq. 10) limits the ability of using temporal shifts to reduce the capacity fleet of wind and solar PV. For spatial shifts: (i) shifting workloads across locations enables taking advantage of difference in weather conditions and (ii) taking advantage of differences in local resources.

Isolating values of spatial and temporal load shifts (98% CFE)

Let's switch our focus to the scenario with **98% CFE score and technology palette 1**. Similarly, the plots below show results for isolated spatial (left) and temporal (right) load shifting. A notable difference is that the value of spatial flexibility is much smaller (18% of cost savings), while the value of temporal flexibility is larger (14%). At the 98% CFE score, 24/7 consumer complements own portfolio of procured CFE technologies with imports of electricity from the regional grid (as shown above). Thus, temporal shifts are also **responsive to local grid's carbon intensity**, which drives its rising value. Shifting loads across locations still delivers a larger chunk of efficiency gains, but the absolute gains drop with a smaller portfolio of the CFE resources and the option to occasionally rely on grid imports. NB Data for the co-optimized case is in **the supplementary graphics**.

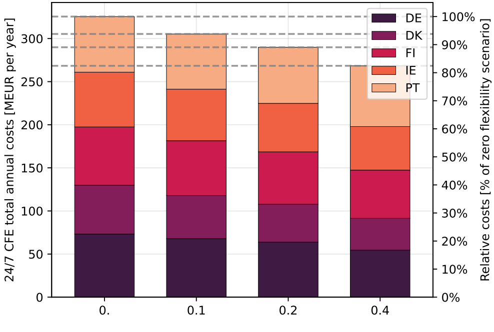

Share of flexible workloads.
Costs reduction in max flexibility scenario: 18% (57.1 MEUR/a)

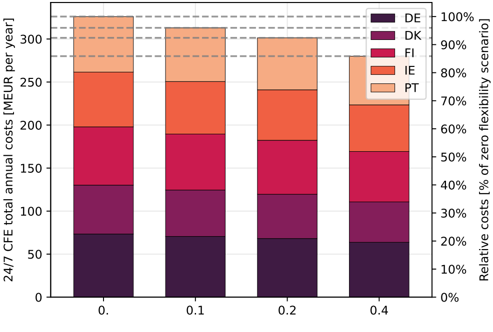

Share of flexible workloads.
Costs reduction in max flexibility scenario: 14% (46.0 MEUR/a)

Isolating values of spatial and temporal load shifts (100% CFE w/ LDES)

Finally, the same perspective for the scenario with **100% CFE score and palette 2 (with LDES)**. Once a LDES system is added to the technology mix of the 24/7 consumer, it can store excess generation from variable renewable resources for extended periods. It helps matching demand with CFE around-the-clock with considerably fewer resources everywhere, including the resource-constrained locations. Co-optimization of spatial load management and LDES brings additional synergies. Long duration storage helps harvesting renewable electricity in the best locations and spatial flexibility *indirectly* opens access to the cheaper clean electricity for all locations. We discuss this effect in more detail in the analysis of **redistribution of average data center loads**.

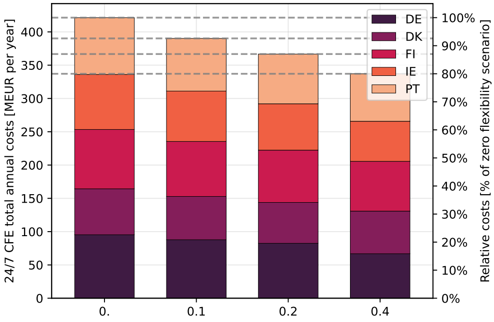

Share of flexible workloads.
Costs reduction in max flexibility scenario: 20% (84.3 MEUR/a)

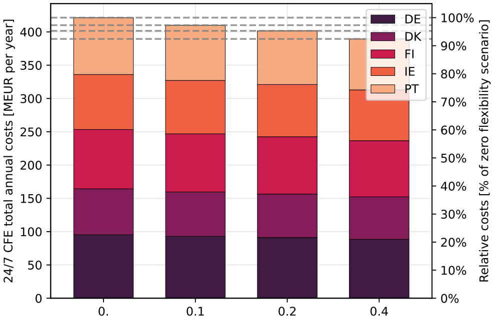

Share of flexible workloads.
Costs reduction in max flexibility scenario: 8% (32.0 MEUR/a)

Insights from time-series data for optimized space-time load shifts

The time-series data for space-time load shifts reveals several distinct patterns in the load-shifting. Let's take a look at the selected scenario (see bottom right for details). The plot below shows **the hourly spatial load shifts** for a data center in Ireland. **Negative** values mapped to blue color represent loads "received" from other locations in a given hour, while **positive** values mapped to red color represent loads "sent" away.

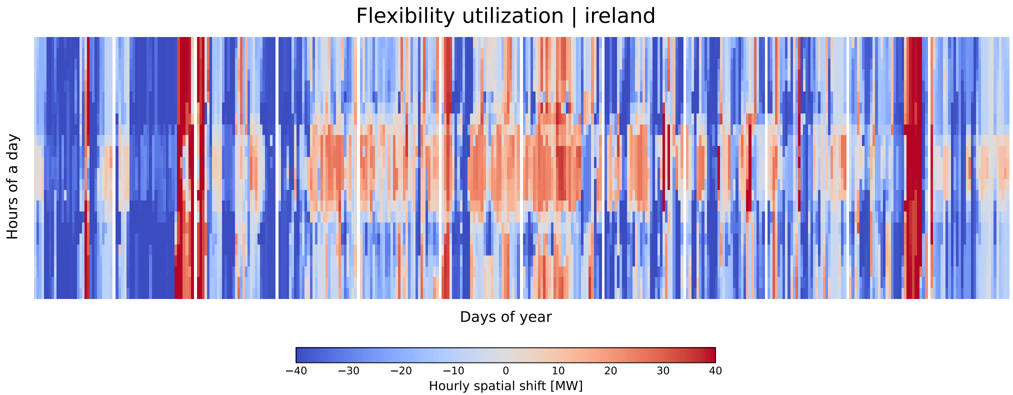

Insights from time-series data for optimized space-time load shifts

Another perspective on the **hourly spatial load shifts** for the same scenario (i.e., the same model run) but another location: a data center in Germany. Similarly, the **negative** values mapped to blue color represent loads "received" from other locations in a given hour, while **positive** values mapped to red color represent loads "sent" away.

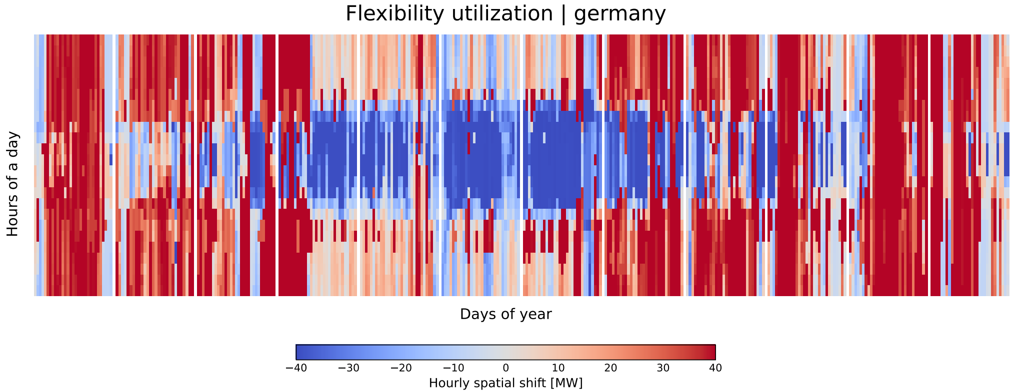

Insights from time-series data for optimized space-time load shifts

The heatmap plots above reveal several insightful observations on utilization of spatial flexibility. The spatial shifts of load have two distinct utilization patterns:

1. **A stochastic pattern:** In the European energy system, the hourly profiles of wind power generation have a low correlation over long distances due to different weather conditions. Spatial flexibility allows the system to take advantage of these differences: spatial flexibility enables "load arbitrage" between locations with different weather conditions. These load shifts are notable by the "vertical stripes" of a sudden color change (i.e., directions of spatial shifts) in the heatmaps above.

   In the Annex, we provide **energy balance plots** illustrating an example this behavior: a data center located in Ireland experiences a tough situation on 03-04 March, due to calm days in the region and low feed-in of the wind generators procured with PPAs. Imports from the regional grid do not help, because the local electricity mix is dirty in this period and data center's planned CFE score is 100%. However, a data center in Denmark experiences good wind conditions on 03-04 March and has excess generation of CFE from the procured porfolio. Thus, a load shift between the two data centers helps to resolve the situation.

   To cover the load with CFE around-the-clock in a scenario with limited load flexibility (10% of flexible workloads instead of 40%—see **this scenario** in the Annex), a data center in Ireland has to procure a **much bigger portfolio** of solar PV and battery storage. Furthermore, a data center in Denmark has to curtail the excess generation of clean electricity (unless it is sold to the regional grid). Overall, load flexibility facilitates a **better utilization** of locally procured resources by reducing the volume of renewable curtailment (**more detail** in the Annex).

Technische Universität Berlin

2. **A daily/seasonal pattern:** other load shifts are caused by the **differences in quality of local resources**. These load shifts are notable by the structured shapes a color change (i.e., directions of spatial shifts) in the flexibility utilization heatmaps.

   The quality of local resources, i.e., the average capacity factors of wind or solar PV in a given region, translates into the levelised costs of electricity (LCOE) for renewable generators. When spatial load shifting is possible, a rational 24/7 consumer can adjust their own procurement strategy by contracting generators in better locations (lower LCOE) and co-optimizing spatial loads shifts accordingly.

   The heatmaps above illustrate this behavior well: a data center located in Ireland—a region with poor solar resources—tends to shift loads away during the daytime from the mid-spring till mid-autumn. Instead, a data center located in Germany—a region with better solar resources—tends to receive loads during this period. It works just about reciprocally for wind-related load shifts: a data center in Germany benefits from having partners in Denmark and Ireland, the two very windy regions in Europe.

   In the supplementary graphics, we provide an example of a **data center in Germany** in the first week of May, where spatial load shifts have a clear daily profile (as can also be seen in the heatmap).

Let us switch our attention to the temporal load management (while staying with the same scenario, see bottom right for details). The plot below shows the **hourly temporal load shifts** for a data center in Portugal. **Negative** values mapped to blue color represent "increase" of a load, i.e., workloads are shifted to a given hour from other times, while **positive** values mapped to red color represent "decrease" of a load, thus workloads are shifted away from a given hour to another time.

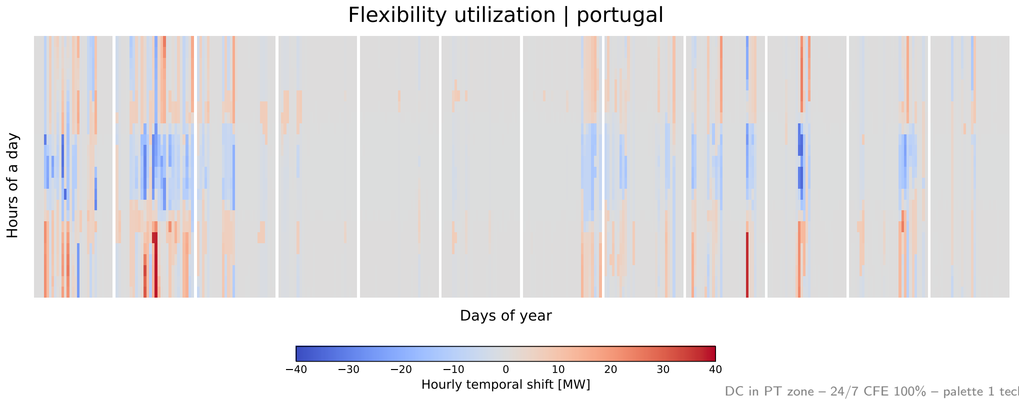

DC in PT zone – 24/7 CFE 100% – palette 1 tech – spatial & temporal shifts – 40% flexible workloads

Insights from time-series data for optimized space-time load shifts

Another perspective on the **hourly temporal load shifts** for the same location (data center in Portugal) but another CFE score: 98%. The color mapping stays as above.

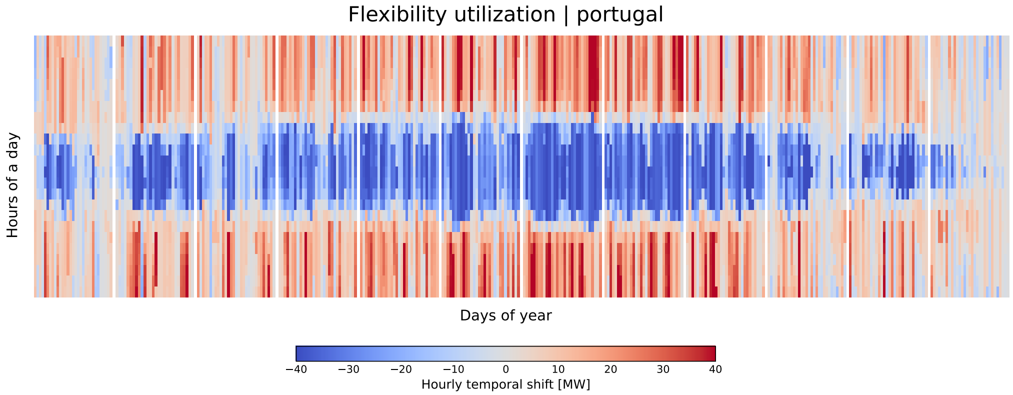

Hourly CFE score of supply from grid

The (modelled) **hourly CFE score of electricity supply from grid** for Portugal.

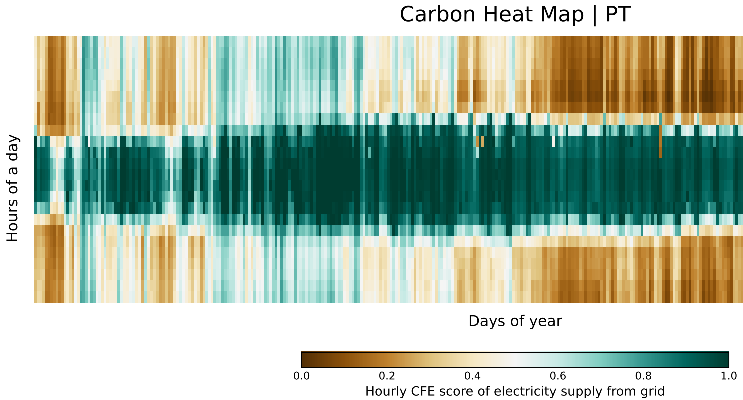

What can we learn about the temporal flexibility usage from the utilization heatmap plots above?

1. In the scenario with a 100% CFE score, data centers mainly rely on spatial flexibility utilization; the usage of temporal flexibility is comparably small.

   The potential for carbon-aware temporal load shifting is created by the variability of the regional grid emission intensity. However, to achieve the 100% CFE score, the 24/7 consumer relies mainly on its own procured generators and storages, i.e., there is (nearly) no imports from the regional grid. This effect was explained in detail in our **previous research** (pp. 33, 41, 47).[2] Thus, the variability of grid emission intensity does not have much influence.

   Temporal flexibility could also be helpful in aligning the demand with the generation of procured renewable generators; however, the potential of using temporal shifts to reduce the capacity fleet is limited by (i) the daily compute usage conservation rule (eq. 10) and (ii) the fact that temporal and spatial load shifts are subject to a shared set of computing capacity constraints (eq. 13a-13c), whereas spatial shifts bring more efficiency gains in this scenario, as shown **above**.

[2] In fact, in the discussion section of the previous study, we show that 24/7 consumers can occasionally rely on grid imports in the 100% CFE case, but that requires certain conditions to be met.

Insights from time-series data for optimized space-time load shifts

2. The results for the scenario with 98% CFE score show a different trend: temporal flexibility is actively utilized to shift load from night to mid-day hours, and this pattern has a seasonal profile.

   At the 98% CFE score, 24/7 consumer complements own portfolio of procured CFE technologies with imports of electricity from the regional grid. Thus, temporal shifts are **responsive to the regional grids' carbon intensity**, i.e., workloads are shifted to "greener" times. The carbon content of the regional grid correlates with a profile of solar PV feed-in, what gives the shape for the temporal flexibility utilization.

   The hourly CFE score of the electricity supply from the regional grid in Portugal is provided above, the **data other locations** is the the Annex.

   We provide supplementary graphics illustrating utilization of temporal flexibility in different contexts with energy balance plots for selected locations and time frames. For a **data center in Portugal (the first week of May)**, the energy balance plot shows the co-optimized utilization of temporal and spatial load flexibility. The temporal load shifts help to align data center's demand profile with procured solar PV generation and electricity imports in hours when the regional grid has a high CFE score. For a **data center in Ireland (the first week of December)**, the temporal load flexibility is used to minimize consumption during a difficult period of low renewable energy feed-in.

Insights from time-series data for optimized space-time load shifts

We wrap up our time-series results section with a brief look at the **hourly spatial load shifts** for the scenario when LDES is added to the technology mix. Spatial flexibility utilization has a more "binary" and complex pattern. This can be attributed to optimization across the *nodes x time periods graph* of all flexibility elements' dispatch decisions and the synergies among them. In particular, co-optimized utilization of spatial shifts and LDES system enables harvesting renewable resources in the best locations, storing it over long periods and providing an access to low-cost carbon-free electricity for all data centers when it is needed the most.

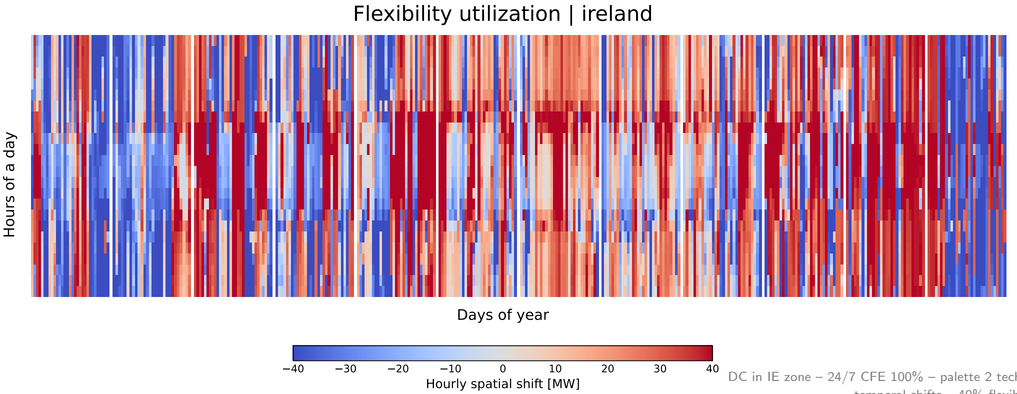

DC in IE zone – 24/7 CFE 100% – palette 2 tech – spatial & temporal shifts – 40% flexible workloads

Economically efficient redistribution of data center average loads (100% CFE)

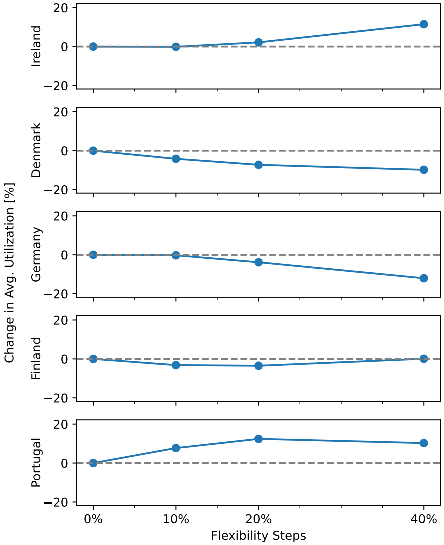

The analysis of the time-series data for space-time load shifts revealed interesting insights. However, a natural question to ask based on the results above is whether space-time flexibility also leads to the *net shift* of loads? In other words: *Does load flexibility facilitate a redistribution of the average utilization of data centers?*

A plot on the left summarised the time-series data and shows changes in the **average utilization** of data centers in each of the five locations as a function of load flexibility.

A interesting aspect of this result is that when 24/7 procurement is done with onshore wind, solar PV and battery storage, the increase of data center average utilization does not occur in locations with "greener" backgrounds grids, such as DK or FI. Instead, the average utilization of data centers increases in locations with **good renewable resources** that have generation profiles distinct to other locations. Such locations include Portugal (notable for its excellent solar resources, resulting in a high capacity factor for solar PV) and Ireland (notable for its good wind resources and a generation profile that is also uncorrelated to the four continental locations).

Economically efficient redistribution of data center average loads (100% CFE w/ LDES)

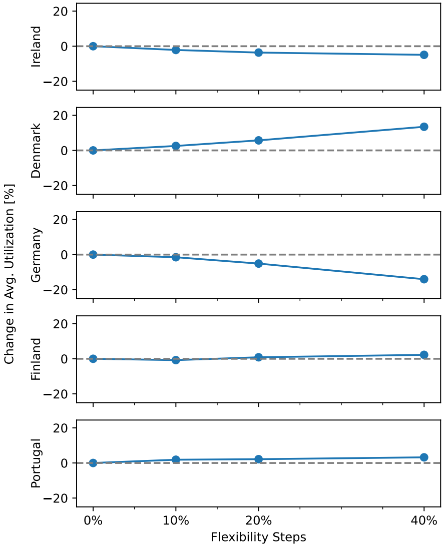

When LDES is added to the technology mix, the changes in data center average loads have another pattern. The average utilization increases only for the data center in Denmark, which has **the best wind conditions** among the five locations. This observation supports the findings on the synergy between space-time load flexibility and LDES in the context of 24/7 CFE shown in the sections on **24/7 costs**, **isolated values of space-time flexibility**, and **time-series analysis**.

Overall, the space-time load-shifting flexibility enables taking advantage of differences in local resource quality, harvesting renewable electricity in the best locations and indirectly **opening access to it for all locations**. When paired with the LDES, the efficiency gains of this effect become even larger. The LDES allows storing CFE over long periods, thus overcoming the restrictions of battery storage and the daily usage conservation rule of the temporal shifts.

Finally, the space-time load-shifting flexibility facilitates the **economically efficient redistribution of loads**, helping data centers to match demand with carbon-free electricity around-the-clock in a more cost-effective way.

Further remarks

- In the **previous study**, we showed that 24/7 carbon free energy matching results in a **notable and systematic reduction of emissions** both for participating consumers and in a regional grid (system-level emissions).

- While preparing this study, we observed that load flexibility *per se* does not decrease emissions further below relative to the high 24/7 CFE procurement baseline; however, flexibility makes achieving CFE targets and the associated system effects **more cost-effective**.

- This effect takes place because 24/7 participants with higher CFE scores rely more on their own portfolio of CFE resources and less on grid imports. For the high CFE scores considered in this study, imports from the local grid are possible; however, the hourly CFE score of imported electricity has to be high enough to match the CFE target. Thus, load flexibility is mainly used to optimize resources for matching demand with carbon-free electricity around-the-clock, which leads to the system effects of 24/7 procurement, such as lower system-wide emissions, at reduced procurement costs.

- Test model runs done without the 24/7 procurement constraint (eq. 15), i.e., simulating a case when data centers cover demand purely with grid purchases without any policy regarding the origin of electricity, result in active shifting flexible loads to "greener" times and locations (driven by the merit-order economics). In that scenario, the increased load flexiblity would be *ceteris paribus* responsible for a significant reduction in system-level emissions.

Further remarks

- As shown in this study, space-time load shifting makes clean electricity more accessible and gives flexible consumers more options for matching demand with carbon-free electricity around-the-clock. As a result, data center operators and other commercial and industrial consumers with flexible demands can **achieve high degrees of CFE matching at lower costs**, and may be interested in joining the **24/7 CFE movement**. Therefore, the system decarbonization impact associated with the 24/7 CFE procurement could be amplified with **greater participation** while requiring **fewer resources**.

- There is a number of initiatives to improve the European Guarantee of Origin (GO) mechanism—the largest standardized market for Energy Attribute Certificates (EACs) in the world—that has currently no recognised system of verifying renewable electricity supply on an hourly basis. For example, the **EnergyTag** initiative is developing a framework for adding a timestamp to EACs, which will make them more reflective of the physical availability of clean energy and allow companies trading CFE credits on hourly basis. ENTSO-E also published a **position paper** highlighting the need for the 24/7 GO system. When a market for the 24/7 GOs is created, the space-time load-shifting flexibility provided by data centers will theoretically decrease demand for GOs in hours when carbon-free electricity is the most expensive, which in turn will reduce certificate costs for **all electricity consumers** interested in sustainability goals, **regardless of their ability to shift loads**.

Conclusions

Conclusions

Conclusion 1: Demand flexibility enables **better access to clean electricity** and creates **more options** for consumers to match demand with carbon-free electricity around-the-clock.

Conclusion 2: Some flexible electricity consumers, such as data centers, can shift computing jobs and associated power loads in both time and location. These mechanisms facilitate the **efficiency and affordability** of 24/7 CFE procurement. The co-optimized space-time load-shifting can reduce the costs of 24/7 CFE by up to 34%, depending on the level of flexibility and technologies available.

Conclusion 3: Demand flexibility is **especially helpful for resource-constrained locations** where hourly matching with 24/7 CFE is difficult.

Conclusion 4: Space-time load-shifting facilitates **economically efficient redistribution of loads** to locations with good carbon-free resources. When paired with long-duration energy storage, the efficiency gains of this effect are even larger.

Conclusion 5: In the European energy system, the hourly profiles of wind power generation have a low correlation over long distances due to different weather conditions. Spatial load flexibility enables the system to move load to locations when and where there is high wind generation, thus **saving costs of energy storage** and **reducing curtailment** of excess generation.

The value of space-time load-shifting flexibility for 24/7 carbon-free electricity procurement

This study is done in a spirit of open and reproducible research:

https://github.com/PyPSA/247-cfe

A fixed link to the complete pack of results for this study:

https://doi.org/10.5281/zenodo.8185850

For questions and collaboration inquiries, please contact

Dr. Iegor Riepin, iegor.riepin@tu-berlin.de

Prof. Tom Brown, t.brown@tu-berlin.de

Suggested citation: Riepin, I. & Brown, T., *The value of space-time load-shifting flexibility for 24/7 carbon-free electricity procurement*, Department of Digital Transformation in Energy Systems TU Berlin, 26 July 2023. Zenodo. DOI: doi.org/10.5281/zenodo.8185850

Annex A: Methodology

Energy system model

- This study is done with a modified version of **PyPSA-Eur** – an open optimization model of the European energy system.
- PyPSA-Eur offers an automated and configurable software pipeline enables scientific workflow from freely available and open raw input data to optimized energy system.
- The model is suitable both for operational studies, as well as generation and transmission expansion planning studies.
- PyPSA-Eur is an open-source project:
  - PyPSA-Eur on GitHub
  - Documentation
  - Feature summary

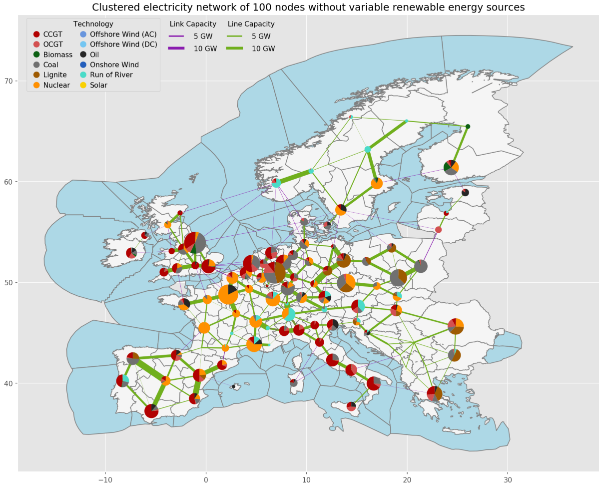

Image: github.com/PyPSA/pypsa-eur

Mathematical model of 24/7 CFE procurement

Technische Universität Berlin

- The mathematical model of 24/7 CFE procurement is based on the former work of authors: **System-level impacts of 24/7 carbon-free electricity procurement in Europe** published in October 2022. The study included mathematics additional to the PyPSA-Eur model to encode a situation when a fraction of C&I consumers in a selected European countries commit to the 24/7 CFE goals. The resulting problem optimized investment and operational decisions to meet projected electricity demand for the 24/7 CFE consumers, as well as the demand of other consumers in the European electricity system, while meeting all relevant engineering, reliability, and policy constraints.

- In this study, we enhance the mathematical model of 24/7 CFE procurement by considering demand flexibility provided by data centers. The load flexibility involves **temporal** (computing jobs scheduling) and **spatial** (computing jobs migration) load shifting.

- Thus, a data center operator (i.e., a flexible consumer following 24/7 CFE goal) can meet a given CFE target by either procuring energy generation and storage assets directly, and buying electricity from a local grid in hours when electricity mix is sufficiently clean (like in the previous study), as well as utilize spatial and/or temporal flexibility to achieve hourly matching of demand with clean electricity more efficiently.

Matching electricity supply and demand: a case of inflexible consumer

The model optimizes a portfolio of carbon-free generation and storage technologies procured by the C&I consumers that commit to 24/7 CFE goal. The portfolio assets have to be located in the same market zone.

The hourly demand of 24/7 participating consumer d_t for hour t can be met by a combination of the following:

- dispatch $g_{r,t}$ of procured carbon-free generators $r \in CFE$
- dispatch $\bar{g}_{s,t}$ of procured storage technologies $s \in STO$ (requires charge $\underline{g}_{s,t}$)
- imports of electricity from the grid im_t.

$$\sum_{r \in CFE} g_{r,t} + \sum_{s \in STO} \left(\bar{g}_{s,t} - \underline{g}_{s,t} \right) - ex_t + im_t = d_t \qquad \forall t \quad (2)$$

NB: the excess from the local supply ex_t can either be sold to the grid at market prices or curtailed.

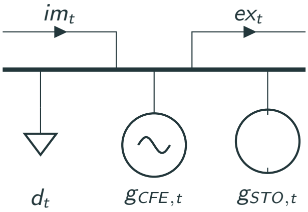

The **24/7 CFE matching** is modelled with a constraint (3), which matches demand of participating consumers with carbon-free resources on an hourly basis. The constraint ensures that sum over generators from procured CFE resources $r \in CFE$, discharge and charge from storage technologies $s \in STO$, as well as import from the grid im_t multiplied by the grid's CFE factor CFE_t must be higher or equal than a certain **CFE score** x multiplied with the total load d_t:

$$\sum_{r \in CFE, t \in T} g_{r,t} + \sum_{s \in STO, t \in T} \left(\bar{g}_{s,t} - \underline{g}_{s,t} \right) - \sum_{t \in T} ex_t + \sum_{t \in T} CFE_t \cdot im_t \geq x \cdot \sum_{t \in T} d_t \tag{3}$$

In any given hour, a data center's energy profile takes one of the following forms:

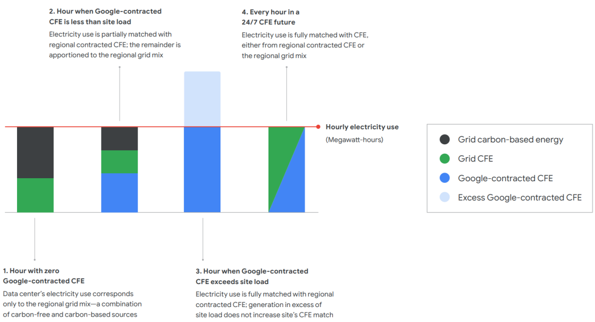

The **CFE score** x [%] measures the degree to which hourly electricity consumption is matched with carbon-free electricity generation within the regional grid. The matric is calculated using both CFE contracted by 24/7 participant, as well as CFE coming from the regional grid mix.

The 24/7 CFE matching concept is aligned with **24/7 CFE: Methodologies and Metrics** paper by Google.

Image: 24/7 CFE: Methodologies and Metrics, Google 2021

24/7 CFE matching: grid CFE factor

The **grid CFE factor** CFE_t in eq. (3) defines the percentage of clean electricity in each MWh of imported electricity from the grid to supply participating 24/7 loads in a given hour. The factor depends on the generation mix in the region where 24/7 participant is located, as well as on the generation mix in other regions from which electricity is imported to the local region ($import_t$).

Using notation on the right, the average cleanness of the rest of the electricity system is:

$$ImportCFE_t = \frac{A_t}{A_t + D_t}$$

The CFE factor of grid supply[a] for a given hour t is:

$$CFE_t = \frac{B_t + ImportCFE_t * import_t}{B_t + E_t + import_t}$$

[a]Generators contracted by 24/7 consumers (C) are excluded from the grid supply.

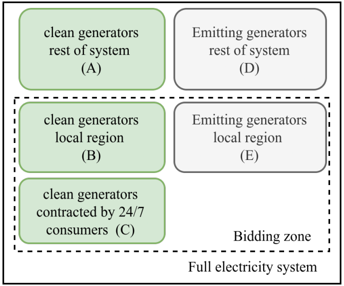

Here we follow **Xu et al. (2021)**

Note that the grid CFE factor is affected by capacity procured by 24/7 consumers. This introduces a nonconvex term to the optimization problem. The nonconvexity can be avoided by treating the grid CFE factor as a parameter that is iteratively updated (starting with $CFE_t = 0 \;\forall t$). In the previous study, we concluded that one forward pass (i.e. 2 iterations) yields very good convergence. This observation holds true also for the optimization problem behind this study with multiple 24/7 consumers.

24/7 CFE matching: excess CFE

The **excess CFE** represents generation from the procured resources above consumption of the 24/7 participant in a particular hour. The excess CFE **is not counted toward the CFE score** – and thus it is subtracted on the left-hand side of the eq. (3). While it does not contribute to the CFE Score, excess CFE could potentially be stored (using batteries) and shifted to another hour, sold to the regional grid at **market prices**, or curtailed.

The total amount of CFE exported to the regional grid is constrained to a certain level on an annual basis. The export limit (*ExLimit*) is set to 20% of annual 24/7 participating consumer's demand. Thus, constraint (4) gives the 24/7 participant flexibility to sell electricity to the regional grid, while avoiding the situation that sales to the grid become significantly larger than CFE supply to own demand.

$$\sum_{t \in T} export_t \leq ExLimit \cdot \sum_{t \in T} d_t \tag{4}$$

The **market prices** are derived from the dual variable of each zone's **energy balance constraint**. An infinitely small relaxation of the constraint, i.e., one unit of load less to be met, returns the marginal costs of providing that unit, which can be used as the electricity price indicator in a competitive market.

Spatial load shifting problem 1/3

We introduce a concept of **spatial load management system** that allows for shifting load across locations. The load shifts take place via **virtual links** – ICT-based pathways between data centers.

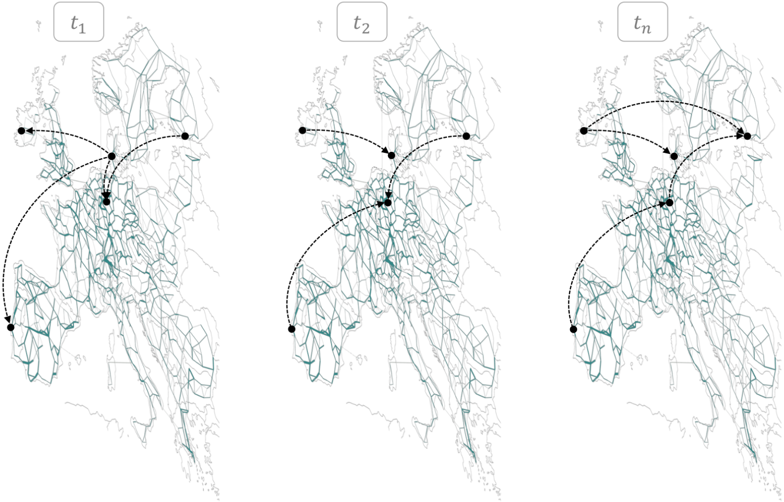

Spatial load shifting problem 2/3

We introduce a concept of **spatial load management system** that allows for shifting workloads across locations. The load shifts take place via **virtual links** – ICT-based pathways between data centers.

Here we follow a mathematical formulation of virtual links proposed by Zhang & Zavala (2022). Let Θ be the set of all virtual links; let $\delta_\vartheta \in \mathbb{R}_+$ be load shifts (flows via virtual pathways); and let N_{DC} be the set of data centers (flexible consumers). We can define $\Theta_n^{snd} := \{\vartheta \in \Theta | snd(\vartheta) = n\} \subseteq \Theta$ and $\Theta_n^{rec} := \{\vartheta \in \Theta | rec(\vartheta) = n\} \subseteq \Theta$ to be the set of sending and receiving virtual links at node $n \in N_{DC}$.

The nodal energy balance defined for inflexible consumers (eq. 2) is now extended by variables representing shifts of load **across locations**, since the dispatched load at a given node can include shifts to/from other data center nodes:

$$\begin{aligned}&\sum_{r\in CFE} g_{r,n,t} + \sum_{s\in STO}\left(\bar{g}_{s,n,t} - \underline{g}_{s,n,t}\right) - ex_{n,t} + im_{n,t} = \\ &d_{n,t} + \sum_{\vartheta\in\Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta\in\Theta_n^{snd}} \delta_{\vartheta,t} \qquad \forall n \in N_{DC}, t \in T\end{aligned} \tag{5}$$

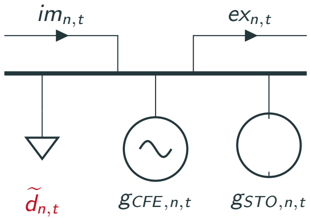

Spatial load shifting problem 3/3

Computing capacity constraints (eq. 6) ensure that the dispatched load at each data center $\widetilde{d}_{n,t}$ does not exceed available capacity (an upper limit, eq. 6b), as well as that a certain data center does not shift load that exceeds flexible jobs share (a lower limit, eq. 6c).

$$\widetilde{d}_{n,t} = d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} \quad \forall n \in N_{DC}, t \in T \tag{6a}$$

$$\widetilde{d}_{n,t} \leq [1+f] \cdot d_{n,t} \quad \forall n \in N_{DC}, t \in T \tag{6b}$$

$$\widetilde{d}_{n,t} \geq [1-f] \cdot d_{n,t} \quad \forall n \in N_{DC}, t \in T \tag{6c}$$

NB spatial load shifts are not subject to any electricity network transmission constraints; as such, the only source of congestion for the virtual links is computing capacity constraints (i.e., availability of flexible workloads).

The 24/7 CFE matching constraint for inflexible consumer (eq. 3) is now defined over a set of data center nodes $n \in N_{DC}$ and is extended on the right-hand side by spatial load shifts. Thus, flexible consumer can benefit from an additional degree of freedom that helps relaxing the constraint for locations and times when providing demand with carbon-free electricity is difficult:

$$\sum_{r \in CFE, t \in T} g_{r,n,t} + \sum_{s \in STO, t \in T} \left(\bar{g}_{s,n,t} - \underline{g}_{s,n,t} \right) - \sum_{t \in T} ex_{n,t} + \sum_{t \in T} CFE_{n,t} \cdot im_{n,t} \geq$$
$$x_n \cdot \sum_{t \in T} \left(d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} \right) \quad \forall n \in N_{DC} \tag{7}$$

Temporal load shifting problem 1/3

To capture temporal flexibility, we introduce a concept a data center **temporal load management system** that allows for shifting load from a given time to another time point in the future.

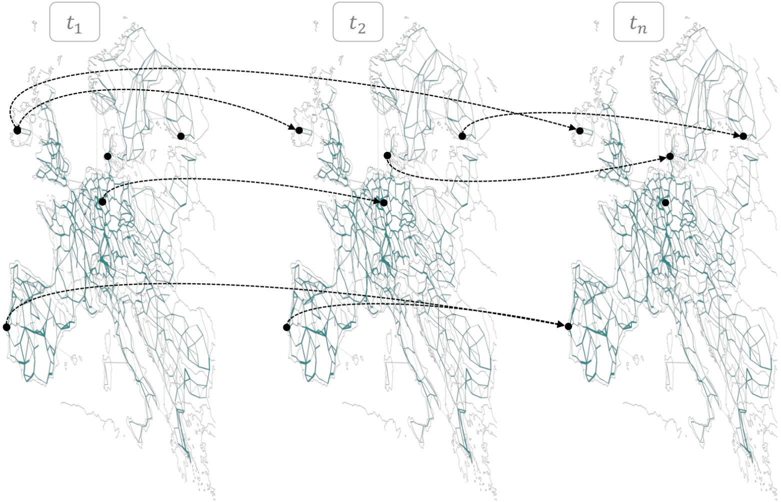

Temporal load shifting problem 2/3

To capture temporal flexibility, we introduce a concept a data center **temporal load management system** that allows for shifting load from a given time to another time point in the future.

Consider a time horizon of our optimization problem $T = \{t_1, t_2, ..., t_T\}$. For simplicity, let us assume that there is a single flexible consumer (i.e., no spatial load shifts) with a demand-side temporal load management mechanism denoted with a singleton set $\{s'\}$. Let variables $\bar{g}_{s',t}, \underline{g}_{s',t} \in \mathbb{R}_+$ be workloads that are rescheduled in time, i.e., shifted from a time t to a later time t'. Thus, the dispatched load $\widetilde{d}_t$ of flexible consumer can deviate from the nominal value $d_{n,t}$ due to temporal load management.

The nodal energy balance is now extended with variables representing shifts of load **across time**:

$$\sum_{r \in CFE} g_{r,t} + \sum_{s \in STO} \left(\bar{g}_{s,t} - \underline{g}_{s,t}\right) - ex_t + im_t = \\ d_t + \sum_{s'} \left(\bar{g}_{s',t} - \underline{g}_{s',t}\right) \qquad \{N_{DC}\}, \forall t \in T \tag{8}$$

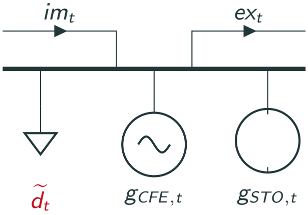

Temporal load shifting problem 3/3

Computing capacity constraints for temporal load shifting problem (eq. 9) ensure that workloads delayed to a given time t do not exceed available cluster capacity (an upper limit, eq. 9b), as well as that only flexible workloads can be shifted in time (a lower limit, eq. 9c).

$$\widetilde{d}_t = d_t + \sum_{s'} \left(\bar{g}_{s',t} - \underline{g}_{s',t} \right) \quad \forall t \in T \quad (9a)$$

$$\widetilde{d}_t \leq [1 + f] \cdot d_t \quad \forall t \in T \quad (9b)$$

$$\widetilde{d}_t \geq [1 - f] \cdot d_t \quad \forall t \in T \quad (9c)$$

We follow Radovanovic et al. (2021) implementing the daily usage conservation rule – an additional constraint to ensure that the cluster-level daily compute usage is preserved when flexible workload is shifted in time:

$$\sum_{t | t \in t(DAYS)} \left(\bar{g}_{s',t} - \underline{g}_{s',t} \right) = 0 \quad \{s'\} \quad (10)$$

The 24/7 CFE matching constraint is also extended to account for the temporal load management mechanism. Consumer with temporally flexible demand benefits from an additional degree of freedom that helps achieving a CFE target by shifting load away from hours when matching demand with carbon-free electricity is expensive:

$$\sum_{r \in CFE, t \in T} g_{r,t} + \sum_{s \in STO, t \in T} \left(\bar{g}_{s,t} - \underline{g}_{s,t} \right) - \sum_{t \in T} ex_t + \sum_{t \in T} CFE_t \cdot im_t \geq$$
$$x \cdot \sum_{t \in T} \left(d_t + \sum_{s'} \left(\bar{g}_{s',t} - \underline{g}_{s',t} \right) \right) \quad \{N_{DC}\} \quad (11)$$

Spatially-temporal load shifting problem 1/3

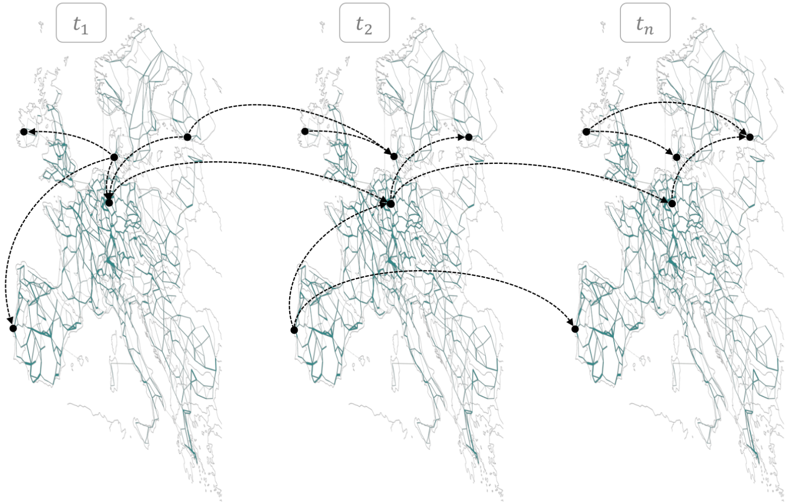

Spatially-temporal load shifting problem 2/3

The temporal and spatial flexibility of electricity demand can be **co-optimized** to help achieve clean electricity targets. The resulting mathematical problem brings together the formulations of spatial and temporal load management systems shown above.

We consider a set of data centers (flexible consumers) $n \in N_{DC}$ located in various locations within the electricity network. Data centers are interconnected with virtual links $\Theta_n^{snd}, \Theta_n^{rec}$ (complete graph). Each data center also has a temporal load management mechanism $S_n^{dsm} := \{s' \in S | dsm(s') = n\}$.

The nodal energy balance is adjusted to account for variables representing load shifts **across space and time**:

$$\sum_{r \in CFE} g_{r,n,t} + \sum_{s \in STO} \left(\bar{g}_{s,n,t} - \underline{g}_{s,n,t} \right) - ex_{n,t} + im_{n,t} =$$
$$d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} + \sum_{s' \in S_n^{dsm}} \left(\bar{g}_{s',n,t} - \underline{g}_{s',n,t} \right) \qquad (12)$$
$$\forall n \in N_{DC}, t \in T$$

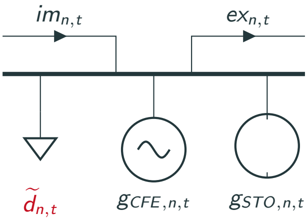

Technische Universität Berlin

Computing capacity constraints (eq. 13) now ensure that the dispatched load at each data center $\widetilde{d}_{n,t}$ does not exceed the limits for each data center $n \in N_{DC}$ considering both spatial and temporal load shifts at each time point t.

$$\widetilde{d}_{n,t} = d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} + \sum_{s' \in S_n^{dsm}} \left(\bar{g}_{s',n,t} - \underline{g}_{s',n,t} \right) \quad \forall n \in N_{DC}, t \in T \tag{13a}$$

$$\widetilde{d}_{n,t} \leq [1+f] \cdot d_{n,t} \quad \forall n \in N_{DC}, t \in T \tag{13b}$$

$$\widetilde{d}_{n,t} \geq [1-f] \cdot d_{n,t} \quad \forall n \in N_{DC}, t \in T \tag{13c}$$

The daily compute usage conservation rule is applied to each data center:

$$\sum_{t | t \in t(DAYS)} \left(\bar{g}_{s',t} - \underline{g}_{s',t} \right) = 0 \quad \forall s' \in S_n^{dsm} \tag{14}$$

Finally, the 24/7 CFE matching constraint is adjusted accordingly. With co-optimization of temporal or spatial load shifting, flexibility can be harnessed to achieve clean electricity targets more efficiently:

$$\sum_{r \in CFE, t \in T} g_{r,n,t} + \sum_{s \in STO, t \in T} \left(\bar{g}_{s,n,t} - \underline{g}_{s,n,t} \right) - \sum_{t \in T} ex_{n,t} + \sum_{t \in T} CFE_{n,t} \cdot im_{n,t} \geq x_n \cdot \sum_{t \in T} \left(d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} + \sum_{s' \in S_n^{dsm}} \left(\bar{g}_{s',n,t} - \underline{g}_{s',n,t} \right) \right) \quad \forall n \in N_{DC} \tag{15}$$

Annex B: Tools and data sources

PyPSA: an energy modelling ecosystem

- pypsa.org project provides a free, user-friendly and performant model environment to support a smooth energy transition around the world.
- The project includes individual packages that enable to go all the way from data processing (e.g., calculating renewable energy potentials or collecting energy assets data) to creating complex energy optimization problems.
- All packages are build in a modular sense so that they may be used independently from each other but interact easily.
- PyPSA development and maintenance is coordinated by the Department of Energy Systems @ TU Berlin **(ENSYS)**.

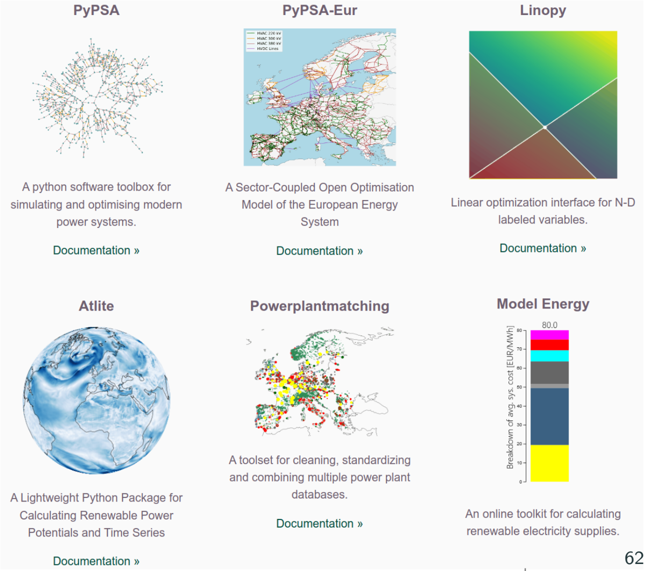

Image: pypsa.org

Data sources: electricity grid

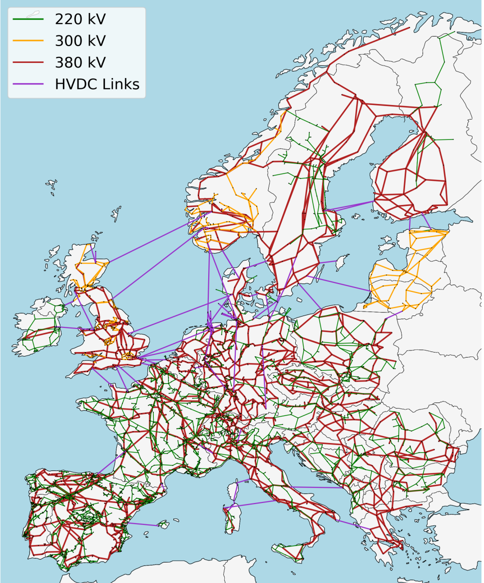

- Grid data contains AC lines at and above 220 kV voltage level, all high voltage DC lines, and substations for the full **ENTSO-E area**.
- Grid data is collected by a modified **GridKit** extraction of the **ENTSO-E Transmission System Map**. GridKit uses spatial and topological analysis to transform map objects from the ENTSO-E interactive map into a network model of the electric power system. The full grid model contains near 6760 lines and 3640 substations.
- The number of nodes fed into optimization model is adjustable, what allows for spatial and topological analysis at **different levels**. The number of nodes can vary between 37 (the number of independent countries / synchronous areas) and several hundred (for computational tractability).

Image: github.com/PyPSA/pypsa-eur

Data sources: power plants

- Existing power generation fleet data is collected with a **powerplantmatching** toolset.
- Powerplantmatching cleans, standardizes and merges the data from multiple open **power plant datasets** to create a combined dataset, which includes all the important information about power plants in Europe in a ready-to-use format for energy system modelling.
- The toolset allows to update the combined data as soon as new input datasets are released.
- Powerplantmatching is an open-source project maintained by TU Berlin team.
  GitHub
  Documentation

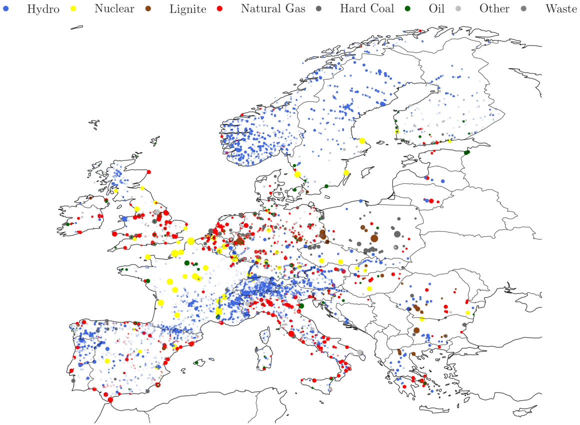

Image: github.com/PyPSA/powerplantmatching

Data sources: technology cost assumptions

- The database of assumptions for energy system technologies (such as capital and operational costs, efficiencies, lifetimes, etc.) is retrieved from the repository **PyPSA/technology-data**.

- The technology-data project compiles information about energy technologies from a variety of sources. The complied dataset has standardized technology names and energy units. All values are linked to original sources.

- technology-data is an open-source project maintained by TU Berlin team.
  [GitHub](#)
  [Documentation](#)

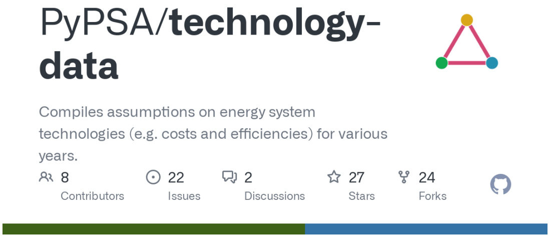

Cost assumptions used in this study originate from the **technology data catalogue** published by The Danish Energy Agency.

Image: github.com/PyPSA/technology-data

Data sources: renewable potentials and time series

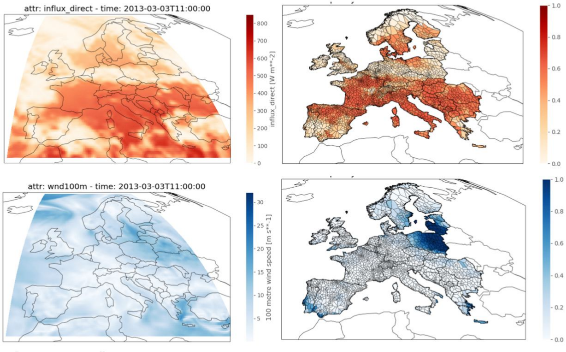

Converting weather data to energy system data

- Renewable power potentials and generation profiles are processed by the **atlite** package, which converts terabytes of weather data (like wind speeds, solar influx) into the data for energy systems modelling.
- With atlite, we process datasets for land cover (CORINE2018), natural protection areas (NATURA2000), and bathymetry (GEBCO2018) to conduct own geospatial land availability analysis.
- The standard data source for renewable time time-series estimation is ECMWF's ERA5 dataset (reanalysis weather data in ca. 30km x 30km and hourly resolution).
- atlite is also an open-source project maintained by TU Berlin team.
  GitHub
  Documentation

Image: atlite.readthedocs.io/

Other background system assumptions

- Electrical demand time-series is based on the **OPSD project**. We assume the same demand profile per bidding zone for 2025 and 2030, as in the representative year 2013.
- We assume 2013 as the representative climate year for renewable in-feed.
- Renewable expansion in the background electricity system is endogenous; we implement renewable generation targets by country that follow the **national energy and climate plans**. For countries w/o a 2025 target, a linear increase from renewable generation in 2020 to 2030 target is assumed. The modelled CO_2 emission intensity of electricity generation in European energy system matches the **estimated values** for 2025/2030.
- National policies and decommissioning plans for coal and nuclear power plants are based on the **Europe Beyond Coal**, and **world-nuclear.org** projects.
- We assume price for EU ETS allowances to be 80 €/tCO_2 for 2025. The price for natural gas is assumed to be 35 €/MWh.[3]

[3] Aligned with natural gas price assumptions in the **REPowerEU Plan** issued by the European Commission in 2022.

Assumptions about technologies available for 24/7 CFE consumers

Year	Technology	CAPEX (overnight cost)	FOM (%/year)	VOM (€/MWh)	Efficiency (per unit)	lifetime (years)
2025	utility solar PV	612 €/kW	1.7	0.01	-	37.5
2025	onshore wind	1077 €/kW	1.2	0.015	-	28.5
2025	battery storage	187 €/kWh	0	-	-	22.5
2025	battery inverter	215 €/kW	0.3	-	0.96	10.0
2025	hydrogen storage[4]	2.5 €/kWh	0	-	-	100.0
2025	electrolysis	550 €/kW	2.0	-	0.67	27.5
2025	fuel cell	1200 €/kW	5.0	-	0.50	10.0

Data is originally retrieved from the **DEA's catalogue for energy technologies**

[4] Underground hydrogen storage in salt cavern

Annex C: Supplementary graphics

Hourly CFE score of supply from grid – DE 2025

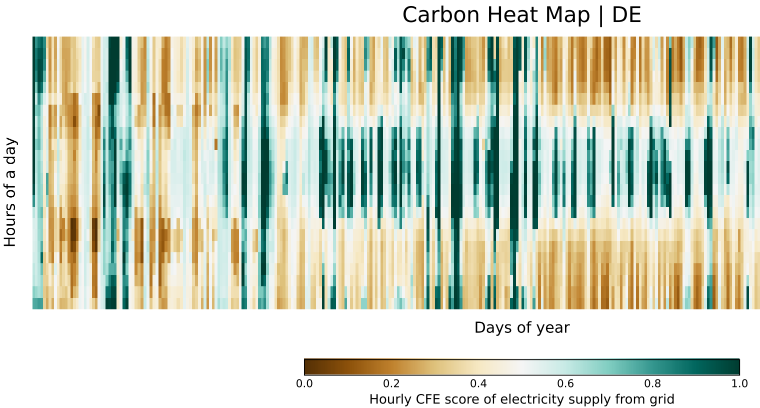

Hourly CFE score of supply from grid – DK 2025

Technische Universität Berlin

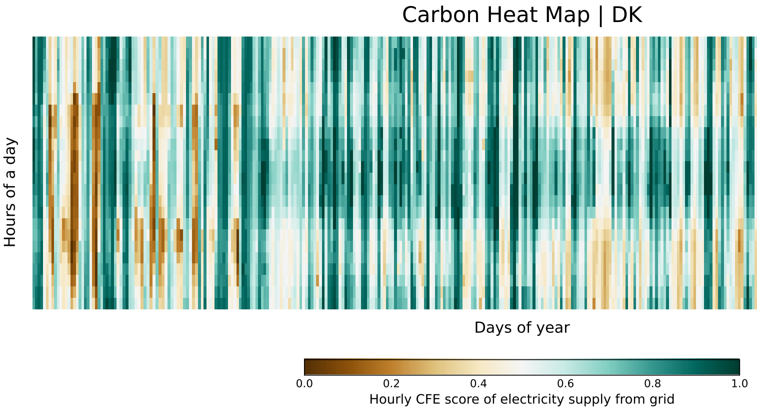

Hourly CFE score of supply from grid – IE 2025

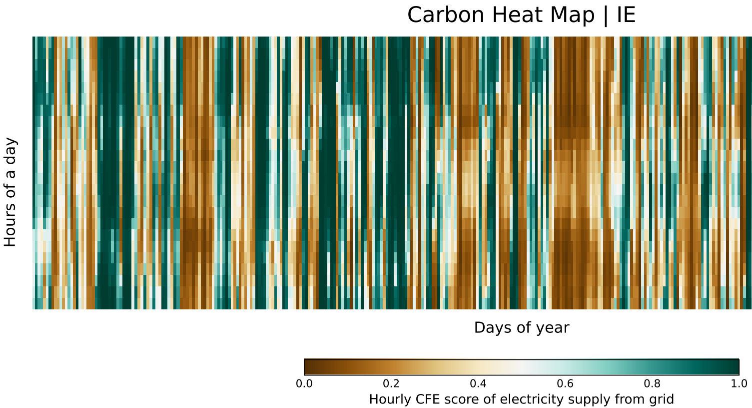

Hourly CFE score of supply from grid – FI 2025

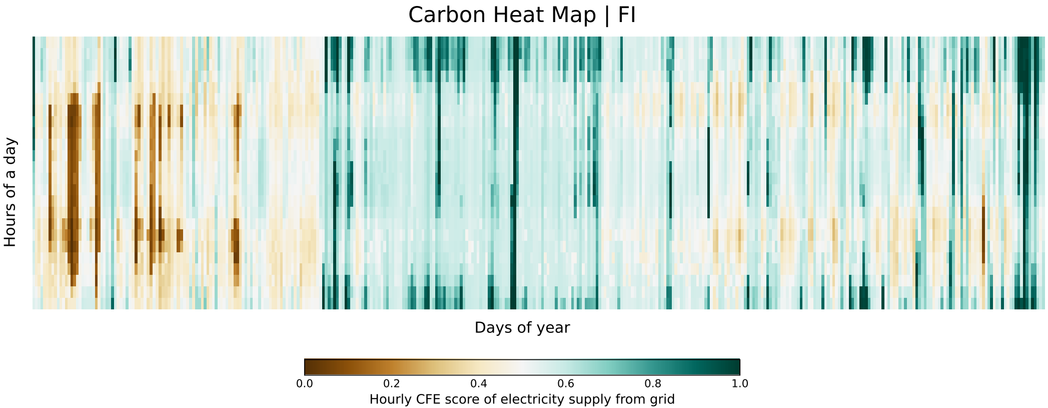

24/7 CFE costs as a function of load flexibility: isolated spatial load shifting

Isolating values of spatial and temporal load management for a scenario with **CFE score of 100% and technology palette 1**. The plots below show the average costs (per MWh of consumption) **(left)** and the total annual costs (per annum) for achieving 24/7 policy in all locations **(right)** with **spatial** load shifting.

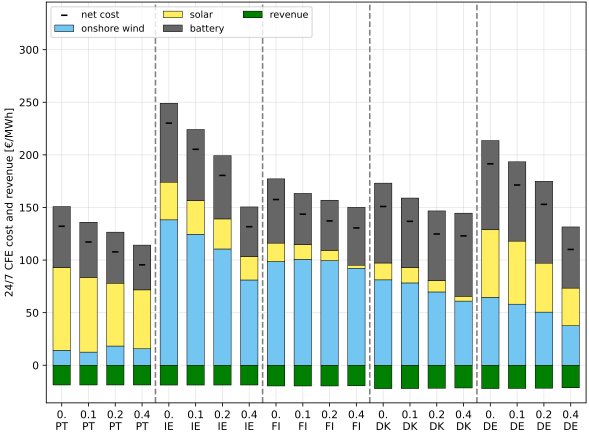

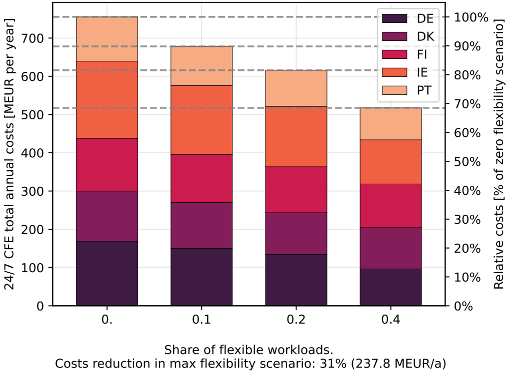

24/7 CFE costs as a function of load flexibility: isolated temporal load shifting

Isolating values of spatial and temporal load management for a scenario with **CFE score of 100% and technology palette 1**. The plots below show the average costs (per MWh of consumption) **(left)** and the total annual costs (per annum) for achieving 24/7 policy in all locations **(right)** with **temporal** load shifting.

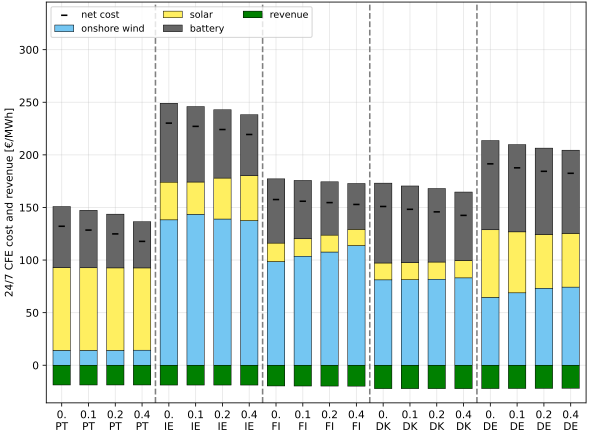

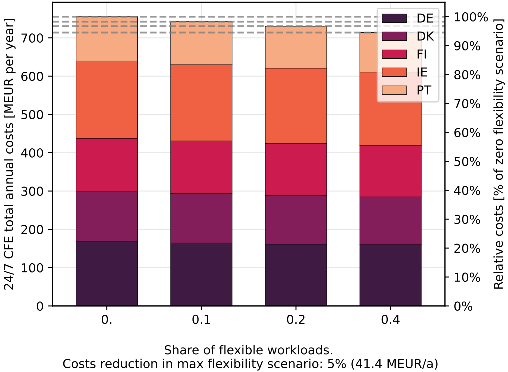

Economic efficiency of co-optimized space-time load shifts 98% CFE and 100% CFE w/ LDES

Technische Universität Berlin

The two plots below complement the discussion on "Isolated values of spatial and temporal load shifts". The plots show the **total annual costs [€/a]** for achieving 24/7 CFE policy in all locations (left y-axis) and their relative representation as a percentage of the zero flexibility scenario's costs (right y-axis). The costs are plotted as a function of the load flexibility potential.

Left panel: 98% CFE scenario; **right panel**: 100% CFE w/ LDES scenario;

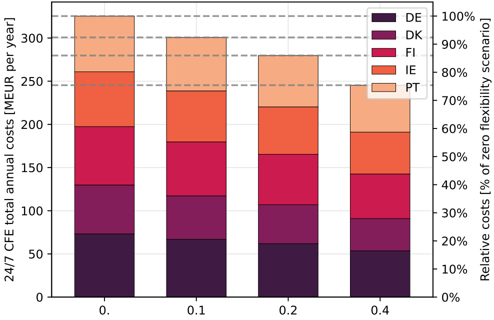

Share of flexible workloads.
Costs reduction in max flexibility scenario: 25% (80.1 MEUR/a)

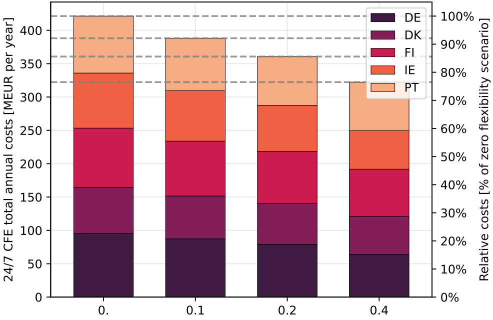

Share of flexible workloads.
Costs reduction in max flexibility scenario: 24% (99.0 MEUR/a)

Reduction of renewable energy curtailment

The plots below show the absolute amount of energy curtailment [GWh] from the portfolio of CFE generators procured by a data center operator as a function of load flexibility. The results are displayed per share of flexible loads $f = \{0\%, 10\%, 20\%, 40\%\}$.

Left panel: technology palette 1 (no LDES); **right panel**: technology palette 2 (with LDES)

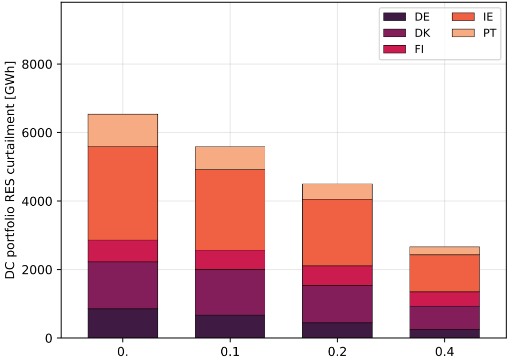

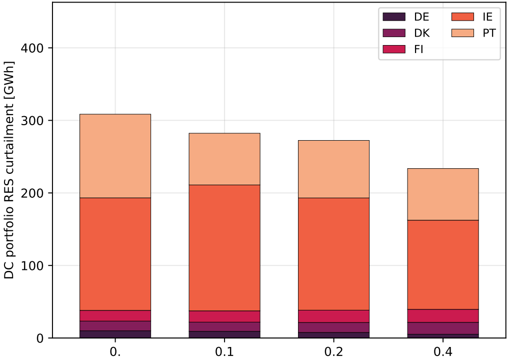

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Ireland.
The first week of March.
40% of flexible workloads.
100% CFE score.

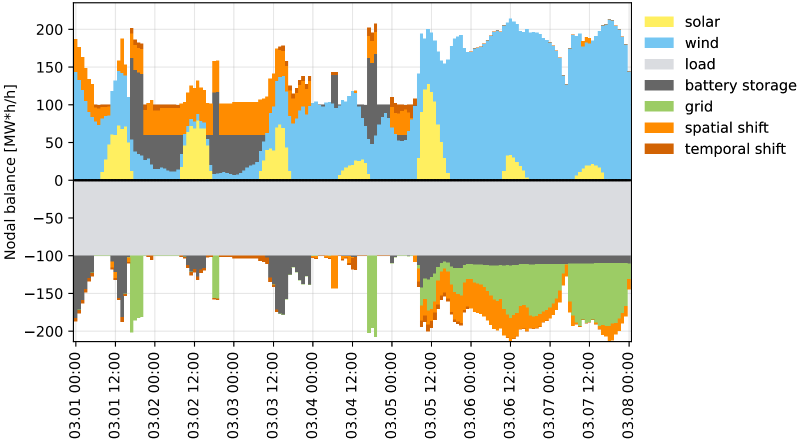

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Ireland.
The first week of March.
10% of flexible workloads.
100% CFE score.

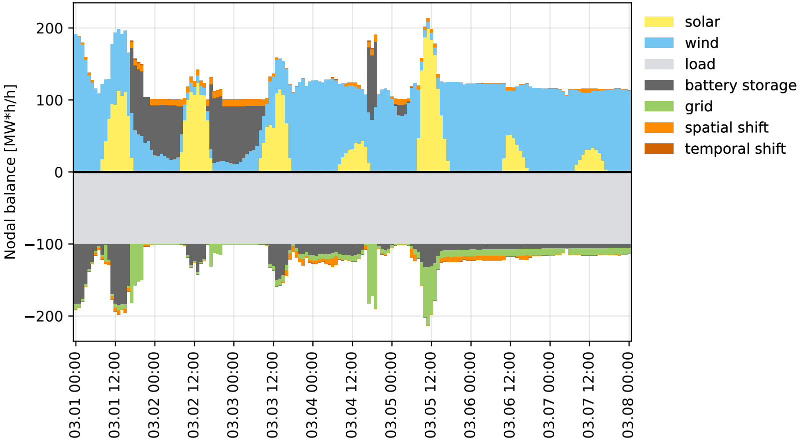

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Denmark.
The first week of March.
40% of flexible workloads.
100% CFE score.

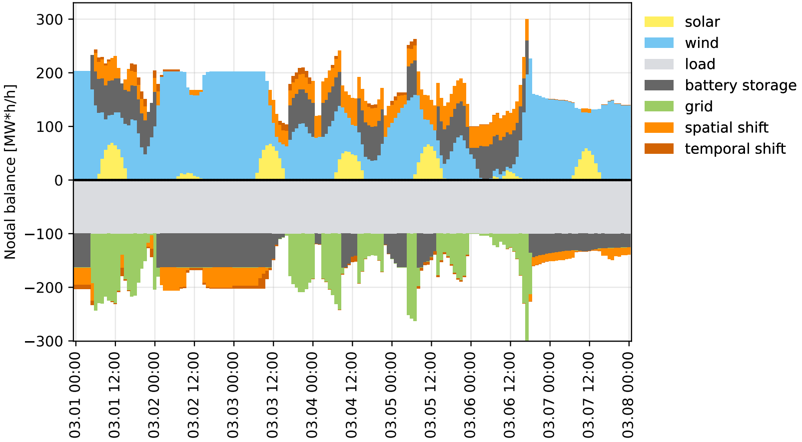

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Denmark.
The first week of March.
10% of flexible workloads.
100% CFE score.

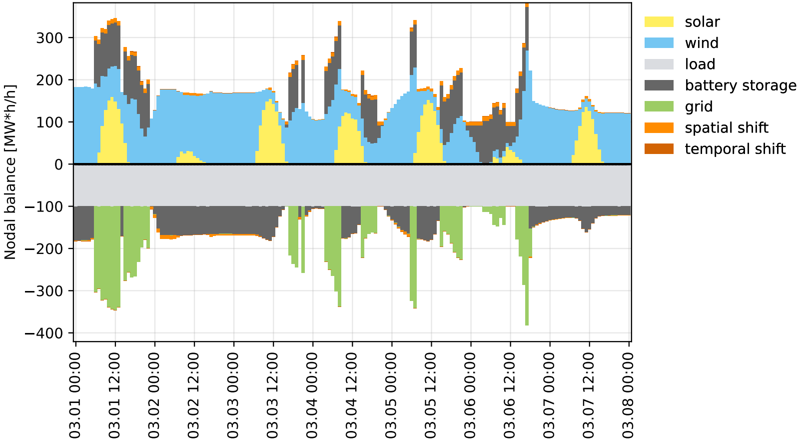

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Germany.
The first week of May.
40% of flexible workloads.
100% CFE score.

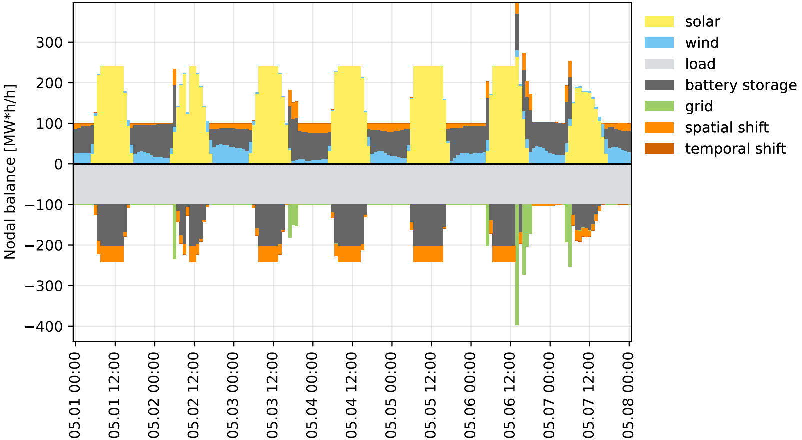

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Portugal.
The first week of May.
40% of flexible workloads.
98% CFE score.

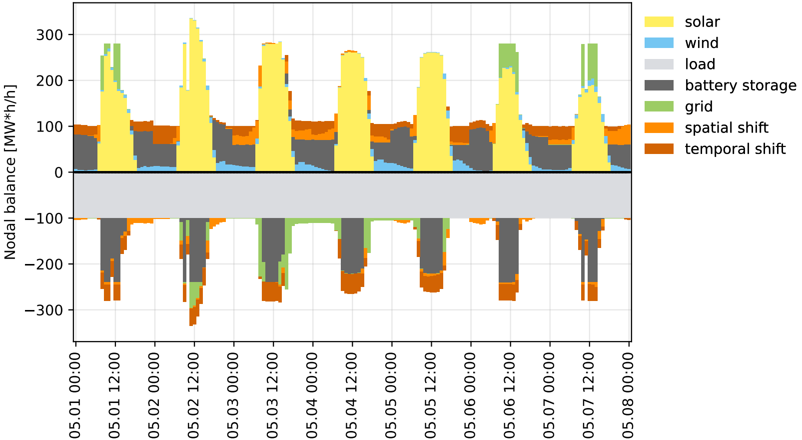

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Ireland.
The first week of December.
40% of flexible workloads.
98% CFE score.

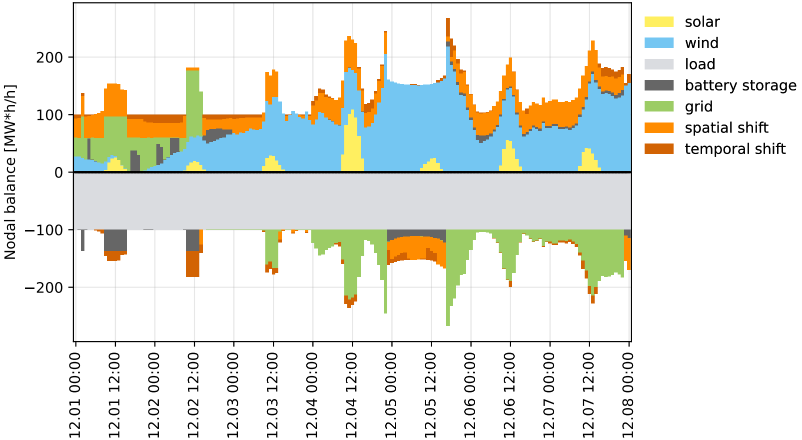

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Ireland.
The first week of March.
40% of flexible workloads.
100% CFE score.
+ LDES (palette 2).

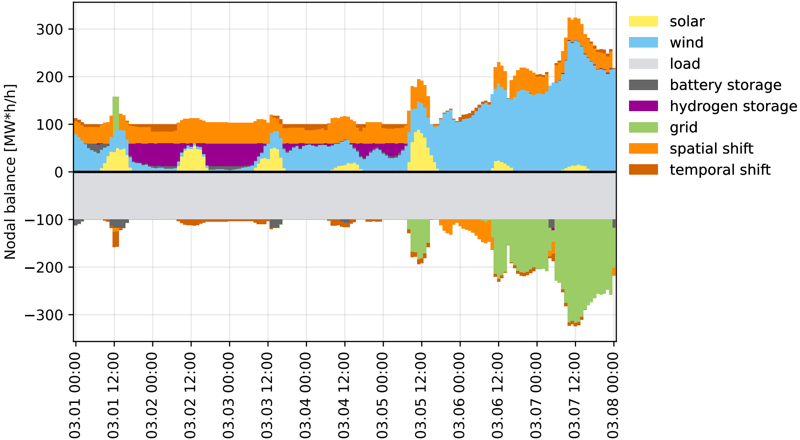

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Ireland.
The first week of March.
0% of flexible workloads.
100% CFE score.
+ LDES (palette 2).

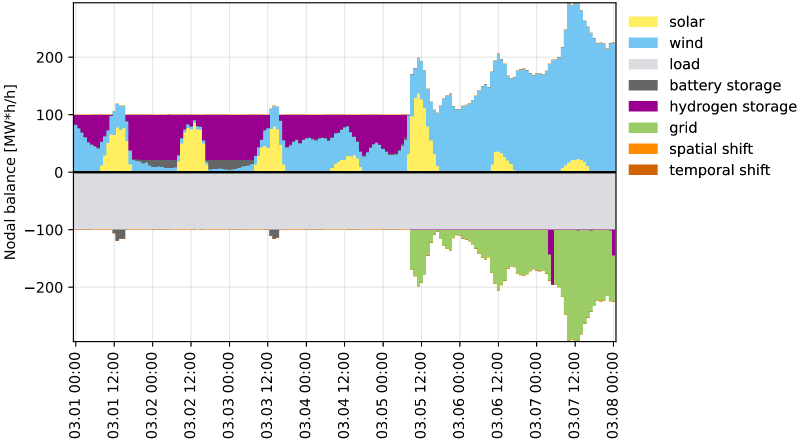

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Denmark.
The first week of March.
40% of flexible workloads.
100% CFE score.
+ LDES (palette 2).

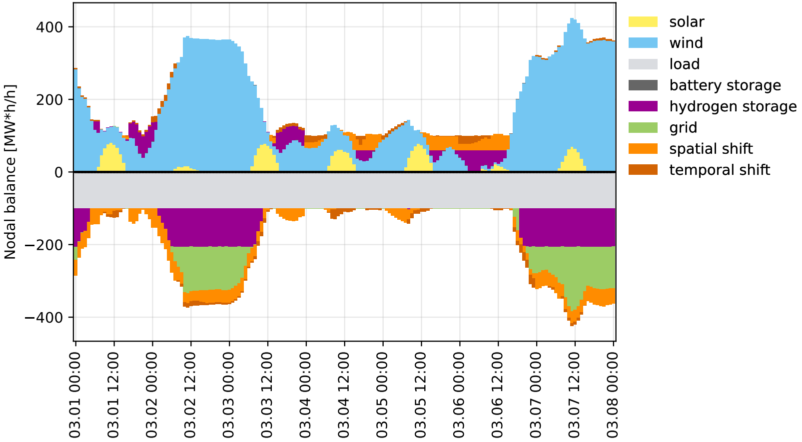

Data center CFE supply and demand

The plot on the right shows the **nodal energy balance** [MW*h/h], i.e., the (cost-optimal) matching of data center consumption with carbon-free energy supply.

Data center in Denmark.
The first week of March.
0% of flexible workloads.
100% CFE score.
+ LDES (palette 2).