

Spatio-temporal load shifting for truly clean computing

Iegor Riepin^{a,1}, Tom Brown^a, Victor M. Zavala^b

^aDepartment of Digital Transformation in Energy Systems, TU Berlin, Germany

^bDepartment of Chemical and Biological Engineering, University of Wisconsin-Madison, USA
Mathematics and Computer Science Division, Argonne National Laboratory, USA

Abstract

Companies with datacenters are procuring significant amounts of renewable energy to reduce their carbon footprint. There is increasing interest in achieving 24/7 Carbon-Free Energy (CFE) matching in electricity usage, aiming to eliminate all carbon footprints associated with electricity consumption on an hourly basis. However, the variability of renewable energy resources poses significant challenges for achieving this goal. We explore the impact of shifting computing jobs and associated power loads both in time and between datacenter locations. We develop an optimization model to simulate a network of geographically distributed datacenters managed by a company leveraging spatio-temporal load flexibility to achieve 24/7 CFE matching. We isolate three signals relevant for informed use of load flexibility: varying average quality of renewable energy resources, low correlation between wind power generation over long distances due to different weather conditions, and lags in solar radiation peak due to Earth’s rotation. We illustrate that the location of datacenters and the time of year affect which signal drives an effective load-shaping strategy. The energy procurement and load-shifting decisions based on informed use of these signals facilitate the resource-efficiency and cost-effectiveness of clean computing—the costs of 24/7 CFE are reduced by 1.29 ± 0.07 €/MWh for every additional percentage of flexible load. We provide practical guidelines on how companies with datacenters can leverage spatio-temporal load flexibility for truly clean computing. Our results and the open-source optimization model can also be useful for a broader variety of companies with flexible loads and an interest in eliminating their carbon footprint.

Keywords: 24/7 carbon-free electricity, Demand side management, Energy modelling, Renewable energy

Contents

1 Introduction	1	3.3 Signal 3: time lag in solar radiation peaks due to Earth’s rotation	8
2 Methods and experimental design	3	3.4 Generalising the results beyond specific load locations	10
2.1 Methods and tools	3	4 Discussion	12
2.2 Datacenter model	4	5 Conclusion	13
2.3 Model scope and parametrisation	6	Acronyms	14
3 Results	7	References	15
3.1 Signal 1: quality of local renewable resources	7	S1 Technology assumptions	17
3.2 Signal 2: low correlation of wind power generation over long distances	8		

Email addresses: iegor.riepin@tu-berlin.de (Iegor Riepin), t.brown@tu-berlin.de (Tom Brown), victor.zavala@wisc.edu (Victor M. Zavala)

¹Corresponding author: Iegor Riepin, TU Berlin, Einsteinufer 25, 10587 Berlin, Germany

1. Introduction

Demand for digital services is growing rapidly. The number of internet users has doubled worldwide over the past

decade, while traffic has grown 25-fold [1]. The significant energy consumption of Information and Communications Technology (ICT) infrastructure is becoming a concern from an ecological standpoint. Global datacenter energy use is estimated to be between 1% and 2% of final electricity demand worldwide in 2020 [2, 3] and is likely to increase rapidly in the future [4]. The Greenhouse Gas (GHG) emissions footprint of the datacenters and data transmission networks was estimated to be 330 million tonnes (Mt) CO₂-eq in 2020, equivalent to 0.9% of global energy-related emissions [1, 5]. For perspective, the national CO₂ emissions of the United Kingdom were around 331 Mt in 2021 [6].

The growing demand for computing drives innovation and economies of scale. Large companies, such as Amazon, Google, and Microsoft are centralizing computing infrastructure into “hyperscale” datacenters. Such large and highly efficient facilities consist of hundreds of server nodes that are managed collectively and geographically dispersed throughout the world [7].

In their role as electricity consumers, hyperscale datacenters have unique characteristics when it comes to demand-side management. Some computing jobs and associated power loads are “flexible”, i.e., the jobs are not time-critical and can be scheduled flexibly without impacting the overall quality of service. Therefore, datacenter operators have the ability to shift a considerable fraction of computing workloads *in time* and *in space* [8]. The temporal load shifting implies re-scheduling of flexible workloads to another time point, i.e., delaying computing job execution. The spatial load shifting implies migration of flexible compute jobs and associated power loads between different physical datacenter locations.

The significant energy consumption of ICT infrastructure, and its anticipated growth, raise a natural question: *How can datacenters make use of their unique spatial and temporal load-shifting capabilities to improve their economics and the environmental footprint?* The topic of this overarching question has been the subject of a growing body of literature over the past decade.

Existing literature – The work by Wierman et al. [9] surveyed the forms, opportunities and challenges for data center demand response options, including aspects of temporal load shifting such as workload delaying or shedding and the geographical load balancing. Several papers discussed methods and policies needed to incentivise participation of datacenters in the power grid’s demand response [10, 11].

Numerous studies investigated ways to harness spatial or temporal load-shifting flexibility to achieve certain economical or environmental goals. Many studies have proposed workload scheduling algorithms for geographical load shifting to reduce datacenter electricity costs [12–17], or looked at the potential of datacenter load migration for mitigating renewable curtailment and GHG emissions in the local electricity grid [18, 19], or aimed to increase the share of renewable energy in datacenter electricity consumption [20–23].

Recently, several papers elaborated a mathematical problem that captures both spatial and temporal load-shifting flexibility. Zhang et al. [24] introduced the concept of virtual links to capture space-time load flexibility provided by geographically-distributed data centers in market clearing procedures. This work was followed by a paper in which the authors demonstrated a market clearing formulation that seeks to remunerate spatio-temporal load shifting for the flexibility service datacenters can offer to the power grid [25].

A great deal of attention in previous research exploring the signals for geographical load shifting has focused on information characterising the state of electricity grids where datacenters operate. These signals include average carbon emissions in the region of operation [18], or locational marginal carbon emissions [26], or locational electricity prices and price differences across datacenter locations [25, 27, 28]. This emerging stream of research work is now finding practical applications. For example, Google recently introduced a Carbon-Intelligent Compute Management system that minimizes carbon footprint and power infrastructure costs by shifting flexible workloads in time across datacenter fleet as a function to the next day’s carbon intensity forecasts [8]. There are several companies providing market data and forecasts on carbon emissions and electricity prices, such as WattTime [29], ElectricityMaps [30], and others.

24/7 Carbon-Free Energy matching – ICT companies invest significantly in renewable energy to reduce their environmental impact, reduce power price volatility, and enhance their brand image. The hyperscale datacenter industry in particular leads in corporate renewable energy acquisition through Power Purchase Agreements (PPAs). With almost 50 GW contracted so far, Amazon, Microsoft, Meta, and Google are the four largest purchasers of corporate renewable energy PPAs [1]. Furthermore, some companies have pledged to eliminate *all* greenhouse gas emissions associated with their electricity use. These com-

panies have set a “24/7 Carbon-Free Energy (CFE)” goal aiming to match electricity demand with clean energy supply on an *hourly basis*. The 24/7 CFE targets were announced by Google and Microsoft for 2030 and Iron Mountain for 2040 [31–33].

One notable feature of the 24/7 CFE goal is that it drives—and requires—a significant volume of energy procurement via *PPAs*. Recent research showed that achieving 24/7 CFE with a 100% target (i.e., achieving 0 gCO₂ emissions per kWh) implies that companies need to rely on their own procured resources, without any electricity procurement from the local grid [34, 35]. This is because the local electricity mix in most regions and at most times does not have a strictly zero carbon content and therefore cannot contribute to 24/7 CFE.

The problem – Companies aiming for 24/7 clean electricity have a unique challenge regarding load-shifting. Because they rely on their own portfolio of carbon-free energy generators instead of buying electricity from a local grid, grid signals such as average carbon emissions or locational electricity prices no longer play a significant role for informed load-shifting. This effect was illustrated with the help of energy system modelling in the previous work of the authors [34, 35]. In this context, the following question arise: *In pursuit of achieving 24/7 Carbon-Free Energy objectives, what signals should datacenter operators focus on to facilitate informed and effective load shifting across space and time?*

Contribution – This is the first paper to examine the role of space-time load shifting in the context of 24/7 carbon-free energy matching. Through energy system modelling, we identify signals for effective shifting of load across space and time. We show that taking these signals into consideration in load shifting and energy procurement strategies can enable companies to achieve significant gains in energy efficiency and affordability of truly clean computing.

For that, we develop a computer model that simulates a datacenter fleet with controlled degrees of flexibility in pursuit of 24/7 CFE objectives. Datacenter fleet optimization is incorporated into a European electricity system model as a mathematical problem. We identify three individual signals that are relevant for effective geographical and temporal load shifting. These signals include (a) varying quality of renewable energy resources (i.e. average capacity factors) across datacenter locations, (b) low correlation between wind power generation over long distances due to different weather conditions, and (c) lags

in solar radiation peak due to Earth’s rotation. The results show that each of these signals plays a role in effective load shifting, although the “weight” of each signal in optimal load shaping strategies depends on the locations of datacenters and time of the year. Furthermore, we demonstrate that load shifting decisions optimized considering these signals facilitate efficiency and affordability of the 24/7 CFE matching. When all three signals are taken into account, the space-time load shifting can reduce the electricity costs of a datacenter fleet operating at 0 kgCO₂/MWh by 1.29 ± 0.07 €/MWh for each percent point of flexible load.

The remainder of the paper is structured as follows: [Section 2](#) introduces a mathematical model of clean energy procurement and gives a brief summary of the methodology, sources of model input data, and the experimental setup. [Section 3](#) presents the results of the experiments and discusses the implications of the findings. [Section 4](#) discusses the results in the context of the existing literature, practical recommendations for incorporating discussed signals into load shaping strategies and gives critical appraisal of the results. [Section 5](#) concludes the paper. All code necessary to reproduce the experiments is published under an open license alongside the paper. See [Data and Code Availability](#).

2. Methods and experimental design

2.1. Methods and tools

The simulations were carried out with the open-source software PyPSA for energy system modelling [36] and Linopy for optimization [37]. The mathematical model is a system-wide cost-minimisation problem. The objective of the model is to co-optimize (i) investment and dispatch decisions of generation and storage assets done by datacenters to meet their electricity demand in line with 24/7 CFE objectives, (ii) space-time load-shifting decisions subject to datacenter flexibility constraints, as well as (iii) investment and dispatch decisions of assets in the rest of the European electricity system to meet the demand of other consumers. The model formulation includes the linear optimal power flow approximation on the transmission network. This work uses a brownfield investment approach, which means that the model includes information about the existing assets of the European electricity system.

The electricity system dispatch and investment problem of this type is a standard in the energy modelling literature [38]. The clean computing model is based on the

framework of 24/7 CFE accounting suggested by Google [39] that was implemented in realms of energy system models by Xu et al. [40] and Riepin & Brown [34]. The space-time load flexibility model is inspired by the work of Zhang et al. [25]. Finally, the space-time load flexibility as a degree of freedom within the 24/7 CFE matching problem is a novel contribution of this work.

2.2. Datacenter model

Datacenters are represented as a subset of demand nodes for electricity with a fixed demand profile, controlled degree of flexibility, and a set of specific constraints ensuring the carbon-free energy matching and utilisation of flexibility in a feasible domain.

The 24/7 CFE matching problem introduces a set of new model components (parameters, variables, constraints) into the power system optimisation problem to represent voluntary yet binding carbon-free energy matching commitments. Here, we model a situation where one company operates a network of datacenters that are geographically distributed but managed collectively. The company is committed to *match every kilowatt-hour of electricity consumption by carbon-free sources in all datacenter locations*. To achieve the target, the company optimises its procurement of carbon-free generation and storage resources, dispatch of procured assets, imports of electricity from the regional grid, spatio-temporal load-shifting flexibility use (schematically illustrated in Figure 1), and when necessary, curtailment of excess generation.

A single datacenter with an inflexible load – First, consider a case of a single datacenter with demand profile d_t for each hour t of the year. The demand profile is assumed to be known in advance and is not flexible. In such case, energy balance constraint requires that hourly

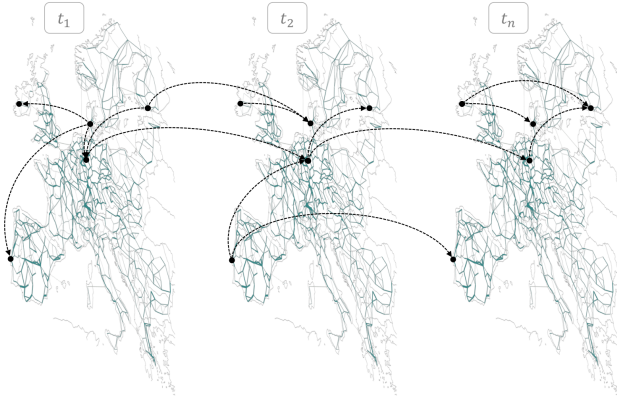


Figure 1: Illustration of spatio-temporal load-shifting optimisation problem faced by a datacenter operator.

demand is met by a combination of the following: (i) dispatch $g_{r,t}$ of procured carbon-free generators $r \in CFE$, (ii) dispatch $\bar{g}_{s,t}$ of procured storage technologies $s \in STO$, and (iii) imports of electricity from the regional grid im_t :

$$\sum_{r \in CFE} g_{r,t} + \sum_{s \in STO} (\bar{g}_{s,t} - g_{s,t}) - ex_t + im_t = d_t \quad \forall t \in T \quad (1)$$

Note that if total electricity yield of renewable generators procured by the datacenter operator exceeds demand in a given hour, the excess electricity ex_t has to be either stored or curtailed.²

The 24/7 CFE matching constraint requires that sum over hourly generation from the contracted generators, plus net dispatch of the storage technologies, plus imports of electricity from the regional grid multiplied by the grid's hourly CFE factor CFE_t minus the excess electricity must be higher or equal than a certain CFE score x multiplied by the total load of a datacenter:

$$\sum_{r \in CFE, t \in T} g_{r,t} + \sum_{s \in STO, t \in T} (\bar{g}_{s,t} - g_{s,t}) - \sum_{t \in T} ex_t + \sum_{t \in T} CFE_t \cdot im_t \geq x \cdot \sum_{t \in T} d_t \quad (2)$$

where *CFE score* x [%] measures the degree to which hourly electricity consumption is matched with carbon-free electricity generation [39]. Thus, equation (2) allows for controlling *the quality score* of the 24/7 CFE procurement by adjusting the parameter x , which was the subject of research in recent literature [34, 40]. Here, we focus on the best quality score ($x = 1$) ensuring that every kilowatt-hour of electricity consumption is met by carbon-free sources at all times.

Note the following properties of the 24/7 CFE matching problem:

1. The contracted generators $r \in CFE$ must be additional to the system and can be sited only in the local bidding zone, known as the requirements for *additionality* and *locational matching*.
2. Excess electricity is not counted toward the CFE score and thus subtracted from the left-hand side of eq. (2).

²In practice, excess electricity can also be sold to the regional electricity market at wholesale market prices. This option is deliberately avoided in this work. When included, optimal usage of flexibility factors in potential revenues from selling excess electricity to the grid.

3. The CFE factor of the regional grid (CFE_t) can be seen as the percentage of clean electricity in each MWh of imported electricity to supply demand of participating consumers in a given hour. To compute CFE_t , we consider both the hourly electricity mix in the local bidding zone and emission intensity of imported electricity from the neighbouring zones. The numerical simulations show that for the perfect 24/7 CFE matching ($x = 1$), the datacenter operator does not rely on electricity imports from the regional grid because local electricity mix must have a strictly zero carbon content for imported electricity to be counted as “carbon-free”. Thus, the datacenter operator has one less degree of freedom to meet the 24/7 CFE matching target. The methodology to calculate the grid CFE factor and linearise resulting nonconvexity is described in detail in the prior work of the authors [35].

Multiple datacenters with flexible load – Here we expand the 24/7 CFE matching problem to a case of multiple datacenters with flexible loads.

Let us assume that a known number of computing jobs with their associated power demand are “flexible”, i.e., power demand can potentially be shifted geographically across datacenter locations, or delayed to other times.³ Thus, the *dispatched load* \tilde{d}_t of a datacenter can deviate from the requested load d_t . The dispatched load \tilde{d}_t can take a value in a feasible domain that is constrained by a datacenter capacity (an upper limit) and the inflexible workloads volume (a lower limit), as illustrated in Figure 2. The range of possible deviations between the dispatched and the requested loads is assumed to lie within a certain *flexibility range* f [%], such as:

$$[1 - f] \cdot d_t \leq \tilde{d}_t \leq [1 + f] \cdot d_t \quad \forall t \in T \quad (3a)$$

$$\tilde{d}_t = d_t + (\bar{\Delta}_t - \underline{\Delta}_t) \quad \forall t \in T \quad (3b)$$

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

To incorporate spatial flexibility, we use a concept of *virtual links* representing non-physical pathways for shifting power loads across datacenter locations. Assume we have a set of datacenters N located in various locations of an electricity network. Let Θ be the set of all virtual links and let $\delta_{\vartheta} \in \mathbb{R}_+$ be the spatial load shifts (flows via virtual

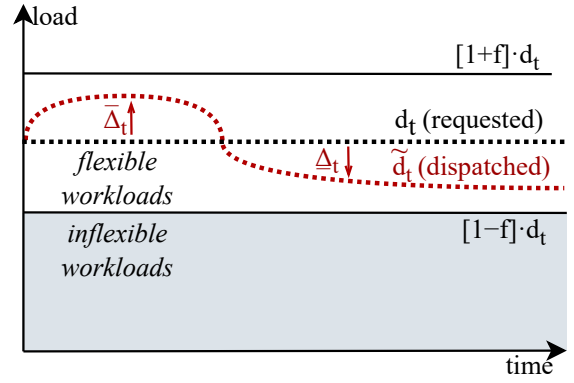


Figure 2: Illustration of load flexibility concept: the absolute value of deviations between dispatched \tilde{d}_t [MW] and requested d_t [MW] loads must fall within a certain flexibility range f [%] for each hour t .

pathways). We can thus 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 for each datacenter $n \in N$.

To incorporate temporal flexibility, we use a concept of *temporal load management* that allows for shifting load from one time point to another in the future.⁴ Let each datacenter have a temporal load management mechanism $S_n^{dsm} := \{s' \in S | dsm(s') = n\}$. If $T = \{t_1, t_2, \dots, t_T\}$ is a time horizon of our optimization problem, we can define $\bar{g}_{s',n,t}, g_{s',n,t} \in \mathbb{R}_+$ to be the amount of load shifted from time t to another time point t' for each datacenter $n \in N$ and $t, t' \in T$.

The nodal energy balance for inflexible consumers (eq. 1) can now be extended by variables representing shifts of load *in space* and *in time*. Thus, the dispatched load \tilde{d}_t of flexible consumer can deviate from the $d_{n,t}$ value due to space-time load shifting.

$$\sum_{r \in CFE} g_{r,n,t} + \sum_{s \in STO} (\bar{g}_{s,n,t} - g_{s,n,t}) - ex_{n,t} + im_{n,t} = d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} + \sum_{s' \in S_n^{dsm}} (\bar{g}_{s',n,t} - g_{s',n,t}) \quad (4)$$

$$\forall n \in N, t \in T$$

Computing capacity constraints (eq. 5) ensure that the dispatched load at each datacenter $\tilde{d}_{n,t}$ falls within the feasible domain considering both spatial and temporal load shifting. Computing capacity limit sets an upper

³The amount of CPU usage on a cluster can be mapped accurately to electricity demand [8].

⁴Zhang et al. [25] show that the temporal load shifting problem can also be formulated via the concept of *virtual links*. This yields the same mathematical problem.

bound (eq. 5b) and inflexible workloads share sets a lower bound (eq. 5c).

$$\begin{aligned} \tilde{d}_{n,t} &= d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} \\ &+ \sum_{s' \in S_n^{dsm}} \left(\tilde{g}_{s',n,t} - g_{s',n,t} \right) \quad \forall n \in N, t \in T \end{aligned} \quad (5a)$$

$$\tilde{d}_{n,t} \leq [1 + f] \cdot d_{n,t} \quad \forall n \in N, t \in T \quad (5b)$$

$$\tilde{d}_{n,t} \geq [1 - f] \cdot d_{n,t} \quad \forall n \in N, t \in T \quad (5c)$$

Note that 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) as defined in eq. (5).

Temporal load shifts have an additional constraint ensuring that the cluster-level compute usage is preserved within a certain time interval. We follow the work of Radovanovic et al. [8] and implement a *daily* compute usage conservation rule:

$$\sum_{t|t \in t(DAYS)} \left(\tilde{g}_{s',t} - g_{s',t} \right) = 0 \quad \forall s' \in S_n^{dsm} \quad (6)$$

Finally, the 24/7 CFE matching constraint defined in eq. 2 can be defined over a set of datacenters nodes N and extended on the right-hand side with variables representing spatial and temporal load shifts:

$$\begin{aligned} &\sum_{r \in CFE, t \in T} g_{r,n,t} + \sum_{s \in STO, t \in T} \left(\tilde{g}_{s,n,t} - g_{s,n,t} \right) \\ &- \sum_{t \in T} ex_{n,t} + \sum_{t \in T} CFE_{n,t} \cdot im_{n,t} \geq \\ &x_n \cdot \left(\sum_{t \in T} \left(d_{n,t} + \sum_{\vartheta \in \Theta_n^{rec}} \delta_{\vartheta,t} - \sum_{\vartheta \in \Theta_n^{snd}} \delta_{\vartheta,t} \right. \right. \\ &\left. \left. + \sum_{s' \in S_n^{dsm}} \left(\tilde{g}_{s',n,t} - g_{s',n,t} \right) \right) \right) \quad \forall n \in N \end{aligned} \quad (7)$$

Overall, the mathematical formulation of the 24/7 CFE matching problem extended with eqs. 4-7 represents a situation where a company pursuing the 24/7 CFE matching target has two additional degrees of freedom—the ability to shift power loads in space and in time—as part of its decision-making process. Co-optimisation of space-time load-shifting flexibility with procurement and dispatch of carbon-free generators can lead to a more resource-efficient and cost-effective solution for carbon-free electricity matching.

2.3. Model scope and parametrisation

Geographical focus of the model covers the entire European electricity system 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). A set of datacenters—electricity consumers committed to 24/7 CFE goals—are placed in several selected zones in each scenario. Each zone have unique weather patterns, renewable potentials, and background electricity mixes. The model's time horizon spans one year. All model runs are done with 1-hourly temporal resolution, i.e., 8760 time steps.

We assume all datacenters to have a load of 100 MW and a flat (baseload) consumption profile. Note that the shape of a consumption profile does not significantly impact the costs of the 24/7 CFE matching [35]. We assume a *complete graph* superstructure of virtual links. Specifically, for every pair of distinct datacenters $n, m \in N$ where $n \neq m$, there exists a bidirectional virtual link $\vartheta_{nm} \in \Theta$ such that $snd(\vartheta_{nm}) = n$ and $rec(\vartheta_{nm}) = m$, and conversely, $snd(\vartheta_{mn}) = m$ and $rec(\vartheta_{mn}) = n$. This configuration ensures that each datacenter is capable of both sending and receiving spatial load shifts to and from every other datacenter in the network.

The model is parametrised for the exemplary year 2025. Input data such as technology cost assumptions, national renewable policies, decommissioning of legacy power plant fleet, and system-wide assumptions are parametrised accordingly. For the 24/7 CFE matching problem, we assume a set of energy technologies that *are commercially available today*: onshore wind, utility scale solar photovoltaics (PV), and battery storage. Cost and other assumptions for these technologies are collected from the Danish Energy Agency [41] and provided in [Supplementary Information](#).

Weather-dependent power potentials for renewable energy technologies, including those of solar PV and onshore wind generators available for the 24/7 CFE matching, are simulated with hourly temporal resolution using the ERA5 reanalysis dataset [42] using the open-source tool Atlite [43]. We assume 2013 to be a representative climate year.

Other electricity system data is processed with an open-source model for the European energy system PyPSA-Eur [44]. All code and input data to reproduce the experiments are available at the GitHub repository under open licenses [45].

3. Results

We present the results as a sequence of case studies of varying levels of complexity, since there are several interacting signals at work. In the first three subsections, we isolate three particular signals that affect the load-shifting. In the final section, we generalize the results across any combination of datacenters.

3.1. Signal 1: quality of local renewable resources

First, we focus on the signal associated with the difference in quality of local renewable resources. Consider a case when datacenters are located in Denmark, Germany, and Portugal (Figure 3: panel a). There are three different types of renewable resources represented by the three sites, with Denmark having a good wind resource, Germany having a moderate wind and solar resources, and Portugal having a good solar resource. The quality of local renewable resources is reflected by the average annual capacity factor for onshore wind and solar PV, as shown in the panels b and c.

Figure 3: panel d shows the modelled cost-optimal capacity of renewable generators and battery storage for 24/7 matching. In the case of inflexible loads, the procurement strategy must ensure a sufficient amount of carbon-free electricity is available round-the-clock *at each site*. Due to the non-dispatchable nature of renewable resources and the relatively high energy costs of battery storage, the cost-optimal portfolio is much larger than the load. To cover the 100 MW load with perfect hourly matching, the consumer would have to procure combined wind and solar PV generators of approx. 900 MW in Denmark and 1300 MW in Germany or Portugal. In each portfolio, the proportion of wind and solar PV reflects the quality of the local resources.

In Figure 3: panel d, steps on the x-axis represent increasing flexible load, allowing for a reduction in the capacity of renewable generators and battery storage necessary for 24/7 matching. By having a degree of load flexibility, a consumer can shift the load from times of scarce renewable generation to times of abundant renewable generation in the same hour, or use the excess renewable generation at one site to cover the load at another site.

The breakdown of costs associated with each procurement strategy is shown in Figure 3: panel e. The panel shows the average cost per MWh of consumption, including renewable generation costs and battery storage

costs.⁵ Demand flexibility enables consumers to better match renewable generation with demand profile, reducing over-procurement and costly battery storage. The costs are reduced in all locations, and especially in locations where hourly matching is most expensive. Modelling results show the costs of 24/7 matching in Germany are at 215 €/MWh when there is no load flexibility, 195 €/MWh when there is 10% flexible loads, and reduced up to 137 €/MWh when there is 40% of flexible loads.

Figure 3: panel f displays the total annual costs for achieving 24/7 matching at all locations (left y-axis) and their relative representation as a percentage of the inflexible scenario costs (right y-axis). This panel summarises the model results: as demand flexibility increases, 24/7 matching becomes more resource-efficient and affordable. The value of additional flexibility also diminishes as the number of flexible loads increases, i.e., the first 10% of flexible loads can reduce costs more than the next 10%. A generalization of this result is presented in Section 3.4.

Finally, Figure 3: panels h-m show optimal hourly load shifts for the three datacenter locations based on a selected scenario with 40% of flexible loads. Several observations can be drawn from the plots regarding the use of spatial and temporal flexibilities.

First, there are clear daily and seasonal patterns of spatial load shifts (panels h, j, l) driven by the quality of local resources, such as wind or solar PV capacity factors. For renewable generators, better resource quality implies higher energy yield per MW of installed capacity, which translates into lower Levelised Costs of Electricity (LCOE). Whenever spatial load shifting is possible, a rational strategy is *to get the most out of the clean energy resources with good quality*. The heatmaps above illustrate this behavior well: a datacenter located in Denmark—a region with poor solar resources—tends to shift loads away from the mid-spring till mid-autumn. Instead, a datacenter located in Portugal—a region with better solar resources—tends to receive loads during this period. It works just about reciprocally for wind-related load shifts: a datacenter in Portugal benefits from having a partner in Denmark, the very windy region in Europe. Germany's datacenter is located between the two extremes, receiving workloads from Denmark during summer daytime and moving them to other locations at other times.

⁵Here we do not consider the option of selling the excess electricity to the regional grid. Excess electricity can thus be stored in the battery, shifted in space or time with load flexibility, or curtailed.

Second, temporal load shifts (panels i, k, m) is used to “flatten” the renewable generation profile. The heatmaps show that the load is consistently shifted to the daytime hours, when solar generation is abundant, and away from the nighttime hours. Note that the temporal flexibility is subject to the daily usage conservation constraint (eq. 6), which limits the temporal flexibility usage for longer periods.

3.2. Signal 2: low correlation of wind power generation over long distances

Second, we focus on the signal associated with low correlation over long distances in wind power generation. Suppose a pair of datacenters are located so that their local wind resources are high quality, but their solar resources are low quality, with increasing distance between them. Taking five pairs as an example, we can consider Ireland and Northern Ireland, England, the Netherlands, and Denmark (Figure 4: panel a).

We can use historical weather data to calculate the Pearson correlation coefficient of the hourly capacity factors for onshore wind generation. Figure 3: panels b and c display the result for the selected regions in Ireland and Denmark, accordingly. There is a drastic falloff in wind feed-in correlation over distances of 300-400 km due to varying weather conditions. This calculation replicates observations in a previous study on the statistical properties of geographically dispersed wind power [46]. It follows that the feed-in of wind power generators procured, for instance, by datacenters in Ireland and Denmark is not correlated.

By using spatial flexibility, datacenter operators can do “load arbitrage” between locations with different weather conditions to take advantage of wind generation’s stochastic nature. In the heatmap of spatial load shifts for selected datacenter pair (Figure 3: panel d), the geographical load shifts resulting from uncorrelated wind feed-in can be spotted by vertical stripes (sudden change of direction) with stochastic patterns. Note that there is no strong seasonal pattern, as it was observed in the previous example, since both locations share similar quality of wind and solar resources.

Figure 3: panel e illustrates the relationship between the distance between datacenter pairs and the resultant cost savings. The results show that the cost savings are most noticeable for steps of 300-400 km, which is consistent with the falloff in wind generation correlation. As the load is more flexible, more load arbitrage is possible, re-

sulting in cost savings; if the load is inflexible, the datacenter pair cannot harvest the benefits of different local weather conditions. In distances greater than 300-400 km, wind feed-in correlations are already low, which leads to no significant cost savings. Distances over 1000 km and high flexibility levels result in a slight increase in cost savings due to the time lag between solar generation peak times. The effect is visible on the heatmap of spatial load shifts. We discuss this signal in detail in the next section.

3.3. Signal 3: time lag in solar radiation peaks due to Earth’s rotation

Third, we focus on the signal associated with the time lag in solar radiation peaks due to Earth’s rotation. Consider a case when datacenters are located in Denmark, Greece, and Portugal (Figure 5: panel a). In this scenario, Danish datacenter has access to high-quality wind resources, whereas Greek and Portuguese datacenters have access to good solar resources.

The Pearson correlation of the simulated hourly capacity factors for solar photovoltaic generation (Figure 5: panels b and c) show that solar generation has good correlation over long distances, in contrast to wind. However, the datacenters in Greece and Portugal are approx. 2700 km apart, which creates potential for the use of spatial flexibility to take advantage of the time lag in solar generation peaks. The difference in solar photovoltaic hourly capacity factors between the two selected locations is shown in Figure 5: panel d. The panel illustrates the time lag in solar generation peaks due to Earth’s rotation. The peak of solar generation in Portugal is delayed by several hours compared to Greece. For comparison, the difference in wind generation hourly capacity factors between the two selected locations is shown in Figure 5: panel e. As expected, low correlation of wind feed-in over long distance results in a stochastic pattern.

Panels f, g and h in Figure 5 show the cost-optimal portfolio of renewable resources and battery storage for 24/7 matching, cost breakdown, and total annual costs of 24/7 matching strategy as a function of load flexibility. Similar to the previous examples, the cost-optimal portfolio of resources necessary for 24/7 matching and associated costs are reduced with increasing flexible loads, as datacenters can get the most out of the clean energy resources with good quality, as well as take advantage of the time lag in solar generation peaks. In the case of the Danish datacenter, benefiting from having partners in solar-rich locations, the costs of 24/7 matching strategy is re-

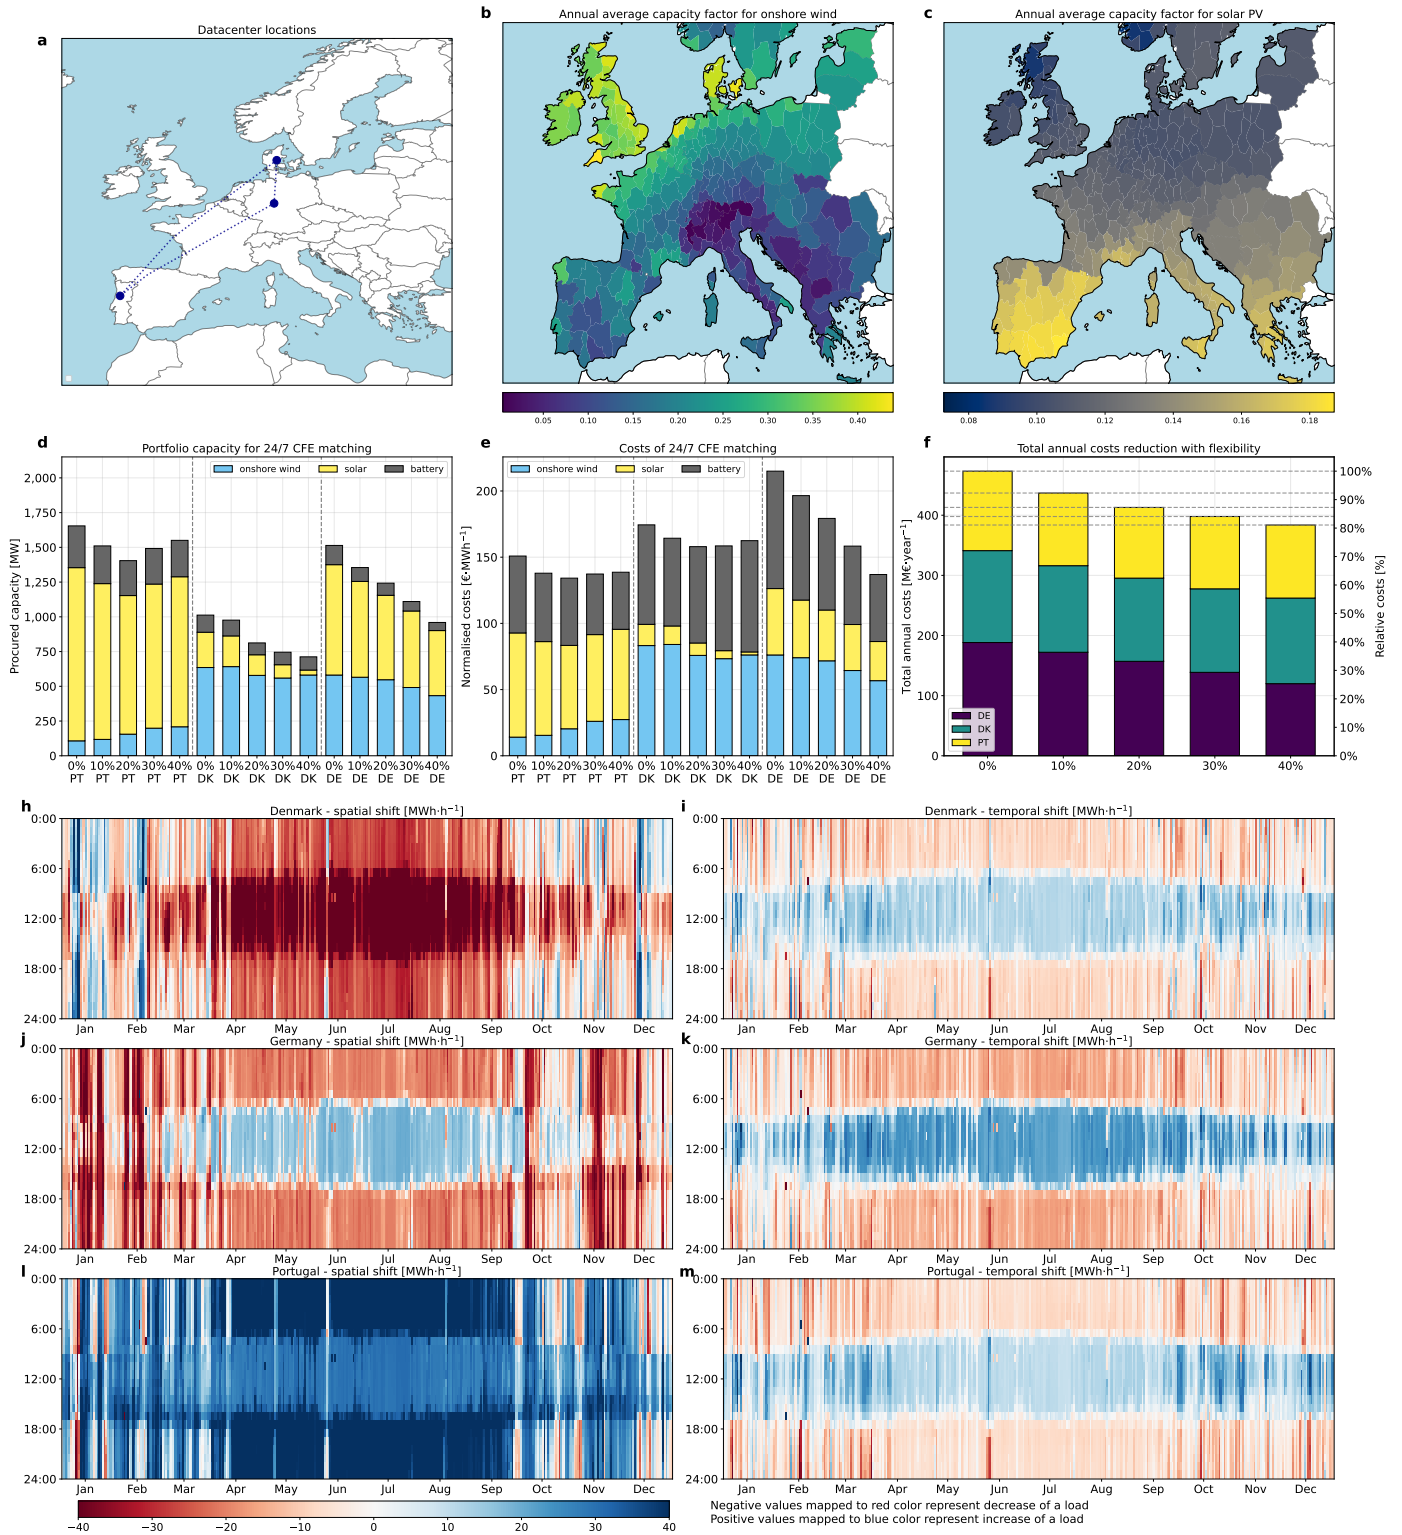


Figure 3: Illustration of the signal 1: quality of local renewable resources. **a**: Assumed datacenter locations: Denmark, Germany, Portugal. **b,c**: Annual average capacity factor of onshore wind and solar photovoltaic. Data is simulated using the ERA5 reanalysis dataset for weather year 2013 and aggregated to 256 regions in Europe. **d**: Cost-optimal portfolio of renewable resources and battery storage sufficient for 24/7 matching. Steps on x-axis represent increasing share of flexible load. **e**: Cost breakdown of 24/7 matching strategy. **f**: Total annual costs of 24/7 matching strategy as a function load flexibility. Relative axis is normalized to the costs of inflexible load. **h,j,l**: Hourly spatial load shifts for the three datacenter locations. Color mapping represents the quantity of load “received” from other locations (positive values) or “sent” away (negative values). **i,k,m**: Hourly temporal load shifts for the three datacenter locations. Color mapping represents the quantity of load shifted to a given hour from other times, or shifted from a given hour to another time.

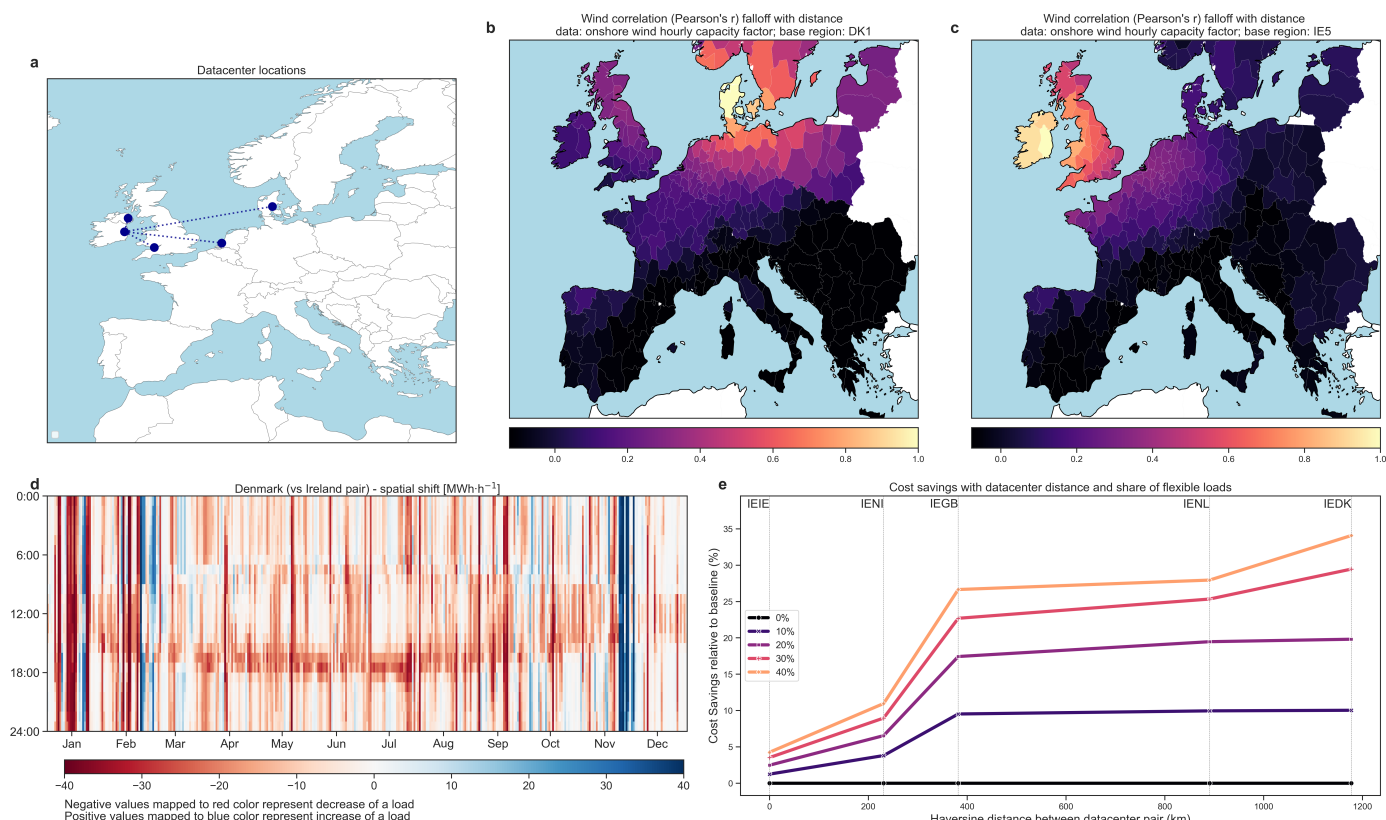


Figure 4: Illustration of the signal 2: low correlation of wind power generation over long distances. **a**: Assumed datacenter locations: pairwise connections across regions with similar quality of renewable resources. Here: Ireland with Northern Ireland, or England, or the Netherlands, or Denmark (-west zone). **b,c**: Pearson correlation of hourly capacity factor for onshore wind generation, if Denmark (-west zone, panel b) or selected region of Ireland (panel c) are taken as basis. As a result of different weather conditions, wind feed-in has a noticeable correlation falloff over distances of 300-400 km. Data is simulated using the ERA5 reanalysis dataset for weather year 2013 and aggregated to 256 regions in Europe. **d**: Hourly spatial load shifts for the selected scenario and datacenter; here datacenter is located Denmark (and another one is located Ireland). Color mapping represents the quantity of load “received” from other locations (positive values) or “sent” away (negative values). **e**: Cost savings of 24/7 matching with increasing distance between datacenters. Costs are normalized to the cost level of inflexible load.

duced from 174 €/MWh to 157 €/MWh with 10% of flexible loads, and down to 106 €/MWh with 40% of flexible loads.

Finally, Figure 5: panels i and j show optimal hourly load shifts for the selected datacenter locations in Greece and Portugal, accordingly. The heatmaps illustrate how spatial flexibility can be used to benefit from the time lag in solar generation peaks. The datacenters in Greece and Portugal receive loads from Denmark between midspring and midfall, with a clear daily pattern caused by the time lag in solar generation peaks. Additionally, the heatmaps show that during the winter months, when solar generation is scarce, both datacenters tend to shift loads to Denmark.

3.4. Generalising the results beyond specific load locations

Finally, we generalise the results beyond specific load locations and to provide a more comprehensive understanding of the relationship between load flexibility and the

costs of 24/7 CFE matching.

Consider that datacenters can be located in any of the following eight countries: Denmark, Ireland, the Netherlands, Germany, Latvia, Greece, Portugal, and France. We selected these countries to represent a variety of renewable resources and geographical locations.⁶ The analysis is carried out by modeling all combinations of three datacenters in the selected countries. With three countries selected from eight, and the order of the countries being irrelevant, there are 56 possible combinations. As a result, the scenario space includes combinations of datacenters with different quality of renewable resources, as well as different distances between them. For each combination of datacenters and level of flexibility, we con-

⁶The analysis could be expanded to include additional locations, but the number of combinations would be significantly increased. These locations are a reasonable compromise between generalisation quality and computational demands.

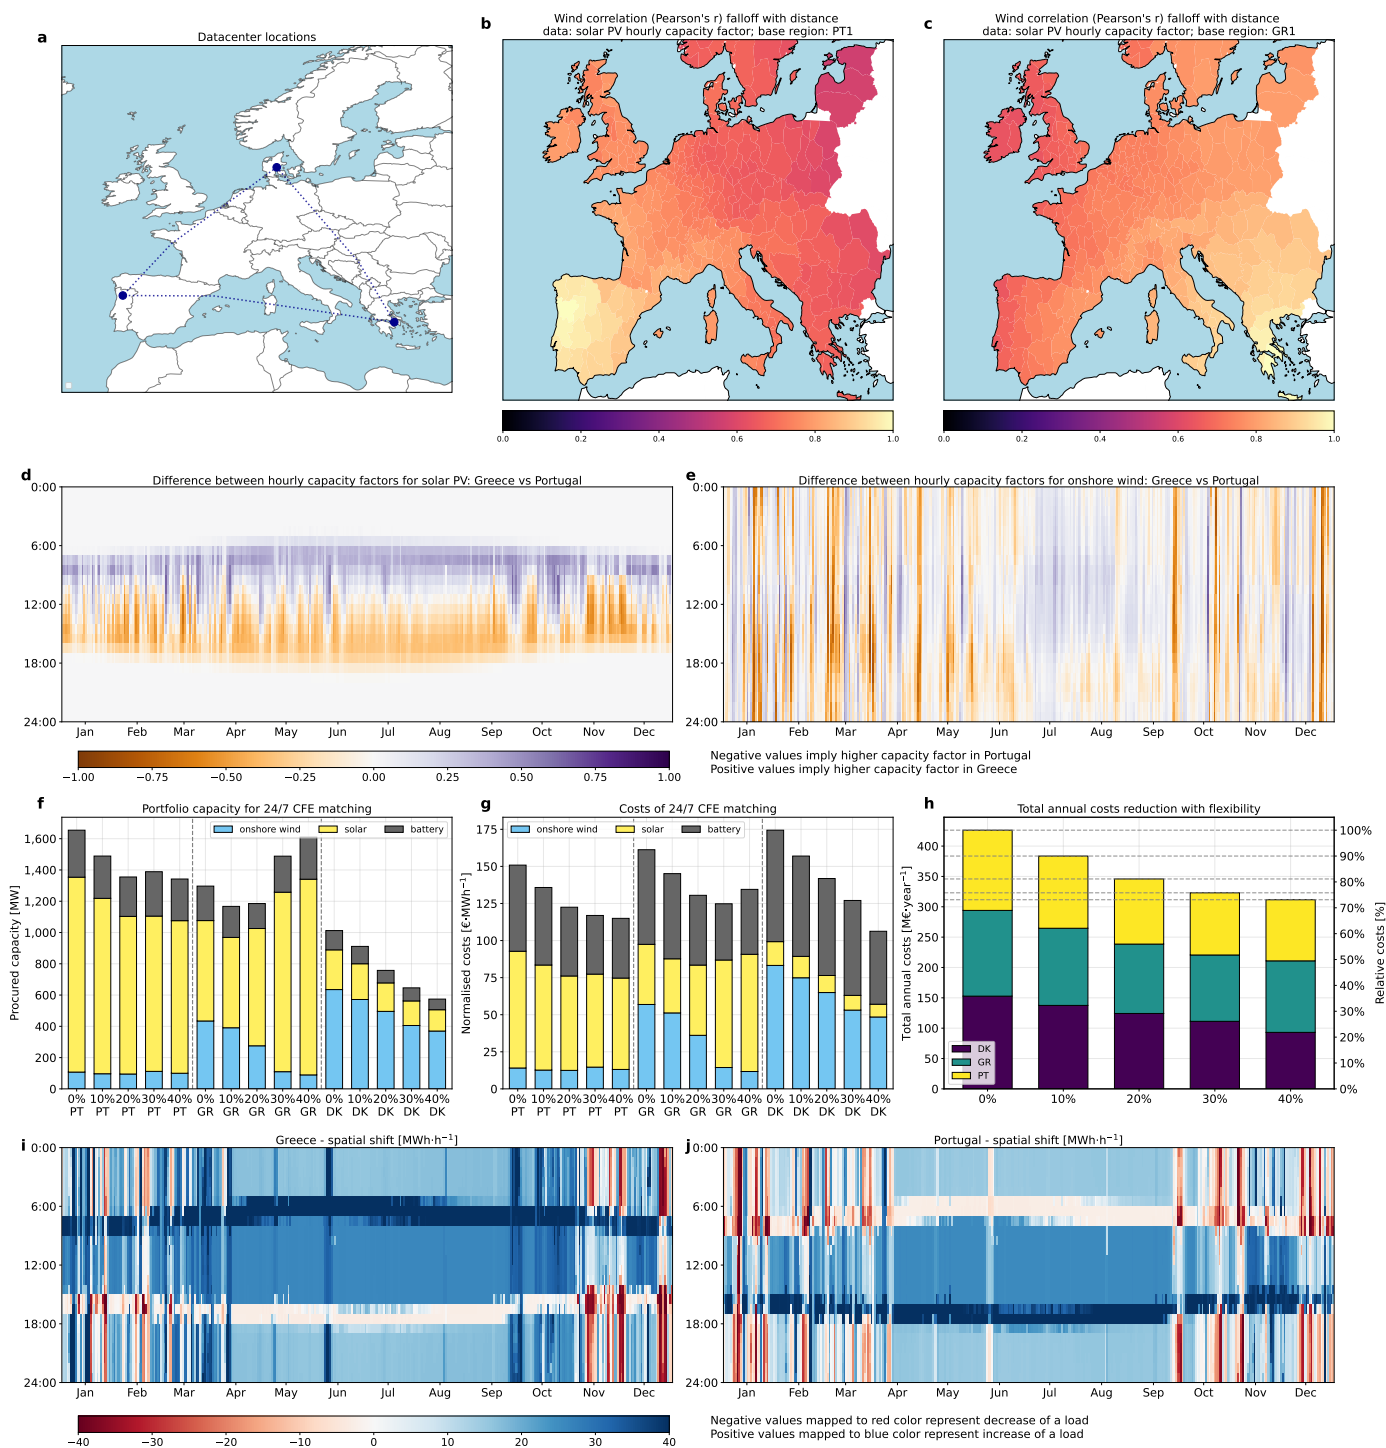


Figure 5: Illustration of the signal 3: time lag in solar radiation peaks due to Earth's rotation. **a**: Assumed datacenter locations: Denmark, Portugal, Greece. **b,c**: Pearson correlation of hourly capacity factor for solar photovoltaic generation, if selected region of Portugal (panel b) or region of Greece (panel c) is taken as basis. Solar generation remains highly correlated over long distances, in contrast to wind generation. **d**: Difference in solar photovoltaic hourly capacity factors between two selected locations: Greece and Portugal. The two locations are approx. 2700 km apart, which results in a noticeable lag in solar generation peaks due to Earth's rotation. **e**: Difference in wind generation hourly capacity factors between two selected locations: Greece and Portugal. As expected, low correlation of wind feed-in over long distance results in stochastic pattern. **f**: Cost-optimal portfolio of renewable resources and battery storage sufficient for 24/7 matching. Steps on x-axis represent increasing share of flexible load. **g**: Cost breakdown of 24/7 matching strategy. **h**: Total annual costs of 24/7 matching strategy as a function load flexibility. Relative axis is normalized to the costs of inflexible load. **i,j**: Hourly spatial load shifts for the selected datacenter locations: Greece (panel i) and Portugal (panel j). Color mapping represents the quantity of load "received" from other locations (positive values) or "sent" away (negative values).

duct the same analysis as in the previous sections, i.e., the model solution includes information about the cost-optimal portfolio of renewable resources and battery storage for 24/7 matching, procurement policy cost breakdowns, and optimal utilisation of spatial and temporal load flexibility.

The results of the analysis are presented in Figure 6. The figure shows the relationship between the level of flexibility and the costs of 24/7 matching for all 56 combinations of datacenters. The normalised costs of 24/7 matching lie in a broad range from 240 €/MWh to ca. 160 €/MWh for inflexible load scenario. The lower bound represents combination Denmark-Greece-Portugal, since these locations share excellent quality of wind and solar resources. The upper bound represents a combination of Germany, Ireland, and Latvia, since none of the three locations have a good solar resource, and only Ireland has a good wind resource. As a result, 24/7 CFE matching is especially expensive in two out of three locations – Germany and Latvia.

Several conclusions can be drawn from the analysis. First, space-time load-shifting flexibility facilitates the efficiency and affordability of 24/7 CFE matching, regardless of the location of the datacenters. Based on specific datacenter locations, the optimal use of space-time flexibility may be driven by one, or a combination, of the signals discussed above; however, in general, flexibility is beneficial to all locations.

Second, the costs of 24/7 CFE matching are reduced by 1.29 ± 0.07 €/MWh for every additional percentage of flexible load. The low standard error suggests that the relationship is robust across different combinations of datacenters. This finding implies that datacenters can make use of the available flexibility to reduce the costs of 24/7 CFE matching, regardless of which of the three signals is dominant for the specific locations. For example, if renewable energy quality is similar across all locations, and solar PV has a high LCOE, the two related signals have little impact on optimal space-time shifting and 24/7 CFE procurement strategy. However, when this occurs, the company operating the datacenters can use all the available load flexibility to “arbitrage” on low correlation of wind feed-in, and factor this into the procurement strategy.

Third, the overlap between the mean trend and linear regression line suggests that the value of additional flexibility diminishes as the number of flexible loads increases, i.e., the first 10% of flexible loads can reduce costs more

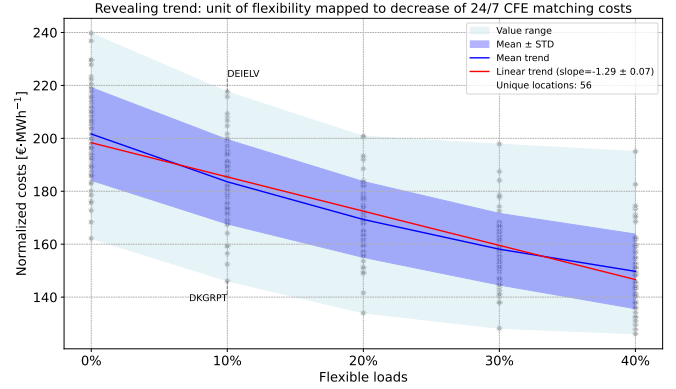


Figure 6: Relationship between the level of flexibility and the costs of 24/7 matching for 56 combinations of datacenter locations. The costs are normalised to the cost level of inflexible load.

than the next 10% of flexible loads. The latter is consistent with the observations made in previous sections.

4. Discussion

Our results in perspective – This work contributes to and builds on two streams of literature: one is concerned with the flexibility and environmental impact of datacenters, the another with the means, costs and implications of 24/7 carbon-free energy matching.

Regarding the first literature stream, our findings are in line with those of existing model-based studies on the flexibility potentials offered by geographically distributed datacenters: load-shifting flexibility facilitates the integration of renewable energy sources and reduces the carbon footprint of datacenters [18–23, 26]. Although opportunities and challenges of datacenter flexibility have drawn considerable attention in the literature, space-time load-shifting flexibility has not been studied in the context of 24/7 CFE matching pursued by companies aiming to run their facilities at $0 \text{ gCO}_2/\text{kWh}$. Since hourly matching of consumption with carbon-free electricity requires and implies renewable energy procurement through PPAs and thus less electricity procurement from the grid, we argue that grid signals such as average carbon emission intensity or locational electricity prices have low value for informed use of load-shifting flexibility in the context of the 24/7 CFE matching. In order to address this challenge, we identify three signals that can enable companies to shift load across space and time effectively. Moreover, we demonstrate how companies can achieve significant gains in energy efficiency and affordability of 24/7 CFE matching by including these signals into their energy procurement and load-shaping strategies.

Regarding the second literature stream, our results indicate that space-time load-shifting increases access to clean electricity and provides flexibility to consumers for matching demand with carbon-free electricity. As a result, datacenters and other commercial and industrial consumers who have flexible demands can achieve high degrees of carbon-free energy matching with lower resources and thus be more cost-effective. The system decarbonization impact of 24/7 CFE procurement indicated by the model-based studies [35, 40] could therefore be amplified with greater participation.

Finally, our results follow the general consensus on the benefits of geographical load balancing in highly renewable electricity networks [47]. However, the benefits of connecting dispersed regions at continental scales are always weighed against the high costs and long lead times associated with grid expansion and reinforcement. It is unique to the datacenter industry that geographical load-shifting flexibility can be achieved at low costs and over long distances by harnessing local differences in the quality of renewable resources and renewable generation profiles.

A broader context – The first commitments to 24/7 CFE matching initiated by the datacenter industry can help companies from other sectors by establishing procurement methods and standard practices for 24/7 CFE matching. [40]. In the same way, early efforts to consider potentials of available load flexibility within energy procurement and load-shifting strategies can help a broad range of companies that are seeking to lower their carbon footprints. In the event that datacenter companies, as well as other commercial and industrial consumers with flexible loads can achieve high degrees of 24/7 CFE matching at moderate costs, more actors may be willing to commit to it. Therefore, the system decarbonization impacts associated with the 24/7 CFE procurement [34] could be amplified with greater participation while requiring fewer resources.

Additionally, even though it is not the focus of the present study, our results provide necessary methodological foundation for studying strategic locations for new datacenters by companies interested to reduce or eliminate completely their carbon footprint. It is therefore important to consider the factors facilitating resource-efficient and cost-effective 24/7 CFE matching identified by this study when choosing optimal locations for datacenters.

Critical appraisal and further work – The study design includes simplifications and assumptions inherent to

energy optimization models. The results should thus be seen with caution, i.e., as model-derived insights rather than quantitative projections.

One of the central assumptions of the model experiment is that datacenters achieve a certain degree of flexibility in their workloads. While we control the share of flexible workloads exogenously (see [Methods and experimental design](#)), the actual flexibility of datacenters is not known with certainty and can vary significantly among facilities. As specific data becomes available, the model can be calibrated to reflect the actual flexibility potential.

Additional empirical research is required to quantify the costs and benefits of utilizing demand flexibility in the ICT industry. 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. By including implicit flexibility costs, the benefits of flexibility can be quantified more accurately. An empirical improvement could also address technical aspects and properties of flexible workloads, such as power usage ramping up and down, reliability, and performance constraints.

In the context of 24/7 CFE matching, further research is needed to capture the interaction of promising energy technologies, such as long-term duration energy storage and space-time load-shifting flexibility. This case is particularly relevant, since 24/7 CFE procurement could create an early market and drive deployment of advanced technologies [34, 40]. It was also demonstrated that geographical load shifting complements long-term duration storage [48].

5. Conclusion

This work explores the role of space-time load-shifting flexibility provided by geographically distributed datacenters in achieving 24/7 carbon-free energy (CFE) matching and identifies signals relevant for effective load shifting.

We develop an optimization model to simulate a network of datacenters managed collectively by a company pursuing 24/7 carbon-free energy matching objective. The model provides users with the ability to specify datacenter locations within the European electricity system, along with the flexibility potentials, among other parameters. Through energy system modeling, we isolate three individual signals that are crucial for effective space-time load-shifting: (i) varying average quality of renewable

energy resources across datacenter locations, (ii) low correlation between wind power generation over long distances due to different weather conditions, and (iii) lags in solar radiation peak due to Earth’s rotation.

In relation to the first signal, we demonstrate that leveraging spatio-temporal load flexibility maximizes the utility of high-quality renewable energy resources. Therefore, the optimal load-shaping strategy entails making informed decisions about where and when to shift loads, based on the average capacity factors of procured renewable energy resources at different locations. Our results indicate that energy costs are reduced at all locations, particularly in those where hourly matching is most challenging. For instance, for a data center located in Germany with partners in Denmark and Portugal, the costs of 24/7 CFE matching decrease from 215 €/MWh with no load flexibility to 195 €/MWh with 10% flexible loads, and further to 137 €/MWh with 40% flexible loads.

As for the second signal, we show that spatio-temporal load shifting allows datacenter companies to take advantage of the low correlation between wind power generation over long distances, i.e., to perform “load arbitrage” between locations with different weather conditions. The results show that the cost savings are most noticeable for steps of 300-400 km, which is in line with the falloff in wind generation correlation. Considering a pair of datacenters located approximately 380 km apart in regions with similar wind and solar conditions, the introduction of 10% flexible loads yields a 9.5% reduction in energy costs, whereas 40% flexible loads facilitate a 26.5% cost savings. For distances greater than 400 km, wind feed-in correlations are already low, so further increasing the distance between datacenters does not lead to additional cost savings.

Regarding the third signal, we show how spatio-temporal load shifting can be used to exploit the lags in solar radiation peak due to Earth’s rotation. Thus, companies seeking 24/7 CFE matching can benefit from the time zone differences between datacenters and the hourly capacity factors of solar power generators at different locations. Solar power can be used more efficiently this way, thereby reducing costs associated with carbon-free energy matching.

Finally, we generalize our results to a broader set of datacenter locations across Europe and show the relationship between the load flexibility and the costs of 24/7 CFE matching. The results reveal that the resource-efficiency and cost-effectiveness of hourly matching are improved

by the effective use of space-time load flexibility in all cases. However, the location of datacenters and the time of year affect which of the three signals are most relevant for an effective load-shaping strategy. Further, we show that with every percentage of flexible load added, the costs of 24/7 CFE matching fall on average by 1.29 ± 0.07 €/MWh, which can be interpreted as the marginal value of load flexibility. Our results also illustrate the diminishing returns of additional load flexibility, which is important for companies seeking to co-optimize their long-term energy procurement and short-term load-shifting strategies.

Overall, our results indicate that space-time load-shifting flexibility increases access to clean electricity and gives consumers more options to match demand with carbon-free electricity. Demand flexibility can enhance resource efficiency and reduce costs of truly clean computing. This can amplify the decarbonization effects associated with the voluntary clean energy commitments by encouraging more datacenter companies and other electricity consumers to join the 24/7 CFE movement.

Acknowledgements

We thank the following people for support and fruitful discussions on the art of energy system modelling: Elisabeth Zeyen, Fabian Hofmann, Fabian Neumann, Devon Swezey, Ana Radovanovic, Chris Adams, Kristina Riepin and the research staff of the Scalable Systems Laboratory at the University of Wisconsin-Madison.

Author Contributions

Igor Riepin: Conceptualization – Data curation – Formal Analysis – Investigation – Methodology – Software – Validation – Visualization – Writing - original draft
Tom Brown: Conceptualization – Methodology – Formal Analysis – Software – Supervision – Writing - review & editing
Victor M. Zavala: Conceptualization – Methodology – Formal Analysis – Writing - review & editing

Data and Code Availability

The simulations were carried out with PyPSA – an open-source software framework for simulating and optimizing modern energy systems [49]. All code, input data and results are published under an open license. The code to reproduce the experiments is available at GitHub [45].

Acronyms

CFE Carbon-Free Electricity

GHG Greenhouse gas

ICT Information and Communications Technology

LCOE Levelised Cost of Electricity

PPA Power Purchase Agreement

PV Solar photovoltaics

References

- [1] IEA, *Data Centres and Data Transmission Networks* (Jul. 2023). URL <https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks>
- [2] D. Mytton, *How much energy do data centers use?* (Feb. 2020). URL <https://davidmytton.blog/how-much-energy-do-data-centers-use/>
- [3] E. Masanet, A. Shehabi, N. Lei, S. Smith, J. Koomey, *Recalibrating global data center energy-use estimates*, *Science* 367 (6481) (2020) 984–986. doi:10.1126/science.aba3758. URL <https://www.science.org/doi/10.1126/science.aba3758>
- [4] A. S. G. Andrae, T. Edler, *On Global Electricity Usage of Communication Technology: Trends to 2030*, *Challenges* 6 (1) (2015) 117–157. doi:10.3390/challe6010117. URL <https://www.mdpi.com/2078-1547/6/1/117>
- [5] J. Malmödin, N. Lövehagen, P. Bergmark, D. Lundén, *ICT Sector Electricity Consumption and Greenhouse Gas Emissions – 2020 Outcome* (Apr. 2023). doi:10.2139/ssrn.4424264. URL <https://papers.ssrn.com/abstract=4424264>
- [6] U. office for National Statistics, *Greenhouse gas emissions, UK Statistical bulletins - Office for National Statistics* (Mar. 2023). URL https://assets.publishing.service.gov.uk/media/6424b8b83d885d000fdade9b/2022_Provisional_emissions_statistics_report.pdf
- [7] Y. Sverdlik, *Analysts: There are Now More than 500 Hyperscale Data Centers in the World* (Oct. 2019). URL <https://www.datacenterknowledge.com/cloud/analysts-the-re-are-now-more-500-hyperscale-data-centers-world>
- [8] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care, S. Talukdar, E. Mullen, K. Smith, M. Cottman, W. Cirne, *Carbon-Aware Computing for Datacenters*, *IEEE Transactions on Power Systems* 38 (2) (2023) 1270–1280. doi:10.1109/tpwrs.2022.3173250. URL <https://ieeexplore.ieee.org/document/9770383>
- [9] A. Wierman, Z. Liu, I. Liu, H. Mohsenian-Rad, *Opportunities and challenges for data center demand response*, in: *International Green Computing Conference, 2014*, pp. 1–10. doi:10.1109/igcc.2014.7039172. URL <https://ieeexplore.ieee.org/document/7039172>
- [10] Z. Liu, I. Liu, S. Low, A. Wierman, *Pricing data center demand response*, in: *The 2014 ACM International Conference on Measurement and Modeling of Computer Systems, Sigmetrics '14, Association for Computing Machinery, New York, NY, USA, 2014*, pp. 111–123. doi:10.1145/2591971.2592004. URL <https://doi.org/10.1145/2591971.2592004>
- [11] Z. Zhou, F. Liu, S. Chen, Z. Li, *A Truthful and Efficient Incentive Mechanism for Demand Response in Green Datacenters*, *IEEE Transactions on Parallel and Distributed Systems* 31 (1) (2020) 1–15. doi:10.1109/tpds.2018.2882174. URL <https://ieeexplore.ieee.org/document/8540414>
- [12] L. Velasco, A. Asensio, J. Ll. Berral, E. Bonetto, F. Musumeci, V. López, *Elastic operations in federated datacenters for performance and cost optimization*, *Computer Communications* 50 (2014) 142–151. doi:10.1016/j.comcom.2014.03.004. URL <https://www.sciencedirect.com/science/article/pii/S0140366414000905>
- [13] H. Dou, Y. Qi, W. Wei, H. Song, *Carbon-Aware Electricity Cost Minimization for Sustainable Data Centers*, *IEEE Transactions on Sustainable Computing* 2 (2) (2017) 211–223. doi:10.1109/tsusc.2017.2711925. URL <https://ieeexplore.ieee.org/document/7938380>
- [14] H. He, H. Shen, *Minimizing the operation cost of distributed green data centers with energy storage under carbon capping*, *Journal of Computer and System Sciences* 118 (2021) 28–52. doi:10.1016/j.jcss.2020.11.004. URL <https://www.sciencedirect.com/science/article/pii/S002200020301112>
- [15] L. Rao, X. Liu, M. D. Ilic, J. Liu, *Distributed Coordination of Internet Data Centers Under Multiregional Electricity Markets*, *Proceedings of the IEEE* 100 (1) (2012) 269–282. doi:10.1109/jproc.2011.2161236. URL <https://ieeexplore.ieee.org/document/5989839>
- [16] C. Ren, D. Wang, B. Uргаonkar, A. Sivasubramaniam, *Carbon-Aware Energy Capacity Planning for Datacenters*, in: *2012 IEEE 20th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 2012*, pp. 391–400. doi:10.1109/mascots.2012.51. URL <https://ieeexplore.ieee.org/document/6298199>
- [17] X. Deng, D. Wu, J. Shen, J. He, *Eco-Aware Online Power Management and Load Scheduling for Green Cloud Datacenters*, *IEEE Systems Journal* 10 (1) (2016) 78–87. doi:10.1109/jsyst.2014.2344028. URL <https://ieeexplore.ieee.org/document/6877627>
- [18] J. Zheng, A. A. Chien, S. Suh, *Mitigating Curtailment and Carbon Emissions through Load Migration between Data Centers*, *Joule* 4 (10) (2020) 2208–2222. doi:10.1016/j.joule.2020.08.001. URL <https://www.sciencedirect.com/science/article/pii/S2542435120303470>
- [19] A. H. Mahmud, S. S. Iyengar, *A Distributed Framework for Carbon and Cost Aware Geographical Job Scheduling in a Hybrid Data Center Infrastructure*, in: *2016 IEEE International Conference on Autonomic Computing (ICAC), 2016*, pp. 75–84. doi:10.1109/icac.2016.21. URL <https://ieeexplore.ieee.org/document/7573118>
- [20] X. Wang, Z. Du, Y. Chen, M. Yang, *A green-aware virtual machine migration strategy for sustainable datacenter powered by renewable energy*, *Simulation Modelling Practice and Theory* 58 (2015) 3–14. doi:10.1016/j.simpat.2015.01.005. URL <https://www.sciencedirect.com/science/article/pii/S1569190X15000155>
- [21] K. Kim, F. Yang, V. M. Zavala, A. A. Chien, *Data Centers as Dispatchable Loads to Harness Stranded Power*, *IEEE Transactions on Sustainable Energy* 8 (1) (2017) 208–218. doi:10.1109/tste.2016.2593607. URL <https://ieeexplore.ieee.org/document/7517380>

- [22] Z. Liu, M. Lin, A. Wierman, S. H. Low, L. L. Andrew, [Geographical load balancing with renewables](#), *ACM SIGMETRICS Performance Evaluation Review* 39 (3) (2011) 62–66. doi:10.1145/2160803.2160862. URL <https://doi.org/10.1145/2160803.2160862>
- [23] C. E. Kelly, J. A. Ging, A. Kansal, M. P. Walsh, [Balancing Power Systems With Datacenters Using a Virtual Interconnector](#), *IEEE Power and Energy Technology Systems Journal* 3 (2) (2016) 51–59. doi:10.1109/jpets.2016.2519611. URL <https://ieeexplore.ieee.org/document/7452537>
- [24] W. Zhang, L. A. Roald, A. A. Chien, J. R. Birge, V. M. Zavala, [Flexibility from networks of data centers: A market clearing formulation with virtual links](#), *Electric Power Systems Research* 189 (2020) 106723. doi:10.1016/j.epsr.2020.106723. URL <https://www.sciencedirect.com/science/article/pii/S0378779620305265>
- [25] W. Zhang, V. M. Zavala, [Remunerating space-time, load-shifting flexibility from data centers in electricity markets](#), *Applied Energy* 326 (2022) 119930. doi:10.1016/j.apenergy.2022.119930. URL <https://www.sciencedirect.com/science/article/pii/S0306261922011874>
- [26] J. Lindberg, L. Roald, B. Lesieutre, [The Environmental Potential of Hyper-Scale Data Centers: Using Locational Marginal CO₂ Emissions to Guide Geographical Load Shifting](#), 2021. URL <http://hdl.handle.net/10125/70999>
- [27] L. Rao, X. Liu, L. Xie, W. Liu, [Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment](#), in: *2010 Proceedings IEEE INFOCOM*, 2010, pp. 1–9. doi:10.1109/infcom.2010.5461933. URL <https://ieeexplore.ieee.org/document/5461933>
- [28] N. H. Tran, D. H. Tran, S. Ren, Z. Han, E.-N. Huh, C. S. Hong, [How Geo-Distributed Data Centers Do Demand Response: A Game-Theoretic Approach](#), *IEEE Transactions on Smart Grid* 7 (2) (2016) 937–947. doi:10.1109/tsg.2015.2421286. URL <https://ieeexplore.ieee.org/document/7102766>
- [29] [WattTime](#). URL <https://watttime.org/>
- [30] [Electricity Maps](#). URL <https://www.electricitymaps.com/>
- [31] G. LLC, [24/7 by 2030: Realizing a Carbon-free Future](#) (2020). URL <https://www.gstatic.com/gumdrop/sustainability/247-carbon-free-energy.pdf>
- [32] Microsoft, [Made to measure: Sustainability commitment progress and updates](#) (Jul. 2021). URL <https://blogs.microsoft.com/blog/2021/07/14/made-to-measure-sustainability-commitment-progress-and-updates/>
- [33] [Iron Mountain: Sustainability Goals](#). URL <https://www.ironmountain.com/about-us/sustainability>
- [34] I. Riepin, T. Brown, [On the means, costs, and system-level impacts of 24/7 carbon-free energy procurement](#) (Dec. 2023). doi:10.5281/zenodo.10407830. URL <https://zenodo.org/records/10407830>
- [35] I. Riepin, T. Brown, [System-level impacts of 24/7 carbon-free electricity procurement in Europe](#), Tech. rep., Zenodo (Oct. 2022). doi:10.5281/zenodo.7180098. URL <https://zenodo.org/record/7180098>
- [36] J. Hörsch, F. Hofmann, D. Schlachtberger, T. Brown, [PyPSA-Eur: An open optimisation model of the European transmission system](#), *Energy Strategy Reviews* 22 (2018) 207–215. doi:10.1016/j.esr.2018.08.012. URL <https://www.sciencedirect.com/science/article/pii/S2211467X18300804>
- [37] [Linopy | Linear optimization with N-D labeled arrays in Python](#) (Jan. 2024). URL <https://github.com/PyPSA/linopy>
- [38] [Open Models - wiki.openmod-initiative.org](#). URL https://wiki.openmod-initiative.org/wiki/Open_Models
- [39] G. LLC, [24/7 Carbon-Free Energy: Methodologies and Metrics](#) (2021). URL <https://www.gstatic.com/gumdrop/sustainability/24x7-carbon-free-energy-methodologies-metrics.pdf>
- [40] Q. Xu, A. Manocha, N. Patankar, J. D. Jenkins, [\(report\) System-level Impacts of 24/7 Carbon-free Electricity Procurement](#), Tech. rep. (Nov. 2021). URL <https://zenodo.org/record/6229426>
- [41] [Technology Data for Generation of Electricity and District Heating](#), Tech. rep., Danish Energy Agency (Aug. 2016). URL <https://ens.dk/en/our-services/projections-and-models/technology-data>
- [42] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, J.-N. Thépaut, [The ERA5 global reanalysis](#), *Quarterly Journal of the Royal Meteorological Society* 146 (730) (2020) 1999–2049. doi:10.1002/qj.3803. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803>
- [43] F. Hofmann, J. Hampp, F. Neumann, T. Brown, J. Hörsch, [Atlite: A Lightweight Python Package for Calculating Renewable Power Potentials and Time Series](#) (Jun. 2021). doi:10.21105/joss.03294. URL <https://github.com/PyPSA/atlite>
- [44] P. authors, [PyPSA-Eur documentation](#). URL <https://pypsa-eur.readthedocs.io/>
- [45] I. Riepin, T. Brown, [Code for space-time shifting optimization paper](#) (2023). URL <https://github.com/lrieo/space-time-optimization>
- [46] B. Hasche, [General statistics of geographically dispersed wind power](#), *Wind Energy* 13 (8) (2010) 773–784. doi:10.1002/we.397. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.397>
- [47] D. P. Schlachtberger, T. Brown, S. Schramm, M. Greiner, [The benefits of cooperation in a highly renewable European electricity network](#), *Energy* 134 (2017) 469–481. doi:10.1016/j.energy.2017.06.004. URL <https://www.sciencedirect.com/science/article/pii/S0360544217309969>
- [48] I. Riepin, T. Brown, [The value of space-time load-shifting flexibility for 24/7 carbon-free electricity procurement](#), Tech. rep., Zenodo (Jul. 2023). doi:10.5281/zenodo.8185850. URL <https://zenodo.org/records/8185850>
- [49] T. Brown, J. Hörsch, D. Schlachtberger, [PyPSA: Python for Power System Analysis](#), *Journal of Open Research Software* 6 (1) (2018) 4. doi:10.5334/jors.188. URL <https://openresearchsoftware.metajnl.com/articles/10.5334/jors.188>

Technology	Year	CAPEX (overnight cost)	FOM [%/year]	VOM [Eur/MWh]	Efficiency [per unit]	Lifetime [years]	Source
Utility solar PV	2025	612 €/kW	1.7	0.01	-	37.5	[S41]
Onshore wind	2025	1077 €/kW	1.2	0.015	-	28.5	[S41]
Battery storage	2025	187 €/kWh	-	-	-	22.5	[S41]
Battery inverter	2025	215 €/kW	0.3	-	0.96	10	[S41]

Notes: All costs are in 2020 euros; CAPEX = capital expenditure; FOM = fixed operations and maintenance costs; VOM = variable operations and maintenance costs.

Table S1: Technology assumptions.

Supplementary Information

S1. Technology assumptions

Cost and other assumptions for energy technologies available for 24/7 CFE participating consumers are collected from the Danish Energy Agency [S41]. These assumptions are listed in Table S1. A full list of technology assumptions, including the data for energy technologies in the background energy system, is available via the reproducible scientific workflow in the GitHub repository [S45].