# Leveraging workload flexibility and waste heat from data centers for renewable-dominated energy systems

**Sai Sudharshan Ravi[1], Jonas Schnidrig[2,3], Michaël Aklin[4], , François Marechal,[1,3]**

1 Industrial Process and Energy Systems Engineering Group, École Polytechnique Fédérale de Lausanne, Rue de l'Industrie 17, 1950 Sion, Switzerland

2 CIRAIG, Institute for Sustainable Energy, University of Applied Sciences Western Switzerland, Rue de l'Industrie 21, 1950 Sion, Switzerland

3 CIRAIG, 3333 Queen Mary Rd, Montréal, QC H3V 1A2, Canada

4 PASU Chair, College of Management of Technology, École Polytechnique Fédérale de Lausanne, 1015, Lausanne, Switzerland

**Abstract:**

Digital infrastructure is becoming a significant component of Europe's energy system. The expansion of data centres, driven by cloud computing and artificial intelligence, is creating concentrated electricity demand and recoverable low-temperature heat. At the same time, meeting Europe's climate and energy targets requires rapid decarbonisation and large-scale integration of renewable energy. Here we quantify pan-European energy-system configurations for 2050 using the EnergyScope optimisation framework, which co-optimises energy supply, conversion, storage, and end-uses under harmonised ICT demand assumptions. This study introduces, for the first time, spatially explicit ICT demand and workload flexibility into a multi-energy optimisation model through the modelling-to-generate-alternatives (MGA) approach. Across one hundred scenarios for the EU-27, the United Kingdom, and Switzerland, we assess trade-offs between system costs, greenhouse-gas emissions, renewable curtailment, and regional equity. Results show that data centres account for a substantial share of final electricity demand and represent a major source of urban waste heat. Coupling data centres to district-heating networks reduces primary energy use and emissions, while aligning computing workloads with solar and wind generation lowers curtailment and costs. These findings demonstrate that digitalisation exerts tangible and spatially uneven effects on the energy transition, underscoring the need to explicitly represent data-centre dynamics in decarbonisation planning.

**Keywords:**

*Energy Systems Modelling, Data Centers, District Heating, Renewable Curtailment, Workload flexibility*

# 1. Introduction

Climate and environmental goals in Europe require deep cuts in greenhouse gas emissions and rapid growth of renewables[1],[2]. Key strategies to achieve this maximum deployment of renewables include the introduction of monetary incentives and setting of ambititous national energy targets [3],[4], [5], yet their regional co-benefits and externalities remain unevenly distributed.. A significant drop in the capital cost of renewables has made them financially attractive, which has so far, driven the growth of renewable energy in Europe[6], [7]. This growth of renewables, apart from the intended consequences of reducing greenhouse emissions from the electricity mix, also has some unintened consequences. Beneficial consequences include, reduction in air pollution[8], decreased dependence on fossil fuels[9], [10], and new economic development[11], [12], [13]. While, the burdens include, increased system costs[14], [15], decreased property values[16], [17], adverse impact on ecosystem and wildlife[18], [19], with most of these impacts unevenly distributed across European regions[20], [21].

At the same time, digital services are expanding. Cloud computing, artificial intelligence, and streaming increase demand for computing and create a new class of large and steady energy users near cities and industry [22]. Data centers draw significant electricity and reject low-temperature heat that can be recovered[23], [24]. They are continuously active, they draw electricity to sustain digital services while releasing low-temperature heat that can serve nearby buildings and industries.This multi-faceted link with the urban energy system exemplifies how digital infrastructure and local energy systems can coexist in a mutually beneficial exchange. A whole energy system view is needed to assess these interactions[25], [26], [27]. Electricity, heat, transport, and industrial services are connected through conversion and storage. Data centers sit at this junction. Their electricity demand has both real-time and deferrable components[28], [29].

Data-center electricity demand can be divided into real-time and deferrable components. The temporal flexibility of the latter enables alignment with renewable availability, while rejected heat provides a potential low-temperature supply to district-heating networks. These real-time loads encompass computing and cooling operations that must occur continuously to maintain service reliability such as web hosting, cloud transactions, and critical network functions. These processes operate under stringent uptime requirements and cannot be delayed without disrupting service quality. In contrast, deferrable loads include computational tasks that can be temporally shifted without immediate performance loss, such as data analytics, software updates, or large-scale AI training, which is growing in proportion. Estimates suggest that proportion of such deferrable workloads could be between 30-70% [30], [31], [32]. These flexible tasks can be strategically scheduled to align with periods of renewable energy abundance or low grid stress, allowing data centers to act as controllable loads that support system balancing and renewable integration[28], [33]. On the other hand, their waste heat can supply district heating where temperatures and networks allow it. Thus, the scale and timing of workloads and the degree of heat integration shape system outcomes.

A way to understand and investigate such interactions and outcomes would be to model spatially explicit scenarios of the information and communication technology (ICT) demands. However, despite the abundance of large-scale energy system models, there are very few that have studied the increasing growth of such ICT demands in conjunction with a multi-energy system model perspective, focusing on regional impacts and renewable penetration. Several studies have focussed on technical and economic aspects of renewable capacity, and the transition pathways models that show future outcomes firmly basing on current system states[34], [35], [36]. Some studies have focussed on regionally equitable allocation of renewable resources with Switzerland or Germany[37], [38], however, the ICT demand in such cases were not explicitly modelled. Even fewer studies have modelled the flexibility aspects associated with data centers in a system level, considering different generation profiles, however, the cross-sectoral coupling impacts of such data centers were not considered when flexibility was prioritised[39], [40], [41]. Consequently, the systemic implications of large-scale digitalization remain poorly quantified. While providing interesting insights and great first steps towards system possibilities, existing studies do not constitute a integrated spatially explicit techno-economic analysis on regional interactions between data centers and the electric and heating grids, by often neglecting transmission and storage capacities and renewable curtailment effects.

Here we quantify 2050 outcomes for EU-27 plus the United Kingdom and Switzerland using EnergyScope, a top-down, technology-rich optimization model. We extend the EnergyScope optimization model to explicitly represent ICT demand and data-centre coupling, generating 100 near-optimal system configurations via Modeling-to-Generate-Alternatives (MGA)[42], [43]. The model co-optimizes supply, conversion, storage, and end-uses across sectors including the ICT sector. We represent data centers explicitly with performance and operational choices. These include power usage effectiveness[44], cooling concepts, options for direct heat recovery and heat pumps, and the share of workloads that can shift in time. The scenario set varies by varying alignments across different system-level objectives and the degree of temporal alignment between deferrable workloads and variable renewables. We examine trade-offs across four objectives: minimizing total system costs, minimizing environmental impacts, maximizing renewable electricity generation, and minimizing regional inequality. We analyze patterns across these spatially explicit scenarios to highlight trade-offs and report total system costs, greenhouse gas emissions, and renewable curtailment, and we map how impacts differ across countries with different levels of hyperscale capacity. We show that these aims have often different outcomes that cannot be optimised simultaneously, with largely differing implementation pathways as has also been shown in the case of the MGA study conducted by Trutnevyte et. al.[45] Minimizing system costs leads to regionally concentrated impacts, while minimising regional inequality leads to higher system costs. These globals trends across objectives do seem to be consistent with the study on central european energy electricity system transition[45], however this analysis focuses on sector coupling with ICT demand integration and flexibility pathways, while employment, land use, and particulate emissions remain outside its scope, To our knowledge, this is the first pan-European study to embed data-centre flexibility and heat recovery within a whole-energy-system optimization framework.

# 2. Results

## 2.1. Monte Carlo projection of data demand and implied electricity use

To model the 2050 Europe-wide scenarios that include ICT demand, we first construct a transparent global projection of data consumption from 2018 to 2050 and derive the corresponding envelope of data-centre electricity use. The model assumes anchored exponential growth, capturing both the slowdown of Moore's law–type efficiency improvements, which raises cooling and power requirements, and the rapid expansion of AI-driven computing workloads. Uncertainty is represented multiplicatively to account for the compounding effects of efficiency and usage intensity over time, yielding a more realistic range of possible outcomes. A linear mapping between data traffic and electricity use is applied but can be substituted for alternative conversion functions. Let (A) be the last observed consumption in 2023, equal to 123 zettabytes. For each simulation (s), consumption in year (t) evolves as

$$D_S(t) = A.\exp\{\mu(t - t_0)\}.\varepsilon_s \tag{1}$$

with constant central growth ($\mu$) and a multiplicative factor ($\varepsilon_s$) that represents persistent deviation from the central path. To preserve the observation in 2023, every simulated trajectory is shifted to start at (A). We simulate one thousand trajectories on annual steps from 2023 to 2050.

We use a lognormal distribution for ($\varepsilon_s$) with zero mean on the log scale and a standard deviation ($\sigma$) of 0.20. This preserves positivity and produces right-skewed spreads that widen with time. The choice reflects two counteracting empirical forces. First, historical efficiency gains from hardware and cooling have slowed with the end of Moore's-law-like scaling. Second, usage growth from cloud, streaming, the internet of things, and the current AI cycle is accelerating. Together, these dynamics make demand projections highly sensitive and justify using a probabilistic formulation rather than a single deterministic trajectory.[46]

Each simulated data-demand path is mapped to electricity use with a constant factor of 4.5 TWh per zettabyte, consistent with efficiency factors derived from WEF reports[47].This communication-level factor allows direct use of the ensemble as exogenous demand in the energy-system model. Readers who prefer an efficiency trajectory that changes over time can substitute their own factor while keeping the same simulated data consumption ensemble.

Figure 1 highlights the accelerating growth and increasing uncertainty of global data demand. The ensemble of simulations shows a steep rise driven by the expansion of AI, cloud computing, and connected devices. By 2050, median data consumption reaches about 1,080 ZB, translating to roughly 4,850 TWh of electricity use. The widening spread from around 3,400 to 6,500 TWh, reflects divergent pathways of digitalisation and efficiency progress. This growing dispersion motivates the use of an uncertainty envelope for ICT demand in subsequent system-optimisation scenarios.

The specification is intentionally simple. It anchors to the most recent observation, it exposes two interpretable parameters that drive the spread of outcomes (the central growth rate and the multiplicative uncertainty), and it produces a complete table of inputs that others can reuse or adapt. The modelling choice is consistent with usage-driven outlooks that document the growing importance of behavioural and service-level drivers under a slowdown of device-level efficiency gains.

This table provides the complete set of values required to reproduce the projections and to adapt them for alternative assumptions about growth in the AI era or efficiency trends after the end of Moore-like scaling.

***Table 1:*** *Input data and assumptions used for the Monte-Carlo simulation*

| Item | Value | Notes and rationale |
|---|---|---|
| Observed years | 2018, 2019, 2020, 2021, 2022, 2023 | Input data used for anchoring[47]·[48]18/05/2026 17:39:00 |
| Observed consumption in 2023 | 123 ZB | Anchor at 2023 for all simulated trajectories[47], [48] |
| Time horizon | 2023 to 2050 | Annual time step |
| Central growth rate, ($\mu$) | 0.08 per year | Central usage growth consistent with strong cloud, streaming, IoT and AI demand under slower hardware-level efficiency gains.[46] |
| Multiplicative spread, ($\varepsilon_s$) | Lognormal with ($\mu_{log}$ = 0), ($\sigma_{log}$ = 0.20) | Captures persistent deviations such as faster adoption or slower efficiency; produces widening but bounded right-skewed uncertainty |
| Number of simulations | 1,000 | Provides stable quantiles for plotting and for downstream scenario selection |
| Electricity conversion | 4.5 TWh per ZB | Communication-level factor used in this study; can be replaced by an efficiency trajectory if desired[47] |
| Reporting years for panel (b) | 2030, 2035, 2040, 2045, 2050 | Five-year marks for box plots |

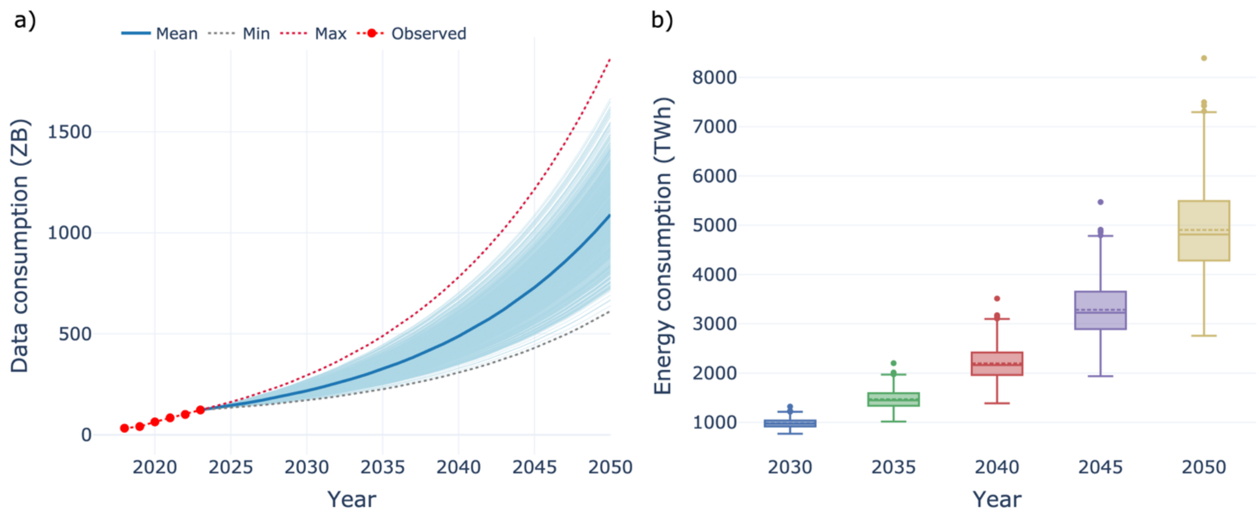

***Figure 1.*** *Monte Carlo projection of data consumption and induced energy use. a, Data consumption from 2023 to 2050 in zettabytes. Light blue lines are 1,000 anchored exponential simulations with 8 percent annual growth and lognormal multiplicative uncertainty with sigma equal to 0.2. The blue line is the mean trajectory. Grey and red dotted lines mark the minimum and maximum simulated paths and the historical observations for 2018 to 2023. b, Box and whisker plots of the implied data center energy use at five-year intervals using 4.5 TWh per zettabyte. Boxes show medians and interquartile ranges. Whiskers and points show the full distribution across simulations.*

## 2.2. Country-level envelopes for data demand and electricity use

We translate the global Monte Carlo ensemble projected to individual countries by scaling each simulated trajectory with country shares of European data-centre activity reported for 2022. These shares come from the Joint Research Centre assessment by Kamiya and Bertoldi[49]. We assume that, in the absence of strong relocation policies, the relative position of the main hubs in Europe remains broadly stable over the horizon considered. This assumption reflects the strong path dependence of data-centre siting, driven by pre-existing fiber connectivity, corporate clustering, and climatic constraints. The ensemble is then mapped to electricity with the factor of 4.5 terawatt-hours per zettabyte like before. Figure 2 displays five-year snapshots from 2030 to 2050 as a layered band between the minimum, mean, and maximum simulated data-demand paths estimated from the global demand projections. For each year, violin plots on the secondary axis show the distribution of implied electricity use. In Figure 2, we present six countries because they hold the highest market shares in Europe and together represent the bulk of European data-centre activity.

The envelopes differ across countries because they combine the common global uncertainty with each country's market share. Germany has the largest envelope. By 2050, the median data demand is about 35 zettabytes with an interquartile range (IQR) of about 20 to 58 zettabytes, which translates to a median electricity use near 160 terawatt-hours and an IQR of about 90 to 260 terawatt-hours. France follows with a 2050 median of about 23 zettabytes and an IQR of about 13 to 38 zettabytes, corresponding to about 103 terawatt-hours with an IQR of about 60 to 171 terawatt-hours. The United Kingdom reaches a 2050 median near 14 zettabytes and an IQR of about 9 to 24 zettabytes, or about 63 terawatt-hours with an IQR of about 40 to 110 terawatt-hours. The Netherlands has a similar level, with a 2050 median close to 13 zettabytes and an IQR of about 8 to 22 zettabytes, or about 60 terawatt-hours with an IQR of about 36 to 100 terawatt-hours. Ireland's 2050 median is near 12 zettabytes with an IQR of about 7 to 21 zettabytes, corresponding to about 54 terawatt-hours with an IQR of about 30 to 95 terawatt-hours. Italy shows a 2050 median near 10 zettabytes with an IQR of about 6 to 18 zettabytes, or about 45 terawatt-hours with an IQR of about 27 to 80 terawatt-hours. These values reflect both the widening of global demand uncertainty over time and the persistent concentration of activity in a few hubs.

The figure highlights two practical points. First, even within the same global demand scenario, national outcomes span wide ranges that matter for planning of grid connections, siting, and district-heating links. Second, the rank order of countries is largely driven by today's market shares, which suggests that policy choices on siting and interconnection can shift national exposure if they change the geography of European data-centre growth. These country-level envelopes provide the spatially disaggregated ICT-demand inputs for the EnergyScope-EU optimization, ensuring that modeled electricity and heat demands reflect realistic concentration patterns

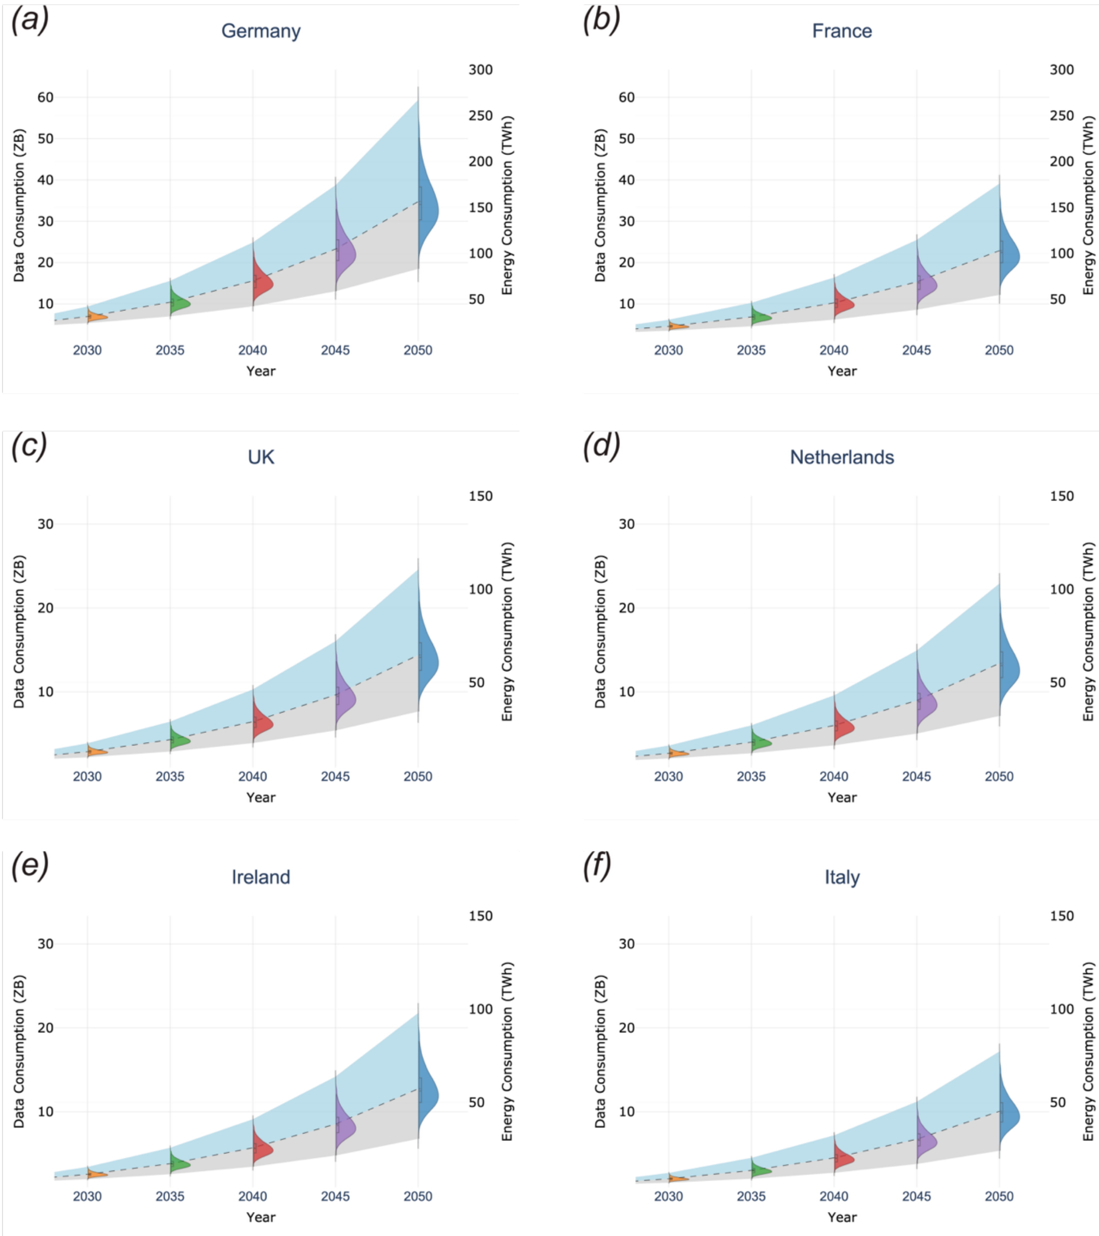


***Figure 2.*** *Country-level projections of data demand and data-centre electricity use. Panels a to f show Germany, France, the United Kingdom, the Netherlands, Ireland, and Italy. For each country, the coloured band marks the range between the maximum and the mean simulated data-demand paths and between the mean and the minimum paths. The dashed line indicates the mean path. Violin plots on the right axis show the distribution of implied electricity use at five-year intervals using 4.5 terawatt-hours per zettabyte. Country shares used for downscaling are taken from the Joint Research Centre report on European data centres for 2022 by Kamiya and Bertoldi.*

# 3. Spatial clustering of data-centre activity and recoverable heat: A Switzerland case study

We illustrate the urban clustering of data-centre activity using Switzerland as a case study (Fig. 3). We compile point locations and facility counts for 24 Swiss urban areas and distribute the 2050 national data-centre capacity across these points in proportion to each city's share of sites. The national total is set to 2.74 GW of installed IT power, obtained from an assumed 9.6 TWh $yr^{-1}$ of 2050 electricity use and a 40 percent average utilization of IT capacity. For each city $i$, we compute the share $s_i$ from the observed number of facilities, allocate $s_i \cdot 2.74\ \mathrm{GW}$ of IT capacity, and estimate recoverable heat as

$$Q_i = s_i \cdot 2.74\ \mathrm{GW} \times utilization \times \eta_{\mathrm{rec}}, \tag{2}$$

with *utilization* equal to 0.40 and heat-recovery efficiency equal to 0.90. We report annual heat in GWh $yr^{-1}$ using 8,760 hours and plot it as a colour scale. Marker size represents the local number of data centres.

The resulting map shows a clear concentration in the Zurich metropolitan area, which accounts for 38 of 114 facilities in the dataset and about one-third of national IT capacity. Under the assumptions above, Zurich's recoverable heat potential reaches about 2,880 GWh $yr^{-1}$ (≈ 330 $MW_{thermal}$ on average). Other notable hubs include Lucerne (14 sites; about 1,060 GWh $yr^{-1}$), Geneva (9 sites; about 680 GWh $yr^{-1}$), Bern (6 sites; about 455 GWh $yr^{-1}$), Basel (5 sites; about 380 GWh $yr^{-1}$), Lausanne (5 sites; about 380 GWh $yr^{-1}$), Lugano (7 sites; about 530 GWh $yr^{-1}$) and Zug (7 sites; about 530 GWh $yr^{-1}$). Smaller clusters appear along the Arc Lémanique and in Ticino. These concentrations align with large urban loads and existing district-heating footprints, which suggests high technical potential for low-temperature heat integration where return temperatures and network capacities are compatible.

The 2050 Swiss electricity demand for data centres from our Monte Carlo envelope was used to set the national scale and was converted to installed capacity with a utilization factor. An assumption here is also that, we consider constant siting proportions over time, so current facility counts proxy the 2050 distribution. Thus, the results of this extrapolation yields merely an indicative urban heat atlas rather than a siting forecast. The map therefore identifies priority districts for thermal coupling and grid interconnection studies, not discrete project locations.

However, such urban clustering of facilities means that recoverable heat from data centers is co-located with dense district heating demand. Under the Swiss case study, the resulting heat potentials are concentrated in existing heating network areas such as Zurich, Geneva, and the Arc Lémanique, which makes integration technically and logistically plausible where return temperatures and network capacities allow such integration. We therefore treat data center heat as a valid supply option for urban district heating and use this Swiss evidence base to parameterize heat recovery within our 2050 optimization model across Europe, assuming similar urban density–network correlations and climate-adjusted temperature profiles derived from Eurostat HDD datasets[50]. In the scenarios that follow, data centers are hence modeled as a sector-coupling interface between electricity and heat, with recoverable heat streams linked to district heating.

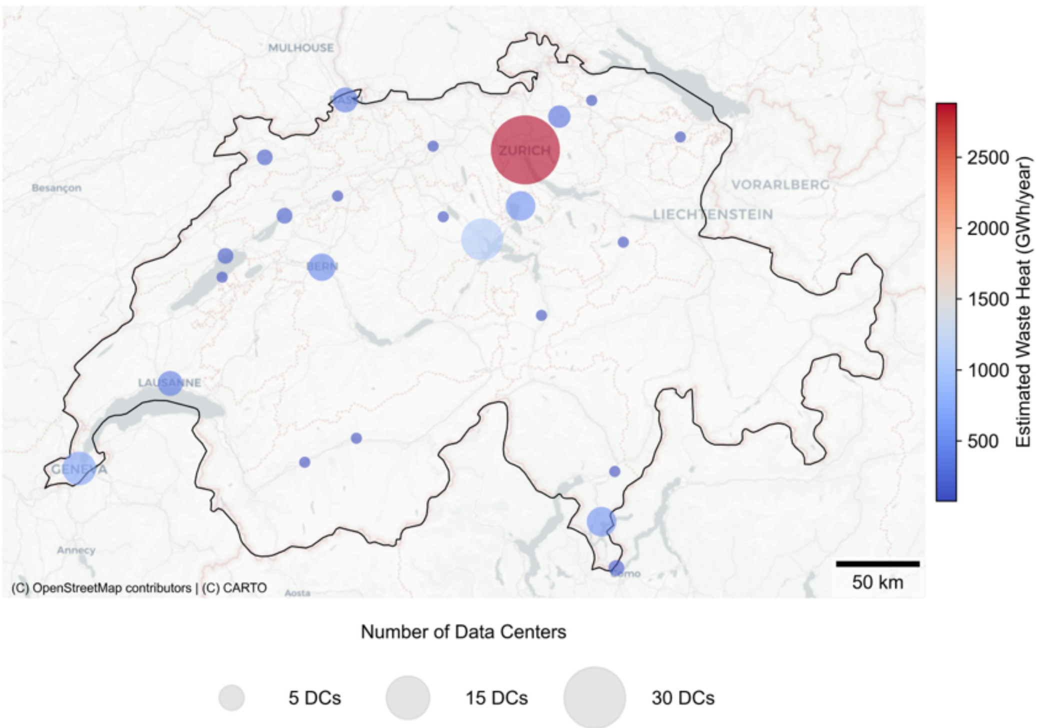


***Figure 3.*** *Swiss data-centre clustering and estimated recoverable heat in 2050. Bivariate map of 24 urban areas. Circle size shows the local number of data centres while colour shows estimated*

*waste heat (GWh $yr^{-1}$) computed from a national total of 2.74 GW of IT capacity, 40 percent utilization, and 90 percent heat-recovery efficiency. Basemap: CartoDB Positron; projection: EPSG:3857; scale bar: 50 km.*

# 4. Modeling-to-generate-alternatives (MGA) with exogenous ICT demand

Within the scope of the multi-energy technology-rich optimisation model used in this study (Energyscope), which is extended to the case of the EU27 countries, including Switzerland and UK, we include the ICT demand projected in the previous steps, exogenously as an additional end-use demand represented through annual totals distributed across representative time slices in 2050. This demand as an additional end-use must be met in every country alongside existing electricity, heat, transport, and industrial services. The model also proportionately delineates this demand by the use of simple factors, across enterprise-related, cloud computing-related and AI-related ICT demands, which will then be met by data centers of different types which incur different cost, and have different heat recovery potentials within the system. The model co-optimizes generation, conversion, storage, and transmission across EU-27 plus the United Kingdom and Switzerland. To explore not only the single cost optimum but a neighborhood of credible designs, we use modeling-to-generate-alternatives (MGA), with a 20 % optimality slack to explore near-cost-optimal system configurations.. MGA is implemented by solving a sequence of optimizations with a small optimality slack relative to the best solution and rotating objectives that emphasize different policy aims. This produces a set of 100 near-optimal scenarios, from which we highlight four reference solutions: minimum system costs, minimum environmental impacts, minimum regional inequality, and maximum renewable electricity. In the figure, these four solutions are shown by a star and colored markers, while the shaded boxes summarize the distribution across all MGA runs.

Across the 100 solutions, wind shows the widest range and largely sets the spread of total capacity. Wind installations vary upto 950 GW, with many solutions around 150 - 400 GW. PV ranges from about 50 to 350 GW. Hydropower additions are limited: run-of-river and dams remain within about 30 to 90 GW and about 60 to 150 GW, respectively. Nuclear lies within about 0 to 250 GW. CCGT with CCS appears as firm capacity between about 120 and 500 GW. Generation mirrors these patterns: wind output spans about 200 to 1,750 TWh per year, PV contributes about 50 to 300 TWh per year, while hydropower and nuclear vary in narrower bands set by resources and policy. The maximum renewables solution sits near the upper edge for wind and PV and approaches the top of the generation ranges. The minimum cost solution is closer to the medians for variable renewables and relies more on firm capacity. The minimum regional inequality solution avoids extreme concentrations. The minimum environmental impacts solution leans toward higher clean capacity within the feasible corridor.

Transmission expansion is required in almost all solutions and clusters in a tight interval of about 29 to 33 GW of additional transfer capacity. Storage energy separates scenarios more clearly. Pumped hydro concentrates near about 33 to 35 TWh, while power-to-X-to-power reaches about 14 to 17 TWh in high renewable cases. Total system costs fall between about 260 and 300 billion euros per year. Generation costs dominate and rise in portfolios with very high variable renewables and storage. Transmission and storage are smaller components in absolute terms but increase toward the high-renewables corner. The minimum environmental impacts solution accepts higher generation and storage spending to cut emissions. Steering toward regional equality raises costs modestly by avoiding extreme spatial concentration.

Adding explicit ICT demand shifts the feasible set toward portfolios with greater variable renewable capacity and increased storage energy capacity of roughly 10–20 TWh, depending on renewable mix, while keeping transmission needs within a narrow band. MGA reveals that several distinct portfolios can meet the same 2050 obligations at near-equal cost. Choices among them affect where capacity is built, how much firm capacity remains, and how far systems move toward high wind and PV. This ensemble view, consistent with earlier European MGA studies, highlights robust planning levers: balanced siting and transmission, storage scaled to renewable shares, and explicit inclusion of ICT demand in long-term strategies

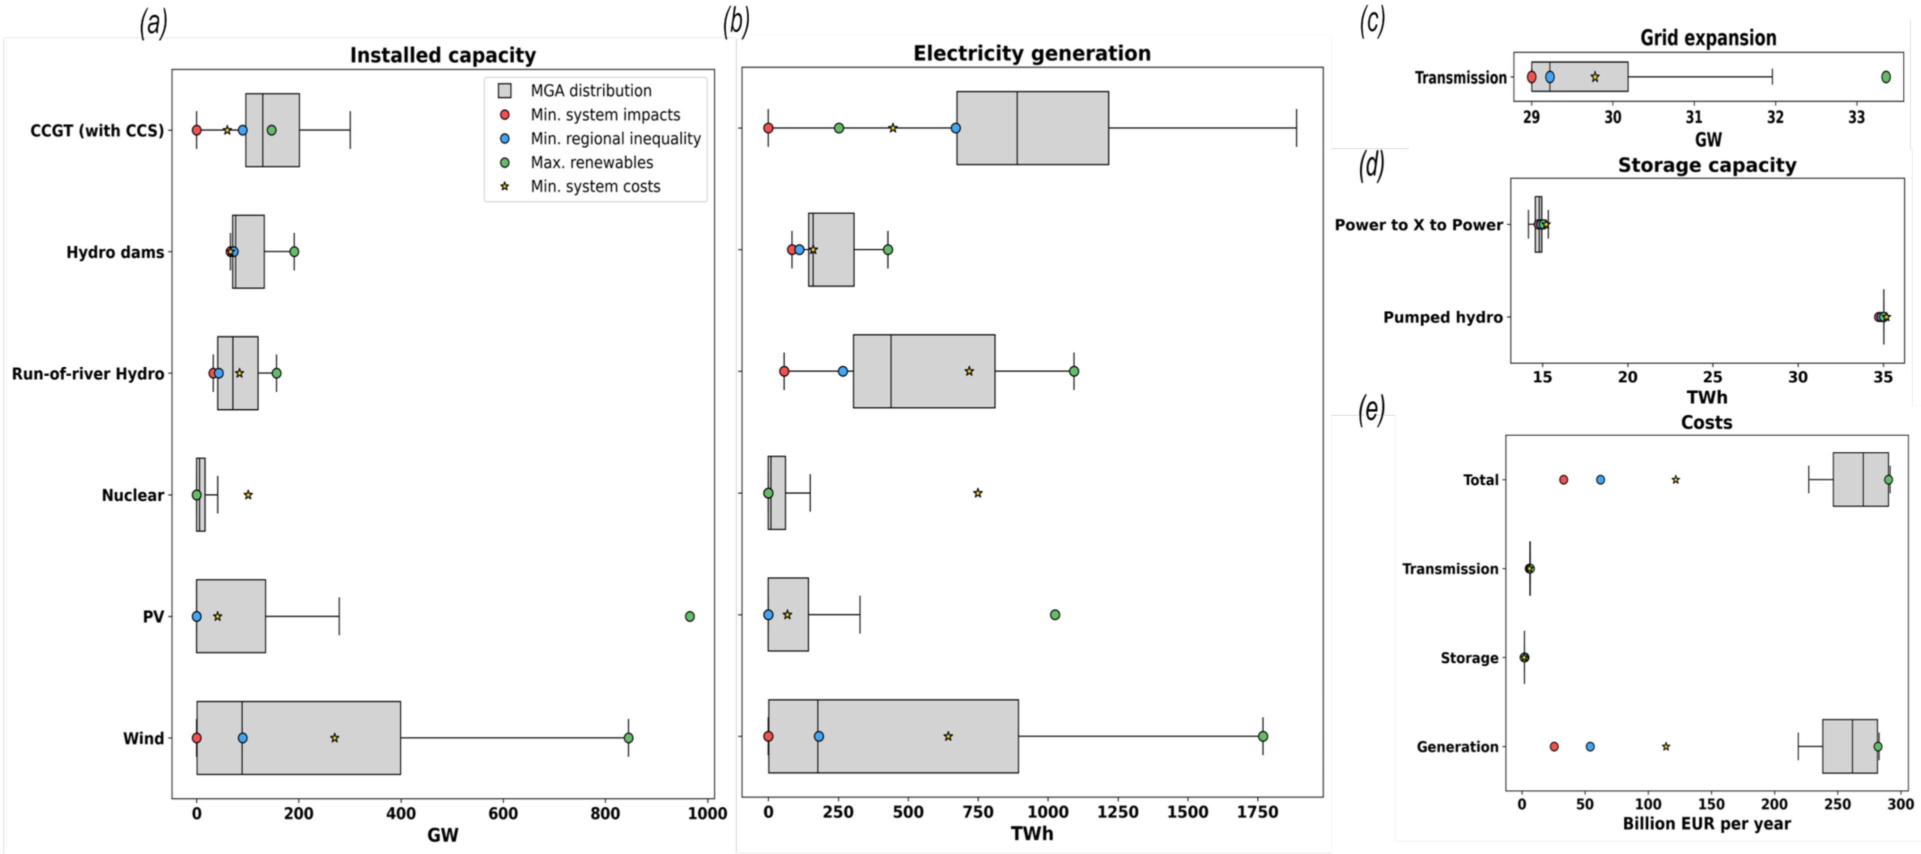


***Figure 4.*** *MGA exploration of 2050 system designs with explicit ICT demand. a, Installed capacity by technology in gigawatts. b, Annual electricity generation in terawatt-hours. c, Additional transmission capacity in gigawatts. d, Storage energy by technology in terawatt-hours. e, Annualized costs for generation, transmission, storage, and total in billion euros per year. Grey box plots summarize the distribution across 100 MGA solutions. Colored markers identify four reference solutions: minimum system costs (star), minimum environmental impacts (red), minimum regional inequality (blue), and maximum renewables (green). The ICT envelope from the Monte Carlo analysis is added exogenously to all runs; a small optimality slack enables exploration of near-optimal designs.*

# 5. Regional distribution of costs, renewables, and emissions across MGA scenarios

After the analysis of MGA results distribution, country-level results across the four objectives (minimum system costs, minimum environmental impacts, minimum regional inequality, and maximum renewable electricity) were analysed to have greater spatial granularity (Fig. 5a–l). In the cost-minimum case, expenditures concentrate in the largest load centres. Germany and France occupy the highest-cost categories on the scale and together account for a substantial share of continental spending, while smaller systems such as Ireland, Portugal, and the Baltic states lie at the lower end of the scale. Renewable output in this case remains close to demand rather than resource extremes: most countries cluster between 100 and 200 TWh $yr^{-1}$, and only a few exceed 200 TWh $yr^{-1}$. Residual emissions are highest in Germany and Poland, which fall within the upper emission quantile, consistent with their large thermal fleets and industrial demand, reflecting the persistence of dispatchable fossil capacity despite strong renewable deployment.

Shifting the objective to minimum environmental impact reduces greenhouse-gas totals across almost all countries. The largest absolute declines occur in Germany, Poland, Italy, and France, where the maps shift down by one to two colour categories relative to the cost minimum, implying emission reductions on the order of 20–40%. The spatial pattern of cost remains broadly similar, but investment moves from firm thermal capacity to renewables and storage, with more countries reaching the 200–250 TWh $yr^{-1}$ range of renewable output.

Prioritising minimum regional inequality spreads both costs and renewable output more evenly. Expenditures decrease in Germany and France and rise in Spain, Sweden, Romania, and Greece, narrowing the implied ratio between the highest- and lowest-cost countries from roughly three-to-one under the cost minimum to about two-to-one. Regional inequality was measured via normalized dispersion in total system costs across countries; future versions could quantify this with explicit metrics such as the Theil or Gini index. Emissions likewise flatten, with Central and Eastern European countries moving out of the **highest-emission categories**. However, this configuration increases total system costs relative to the cost-minimization case.

Maximising renewable electricity pushes several systems to the top of the renewable scale. Spain, Sweden, and Finland reach 300–400 TWh $yr^{-1}$, and large demand centres such as Germany, Italy, and France rise to 200–300 TWh $yr^{-1}$, roughly a two- to three-fold increase compared with their levels in the cost-minimum case. These expansions necessitate proportionally higher storage capacities and lead to increased curtailment without additional flexibility sources. The associated cost maps darken by one to two shades across most countries, reflecting additional storage and grid requirements, while emissions fall most strongly in countries with large residual thermal fleets, notably Germany and Poland.

Taken together, the maps indicate that objective choice changes where Europe builds and operates assets even when all scenarios satisfy the same 2050 demands, including the same explicit ICT loads across these scenarios. Minimising costs concentrates expenditures and emissions in a few hubs; minimising impacts or maximising renewables reduces emissions while increasing spending and shifting investment toward resource-rich regions; minimising regional inequality spreads both costs and benefits more evenly at the expense of higher totals. These distributional differences show two- to three-fold swings in national renewable output, category-level shifts in emissions for Poland and Germany, and a marked narrowing of cross-country cost dispersion that are central for planning and for designing mechanisms to share costs fairly as ICT demands grow.

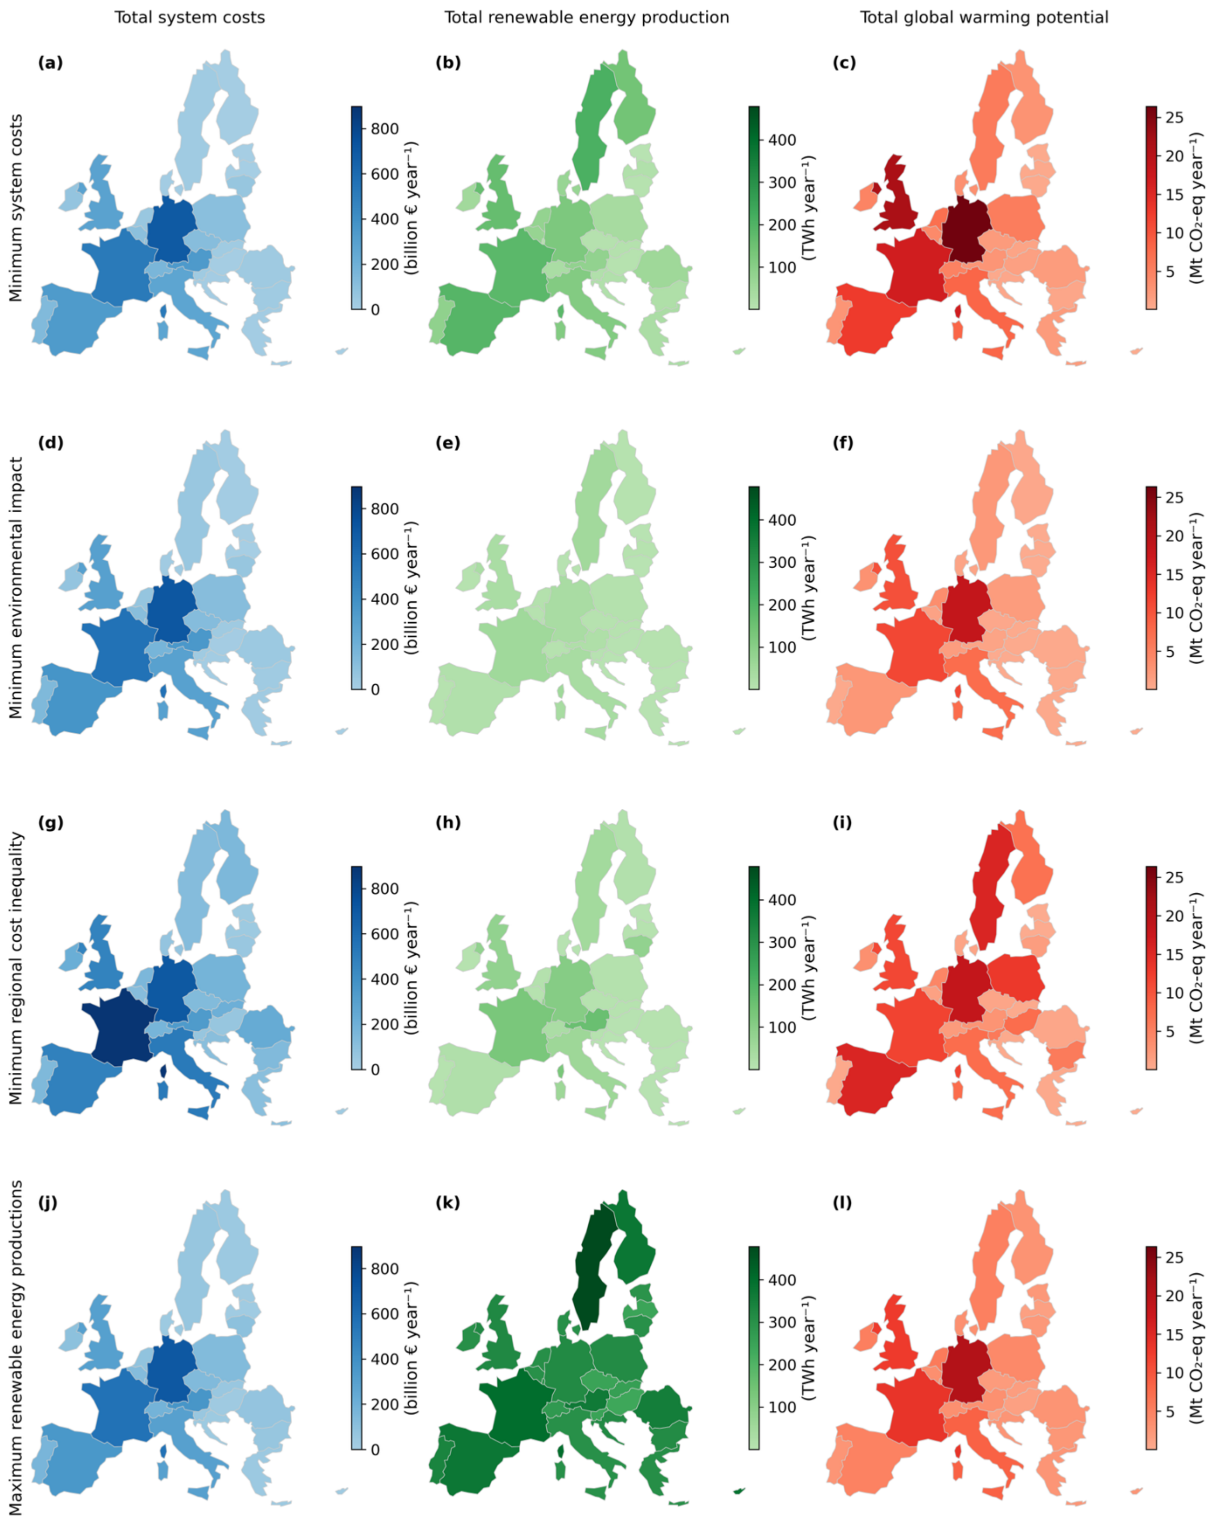

Total system costs
Total renewable energy production
Total global warming potential
Minimum system costs
Minimum environmental impact
Minimum regional cost inequality
Maximum renewable energy productions
(a)
(b)
(c)
(d)
(e)
(f)
(g)
(h)
(i)
(j)
(k)
(l)
(billion € year⁻¹)
0
200
400
600
800
(TWh year⁻¹)
100
200
300
400
(Mt CO₂-eq year⁻¹)
5
10
15
20
25

***Figure 5.*** *Country-level outcomes in 2050 across four MGA objectives. Rows show, top to bottom: minimum system costs, minimum environmental impacts, minimum regional inequality, and maximum renewable electricity. Columns report total system costs (billion euros per year), total renewable electricity production (TWh per year), and total greenhouse-gas emissions (Mt $CO_2$-eq per year). Shading indicates country totals under each objective after adding the same exogenous ICT demand and enforcing whole-energy-system constraints.*

# 5. Workload flexibility and curtailment reduction

While the ICT assets operations were already prelinked to show the effect of sectoral coupling between the heat and electricity sectors before, the ICT demands, were maintained almost at a near-constant rate as they are today, posing a constant demand to the system. However, given the fact that a growing share of ICT demand consists of deferrable workloads., we test whether data centres can act as a flexibility resource by shifting a share of computing to hours with high variable renewable output. We construct a workload flexibility factor (WFF) between 0 and 1. A value of 0 represents today's profile in which demand is largely time-inelastic. A value of 1 represents the upper bound in which the flexible share of workloads is fully aligned with the hourly VRE profile. Intermediate values interpolate linearly between these extremes, maintaining constant annual energy demand. The flexible share captures non-time-critical tasks such as batch analytics, training phases of AI models, data backup, and other deferred computing. Reported ranges in the literature suggest that about 30 to 70% of workloads can be deferred in this way[30], [31], [32]. To enable higher output during VRE peaks, the installed IT capacity must increase. We therefore co-vary the data centre capacity with WFF and quantify how curtailment changes at the European scale.

Figure 6 Panel (a) shows that the VRE curtailment rate declines monotonically as WFF increases. Curtailment falls from just over 9 percent at WFF = 0 to about 6.5 percent at WFF = 1, a reduction of around three percentage points, which corresponds to a relative decrease of roughly one third, at the European level. Panel (b) reports the associated installed IT capacity needed to realise these profiles. Capacity rises steadily from tens of gigawatts at WFF = 0 to several hundred gigawatts at WFF = 1. This increase reflects the need to process more work in a shorter window around renewable peaks while keeping service levels constant for time-critical tasks.

Panel (c) illustrates this mechanism at monthly scale. The light shaded area gives the baseline curtailment under inelastic demand. The darker area shows curtailment with flexible workloads. The dashed lines show data-centre dispatch with (dark-blue) and without (light-blue) flexibility. Without flexibility, dispatch is fairly flat through the year. With flexibility, dispatch peaks in summer months when solar generation is high and drops in the other seasons. The reduction in the shaded curtailment area during summer indicates that shifting deferrable computing into high-VRE months absorbs a sizeable share of otherwise wasted energy. The qualitative result is robust to alternative capacity-scaling rules, although the exact magnitude of curtailment reduction depends on the flexible share and the degree of oversizing.

Hence, dynamic workload scheduling emerges as a credible system-level flexibility lever, enabling deeper renewable integration without proportionally expanding storage or accepting higher curtailment. These results also show that capacity planning for large facilities should jointly consider IT sizing, interconnection with grid, co-ordination with other system actors and heat-recovery opportunities, since the same alignment that reduces curtailment also concentrates waste heat in periods with high renewable output. In our system scenarios, we therefore show that data centres not only as an electricity end-use but also as a flexibility asset whose dispatch can be steered toward VRE-rich hours within the bounds of a realistic deferrable share.

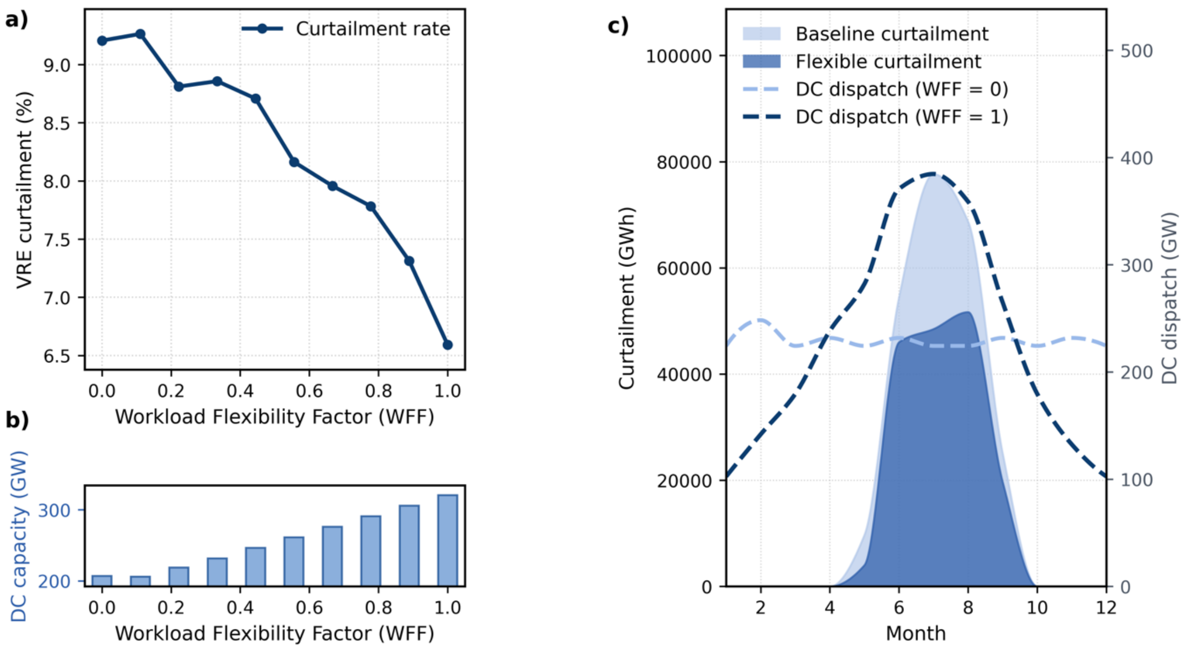


***Figure 6.*** *Data-centre workload flexibility reduces VRE curtailment. a, Curtailment rate as a function of the workload flexibility factor (WFF). Values decrease from just above 9 percent at WFF = 0 to about 6.5 percent at WFF = 1. b, Installed IT capacity required to realise the flexible profiles rises with WFF from tens to several hundred gigawatts. c, Monthly curtailment for the baseline (light) and flexible case (dark), with data-centre dispatch shown as dashed lines. Flexible operation shifts computing into VRE-rich months, lowering curtailment while requiring higher peak IT capacity.*

# 6. Policy to align data-centre flexibility with system needs

Across Europe, system operators are procuring growing volumes of flexibility services, while many large data-centre operators are simultaneously developing private microgrids optimised for reliability and low power-usage-effectiveness (close to 1). This leads to a coordination gap where operators minimise grid-interactive capability while transmission and distribution system operators continue to curtail variable renewables and pay for balancing capacity that data centres could provide. Evidence from market design indicates that dispatch arrangements and remuneration determine utilisation[51]. When the system operator can call flexible capacity with availability payments and verified baselines, curtailment falls more than under purely merchant strategies that follow price spreads and under-deliver during stressed hours.

Data centres are technically suited to these services. They combine uninterruptible power supplies with fast response, controllable cooling loads, on-site generation, uniteruppted power-supply (UPS) and a sizable fraction of deferrable computing with reponse times with UPS as fast as tens of milliseconds[52]. Studies of production traces and operator surveys place the deferrable share between 30 and 70 percent without breaching service levels[30], [31], [32]. In our scenarios, shifting these workloads toward renewable-rich hours and exporting low-temperature heat reduces curtailment and fossil peaking and lowers total system costs by reducing curtailment. Private optimisation of data center operations based on network traffic and bring-your-own-power (BYOP) strategies adopted by data center operators tend to under-size interactive capacity to preserve reported PUE, which forgoes these system benefits.

A policy framework is therefore required that makes participation routine and remunerated. Planning consent for new large facilities should include baseline and telemetry requirements and controllable UPS and HVAC capable of frequency response and load shifting providing operational flexibility with protected state-of-charge limits and fail-safe rules. Markets should offer standard products that match data-centre capabilities: sub-second frequency response from UPS, multi-hour intraday shifting from IT and cooling, and day-ahead scheduling of deferrable workloads into renewable peaks. Payments should combine availability with pay-as-performance energy settlement to ensure dispatch when the system needs it. Interconnection agreements for sites with on-site generation or storage should specify import and export flexibility windows, enable frequency response, and apply locational signals that reward operation during regional renewable peaks. Such mechanisms would align with the design principles of existing European balancing and capacity markets under the Clean Energy Package[53]. In urban areas, participation should be coupled to heat-recovery feasibility so that low-temperature heat and flexible demand are co-optimised. For very large sites, a modest and rising obligation to offer a share of peak load as flexible capacity, with automatic opt-out when service levels are at risk, would ensure system-wide energy and cost benefits while preserving ICT supply reliability. Together, these measures align private incentives with public goals and convert existing UPS, thermal capacity, and deferrable compute into a coordinated resource that reduces curtailment potential up to 9% of the overall VRE generation.

# 7. Methods and Assumptions

## 7.1. Methods:

The analysis builds upon the *EnergyScope* optimization framework[37], [54] a linear techno-economic model that determines cost-optimal configurations of national energy systems under under a set of representative hourly time slices capturing typical seasonal and diurnal variations. Each country $r \in R$ is modelled as an autonomous energy system with its own technologies, resources, and end-use demands $EUD_c(r,t)$ for each energy carrier $c$, technology *i,* and time slice $t$. The objective is to minimize the total annualized system cost

$$\min \boldsymbol{C_{tot}} = \sum_{r \in R} \boldsymbol{C_{inv}}(r) + \boldsymbol{C_{maint}}(r) + \boldsymbol{C_{op}}(r) \tag{3}$$

where the investment ($\boldsymbol{C_{inv}}$), maintenance ($\boldsymbol{C_{maint}}$), and operational ($\boldsymbol{C_{op}}$) costs are computed as:

$$\boldsymbol{C_{inv}}(r) = \sum_{i} c_{inv}(i) \cdot \boldsymbol{F}(r,i) \cdot \tau(i) \tag{4}$$

$$\boldsymbol{C_{maint}}(r) = \sum_{i} c_{maint}(i) \cdot \boldsymbol{F}(r,i) \cdot \frac{1}{n} \tag{5}$$

$$\boldsymbol{C_{op}}(r) = \sum_{j} \sum_{t} c_{op}(j,t) \cdot \boldsymbol{F_t}(r,j,t) \cdot t_{op}(t) \tag{6}$$

All techno-economic parameters ($c_{inv}(i)$, $c_{maint}(i), c_{op}(j)$), efficiencies $\eta(i)$, and availabilities follow the standard *EnergyScope* database.

Energy balances ensure that, for each carrier c, the total domestic production, storage discharge, and imports $\boldsymbol{F_t^+}(r,j,t)$ equal the final energy demand $EUD(r,t)$, storage charging, and exports $\boldsymbol{F_t^-}(r,j,t)$:

$$\begin{aligned}\boldsymbol{EUD}(r,t) = \sum_{i} \eta(i) \cdot \boldsymbol{F_t}(r,i,t) + \sum_{j} \boldsymbol{F_t^+}(r,j,t) - \boldsymbol{F_t^-}(r,j,t) \\ + \sum_{sto} \boldsymbol{F_t^+}(r,sto,t) - \boldsymbol{F_t^-}(r,sto,t)\end{aligned} \tag{7}$$

The *Digital Heatwaves* extension introduces three model enhancements within this existing structure. First, an ICT electricity end-use $\boldsymbol{EUD}(i,r,t)$ is added to the electricity balance exogenously, where *i* represents the type of ICT demand such as enterprise-level ICT demand, AI-type ICT demand and cloud-computing type ICT demand. The demand was disaggregated to these types to maintain granularity across the technological parameters used to meet these demands. Data center as a technology unit for each type, was modelled based using cost parameters based on a cost calculater from Schneider electric[55]. The data centers would in turn pose an electricity demand that will have to be met by the system. Second, a recoverable-heat stream is integrated into the district-heating balance that takes into account the integration of data centers in urban districts. The heat-recovery efficiency was set at 0.9 with recoverable heat temporally coinciding with ICT electricity demand, given their continuous operation profiles[56]. Third, a workload-flexibility factor $\phi(r)$ modifies the temporal distribution of ICT demand by ensuring that annual ICT workloads $W_{ICT}(r)$ is conserved while allowing time-shifting within predefined limits. Flexibility was operationalized by redistributing hourly ICT loads using the normalized workload profiles detailed in Table S7, ensuring that total annual demand remains invariant. To explore alternative near-optimal energy configurations, we apply a *Modelling-to-Generate-Alternatives* (MGA) approach. After identifying the cost-optimal solution $x^*$ that minimizes the total system cost $C(x)$, the optimization is repeated under a relaxed cost constraint (within a 20 % cost slack, ε = 0.2), allowing exploration of the feasible space of "good" solutions whose cost remains close to the optimum:

$$\begin{aligned} \text{Minimize/Maximize} \quad & C(x) \\ \text{s.t.} \quad & C(x) \leq (1+\varepsilon)\, C(x^*), \\ & x \in \Omega, \end{aligned} \tag{8}$$

where Ω denotes all techno-economic and energy balance constraints. The slack parameter is set to 20% in this study similar to the MGA analysis carried out in Trutnevyte et. al.[45], however it must be noted that the choice of slack is also region-specific as has been shown by Fossas et. al.[57] By successively minimizing or maximizing for the selected objectives the MGA framework reveals the diversity of system configurations that achieve comparable performance yet embody distinct technological or spatial strategies. This formulation allows assessing how digitalization, flexibility, and heat recovery affect the resilience and efficiency of European energy systems in 2050.

## 7.2 Input data & assumptions

The methodology is applied the EnergyScope-EU model. The harmonized dataset covering all EU Member States, the United Kingdom, and Switzerland, derived from the detailed European compilation of Meyer (2021)[58] and Multi-cell model developed at University of Liège. Each country is represented as an autonomous energy system characterized by its own end-use demands, renewable potentials, and technology stocks. Final energy demands $EUD_c(r,t)$ are disaggregated into households, services, industry, and transport, with sectoral shares and growth trajectories taken from the European Calculator (EUCalc) model[59] and Eurostat population and GDP projections[50].

Monthly profiles for heating and lighting follow climatic indicators (heating-degree days and daylight duration), ensuring consistent temporal allocation across regions (refer Table S5 and S6). Renewable-energy capacities and potentials are based on the JRC ENSPRESO[60], Hydropower[61], and GeoElec databases[62], complemented by Eurostat[50] sources. Solar, wind, and hydro capacity factors are derived from historical monthly generation statistics and assumed geography-dependent but time-invariant. Biomass and waste potentials scale proportionally with population.

*Table 2: Key input categories and data sources used in the EnergyScope-EU model*

| Input | Description | Main sources | Resolution |
|---|---|---|---|
| Final energy demand | Sectoral end-uses (HH, Services, Industry, Transport) | EUCalc (2018) [59] Eurostat (2015–2050) [50]18/05/2026 17:39:00 | Annual → representative time slices |
| Technology database | $c_{inv}, c_{maint}, \eta, c_p$, .... | EnergyScope-EU (Meyer 2020) [58], JRC-SETIS[63], IEA[64] 18/05/2026 17:39:00 | Country-specific |
| Renewable potentials | Solar, wind, hydro, geothermal, biomass | ENSPRESO [60], JRC Hydropower [61], GeoElec[62], EU Biomass Atlas[65] | Geographic / country level |

| | | | |
|---|---|---|---|
| Emission factors & fuel prices | Fossil and renewable energy carriers | PRIMES 2020[66], ENTSO-E[66]18/05/2026 17:39:00 | Annual, by carrier |
| Climatic data | Heating & cooling degree days, daylight hours | ERA5[67], Eurostat Climate Indicators[68] | Monthly |
| Cross-border topology | Country adjacencies and ferry links | ENTSO-E[66], Google Maps API[69] | Static (2015–2050) |

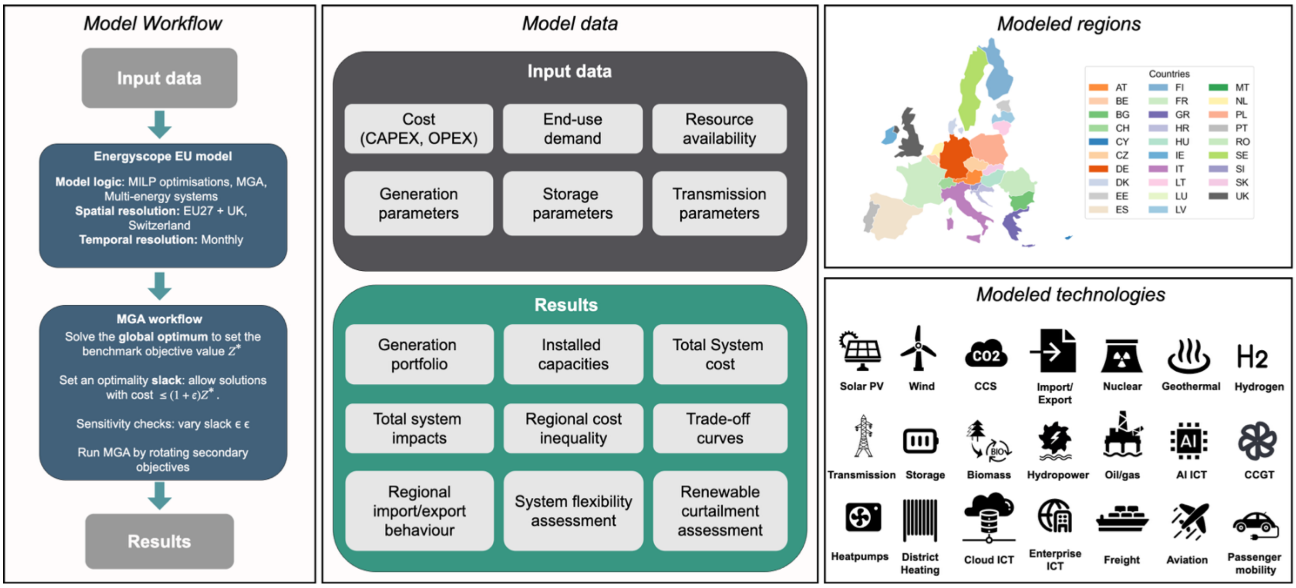


***Figure 7.*** *Overview of the modeling workflow, data, outputs, modeled regions (EU-27+UK+CH), and technology set in EnergyScope-EU with MGA.*

## References:


[1] "The European Green Deal - European Commission." Accessed: Oct. 20, 2025. [Online]. Available: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en

[2] "Reducing carbon emissions: EU targets and policies | News | European Parliament." Accessed: Oct. 31, 2023. [Online]. Available: https://www.europarl.europa.eu/news/en/headlines/society/20180305STO99003/reducing-carbon-emissions-eu-targets-and-policies

[3] *Going climate-neutral by 2050: a strategic long term vision for a prosperous, modern, competitive and climate neutral EU economy*. Publications Office of the European Union, 2019. Accessed: Oct. 20, 2025. [Online]. Available: https://data.europa.eu/doi/10.2834/02074

[4] D. Reiche and M. Bechberger, "Policy differences in the promotion of renewable energies in the EU member states," *Energy Policy*, vol. 32, no. 7, pp. 843–849, May 2004, doi: 10.1016/S0301-4215(02)00343-9.

[5] M. Nicolini and M. Tavoni, "Are renewable energy subsidies effective? Evidence from Europe," *Renewable and Sustainable Energy Reviews*, vol. 74, pp. 412–423, Jul. 2017, doi: 10.1016/j.rser.2016.12.032.

[6] K. Neuhoff, N. May, and J. C. Richstein, "Financing renewables in the age of falling technology costs," *Resource and Energy Economics*, vol. 70, p. 101330, Nov. 2022, doi: 10.1016/j.reseneeco.2022.101330.
[7] M. Xiao, T. Junne, J. Haas, and M. Klein, "Plummeting costs of renewables - Are energy scenarios lagging?," *Energy Strategy Reviews*, vol. 35, p. 100636, May 2021, doi: 10.1016/j.esr.2021.100636.
[8] P. Rafaj *et al.*, "Outlook for clean air in the context of sustainable development goals," *Global Environmental Change*, vol. 53, pp. 1–11, Nov. 2018, doi: 10.1016/j.gloenvcha.2018.08.008.
[9] A. C. Marques, J. A. Fuinhas, and D. A. Pereira, "Have fossil fuels been substituted by renewables? An empirical assessment for 10 European countries," *Energy Policy*, vol. 116, pp. 257–265, May 2018, doi: 10.1016/j.enpol.2018.02.021.
[10] X. Ouyang and B. Lin, "Impacts of increasing renewable energy subsidies and phasing out fossil fuel subsidies in China," *Renewable and Sustainable Energy Reviews*, vol. 37, pp. 933–942, Sep. 2014, doi: 10.1016/j.rser.2014.05.013.
[11] S. Ntanos *et al.*, "Renewable Energy and Economic Growth: Evidence from European Countries," *Sustainability*, vol. 10, no. 8, p. 2626, Aug. 2018, doi: 10.3390/su10082626.
[12] A. Omri, N. Ben Mabrouk, and A. Sassi-Tmar, "Modeling the causal linkages between nuclear energy, renewable energy and economic growth in developed and developing countries," *Renewable and Sustainable Energy Reviews*, vol. 42, pp. 1012–1022, Feb. 2015, doi: 10.1016/j.rser.2014.10.046.
[13] J. Wang, S. Zhang, and Q. Zhang, "The relationship of renewable energy consumption to financial development and economic growth in China," *Renewable Energy*, vol. 170, pp. 897–904, Jun. 2021, doi: 10.1016/j.renene.2021.02.038.
[14] A. Held *et al.*, "How can the renewables targets be reached cost-effectively? Policy options for the development of renewables and the transmission grid," *Energy Policy*, vol. 116, pp. 112–126, May 2018, doi: 10.1016/j.enpol.2018.01.025.
[15] P. J. Heptonstall and R. J. K. Gross, "A systematic review of the costs and impacts of integrating variable renewables into power grids," *Nat Energy*, vol. 6, no. 1, pp. 72–83, Jan. 2021, doi: 10.1038/s41560-020-00695-4.
[16] C. Lang, J. J. Opaluch, and G. Sfinarolakis, "The windy city: Property value impacts of wind turbines in an urban setting," *Energy Economics*, vol. 44, pp. 413–421, Jul. 2014, doi: 10.1016/j.eneco.2014.05.010.
[17] N. Leskinen, J. Vimpari, and S. Junnila, "The impact of renewable on-site energy production on property values," *Journal of European Real Estate Research*, vol. 13, no. 3, pp. 337–356, Apr. 2020, doi: 10.1108/JERER-11-2019-0041.
[18] Y. Song, D. Xue, Q. Wen, H. Ye, and B. Ma, "Renewables' impacts on ecosystems in China," *Science*, vol. 383, no. 6689, pp. 1302–1303, Mar. 2024, doi: 10.1126/science.ado6369.
[19] D. Serrano *et al.*, "Renewables in Spain threaten biodiversity," *Science*, vol. 370, no. 6522, pp. 1282–1283, Dec. 2020, doi: 10.1126/science.abf6509.
[20] B. Bruckner *et al.*, "Ecologically unequal exchanges driven by EU consumption," *Nat Sustain*, vol. 6, no. 5, pp. 587–598, May 2023, doi: 10.1038/s41893-022-01055-8.
[21] C. Ganzleben and A. Kazmierczak, "Leaving no one behind – understanding environmental inequality in Europe," *Environ Health*, vol. 19, no. 1, p. 57, May 2020, doi: 10.1186/s12940-020-00600-2.
[22] N. Jones, "How to stop data centres from gobbling up the world's electricity," *Nature*, vol. 561, no. 7722, pp. 163–166, Sep. 2018, doi: 10.1038/d41586-018-06610-y.

[23] P. Huang *et al.*, “A review of data centers as prosumers in district energy systems: Renewable energy integration and waste heat reuse for district heating,” *Applied Energy*, vol. 258, p. 114109, Jan. 2020, doi: 10.1016/j.apenergy.2019.114109.

[24] J. Whitney and P. Delforge, “Data center efficiency assessment: Scaling up energy efficiency across the data center industry,” Natural Resources Defense Council (NRDC) and Anthesis Group, Aug. 2014. [Online]. Available: https://info.anthesisgroup.com/hubfs/Data-Center-Issue-Paper-final826.pdf

[25] F. Gracceva and P. Zeniewski, “A systemic approach to assessing energy security in a low-carbon EU energy system,” *Applied Energy*, vol. 123, pp. 335–348, Jun. 2014, doi: 10.1016/j.apenergy.2013.12.018.

[26] J. McIntyre and M. Pradhan, “A Systemic Approach to Addressing the Complexity of Energy Problems,” *Systemic Practice and Action Research*, vol. 16, no. 3, pp. 213–223, Jun. 2003, doi: 10.1023/A:1023811922579.

[27] A. E. H. Berjawi, S. L. Walker, C. Patsios, and S. H. R. Hosseini, “An evaluation framework for future integrated energy systems: A whole energy systems approach,” *Renewable and Sustainable Energy Reviews*, vol. 145, p. 111163, Jul. 2021, doi: 10.1016/j.rser.2021.111163.

[28] T. Chen, Y. Zhang, X. Wang, and G. B. Giannakis, “Robust Workload and Energy Management for Sustainable Data Centers,” *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 3, pp. 651–664, Mar. 2016, doi: 10.1109/JSAC.2016.2525618.

[29] Y. Guo, Y. Gong, Y. Fang, P. P. Khargonekar, and X. Geng, “Energy and Network Aware Workload Management for Sustainable Data Centers with Thermal Storage,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 8, pp. 2030–2042, Aug. 2014, doi: 10.1109/TPDS.2013.278.

[30] Y. Cao *et al.*, “Data-driven flexibility assessment for internet data center towards periodic batch workloads,” *Applied Energy*, vol. 324, p. 119665, 2022.

[31] M. T. Takci, M. Qadrdan, J. Summers, and J. Gustafsson, “Data centres as a source of flexibility for power systems,” *Energy Reports*, vol. 13, pp. 3661–3671, 2025.

[32] G. Ghatikar, “Demand response opportunities and enabling technologies for data centers: Findings from field studies,” 2012.

[33] D. Paul, W.-D. Zhong, and S. K. Bose, “Energy efficient scheduling in data centers,” in *2015 IEEE International Conference on Communications (ICC)*, Jun. 2015, pp. 5948–5953. doi: 10.1109/ICC.2015.7249270.

[34] B. Chen, R. Xiong, H. Li, Q. Sun, and J. Yang, “Pathways for sustainable energy transition,” *Journal of Cleaner Production*, vol. 228, pp. 1564–1571, Aug. 2019, doi: 10.1016/j.jclepro.2019.04.372.

[35] S. Bolwig *et al.*, “Review of modelling energy transitions pathways with application to energy system flexibility,” *Renewable and Sustainable Energy Reviews*, vol. 101, pp. 440–452, Mar. 2019, doi: 10.1016/j.rser.2018.11.019.

[36] “Understanding the vicious cycle of myopic foresight and constrained technology deployment in transforming the European energy system: iScience.” Accessed: Oct. 20, 2025. [Online]. Available: https://www.cell.com/iscience/fulltext/S2589-0042(24)02594-X

[37] S. Moret, M. Bierlaire, and F. Maréchal, “Strategic Energy Planning under Uncertainty: a Mixed-Integer Linear Programming Modeling Framework for Large-Scale Energy Systems,” in *Computer Aided Chemical Engineering*, vol. 38, Z. Kravanja and M. Bogataj, Eds., in 26 European Symposium on Computer Aided Process Engineering, vol. 38. , Elsevier, 2016, pp. 1899–1904. doi: 10.1016/B978-0-444-63428-3.50321-0.

[38] K. Hansen, B. V. Mathiesen, and I. R. Skov, "Full energy system transition towards 100% renewable energy in Germany in 2050," *Renewable and Sustainable Energy Reviews*, vol. 102, pp. 1–13, Mar. 2019, doi: 10.1016/j.rser.2018.11.038.

[39] I. Alaperä, S. Honkapuro, and J. Paananen, "Data centers as a source of dynamic flexibility in smart girds," *Applied Energy*, vol. 229, pp. 69–79, Nov. 2018, doi: 10.1016/j.apenergy.2018.07.056.

[40] S. Chen *et al.*, "Operational flexibility of active distribution networks with the potential from data centers," *Applied Energy*, vol. 293, p. 116935, Jul. 2021, doi: 10.1016/j.apenergy.2021.116935.

[41] Y. Zhan, M. Ghamkhari, D. Xu, S. Ren, and H. Mohsenian-Rad, "Extending Demand Response to Tenants in Cloud Data Centers via Non-Intrusive Workload Flexibility Pricing," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3235–3246, Jul. 2018, doi: 10.1109/TSG.2016.2628886.

[42] J. Price and I. Keppo, "Modelling to generate alternatives: A technique to explore uncertainty in energy-environment-economy models," *Applied Energy*, vol. 195, pp. 356–369, Jun. 2017, doi: 10.1016/j.apenergy.2017.03.065.

[43] K. Esser, J. Finke, V. Bertsch, and A. Löschel, "Participatory modelling to generate alternatives to support decision-makers with near-optimal decarbonisation options," *Applied Energy*, vol. 395, p. 126184, Oct. 2025, doi: 10.1016/j.apenergy.2025.126184.

[44] G. A. Brady, N. Kapur, J. L. Summers, and H. M. Thompson, "A case study and critical assessment in calculating power usage effectiveness for a data centre," *Energy Conversion and Management*, vol. 76, pp. 155–161, Dec. 2013, doi: 10.1016/j.enconman.2013.07.035.

[45] J.-P. Sasse and E. Trutnevyte, "Regional impacts of electricity system transition in Central Europe until 2035," *Nat Commun*, vol. 11, no. 1, p. 4972, Oct. 2020, doi: 10.1038/s41467-020-18812-y.

[46] M. Koot and F. Wijnhoven, "Usage impact on data center electricity needs: A system dynamic forecasting model," *Applied Energy*, vol. 291, p. 116798, Jun. 2021, doi: 10.1016/j.apenergy.2021.116798.

[47] "The ICT sector needs to drive energy innovation to handle growing data volumes sustainably. Here's how.," World Economic Forum. Accessed: Oct. 20, 2025. [Online]. Available: https://www.weforum.org/stories/2024/05/data-growth-drives-ict-energy-innovation/

[48] "Data growth worldwide 2010-2028," Statista. Accessed: Oct. 20, 2025. [Online]. Available: https://www.statista.com/statistics/871513/worldwide-data-created/

[49] G. Kamiya and P. Bertoldi, "Energy Consumption in Data Centres and Broadband Communication Networks in the EU," JRC Publications Repository. Accessed: Oct. 20, 2025. [Online]. Available: https://publications.jrc.ec.europa.eu/repository/handle/JRC135926

[50] "Database - Eurostat." Accessed: Oct. 29, 2025. [Online]. Available: https://ec.europa.eu/eurostat/data/database

[51] M. Lynch, G. Longoria, and J. Curtis, "Market design options for electricity markets with high variable renewable generation," *Utilities Policy*, vol. 73, p. 101312, Dec. 2021, doi: 10.1016/j.jup.2021.101312.

[52] D. Fewson, "7 - Switched-mode power supplies," in *Introduction to Power Electronics*, D. Fewson, Ed., Oxford: Butterworth-Heinemann, 1998, pp. 125–148. doi: 10.1016/B978-034069143-4/50009-5.

[53] "CEP." Accessed: Oct. 30, 2025. [Online]. Available: https://www.entsoe.eu/cep/

[54] "Power to the People: On the Role of Districts in Decentralized Energy Systems." Accessed: Oct. 29, 2025. [Online]. Available: https://www.mdpi.com/1996-1073/17/7/1718
[55] "Data Center Capital Cost Calculator." Accessed: Oct. 29, 2025. [Online]. Available: https://www.se.com/ww/en/work/solutions/system/s1/data-center-and-network-systems/trade-off-tools/data-center-capital-cost-calculator/
[56] "High-Performance Computing Data Center Cooling System Energy Efficiency | Computational Science | NREL." Accessed: Oct. 30, 2025. [Online]. Available: https://www.nrel.gov/computational-science/data-center-cooling-system
[57] A. Fossas Tenas, X. Wen, and E. Trutnevyte, "Empirical evidence for slack in Modeling to Generate Alternatives," Oct. 02, 2025, *Social Science Research Network, Rochester, NY*: 5608431. doi: 10.2139/ssrn.5608431.
[58] S. M. P. Meyer, "EnergyScope Europe - Data research and model adaptation," Jan. 2021, Accessed: Oct. 29, 2025. [Online]. Available: https://infoscience.epfl.ch/handle/20.500.14299/174879
[59] "eucalc." Accessed: Oct. 29, 2025. [Online]. Available: http://tool.european-calculator.eu/app/emissions/ghg-emissions/?levers=1j12112ffl11211mp2b111fffffpppppp11f411111e3211r211l21n221
[60] "Joint Research Centre Data Catalogue - ENSPRESO - an open data, EU-28 wide, transparent a... - European Commission." Accessed: Oct. 29, 2025. [Online]. Available: https://data.jrc.ec.europa.eu/collection/id-00138
[61] E. Quaranta *et al.*, "Clean Energy Technology Observatory: Hydropower and Pumped Hydropower Storage in the European Union – 2022 Status Report on Technology Development, Trends, Value Chains and Markets," JRC Publications Repository. Accessed: Oct. 29, 2025. [Online]. Available: https://publications.jrc.ec.europa.eu/repository/handle/JRC130587
[62] "Geoelec." Accessed: Oct. 29, 2025. [Online]. Available: http://www.geoelec.eu/
[63] "SETIS research and innovation data - SETIS - SET Plan information system." Accessed: Oct. 29, 2025. [Online]. Available: https://setis.ec.europa.eu/publications-and-documents/setis-research-and-innovation-data_en
[64] "IEA – International Energy Agency," IEA. Accessed: Oct. 29, 2025. [Online]. Available: https://www.iea.org/data-and-statistics
[65] "DBFZ DataLab - OpenData." Accessed: Oct. 29, 2025. [Online]. Available: https://datalab.dbfz.de/
[66] "ENTSO-E Transparency Platform." Accessed: Oct. 29, 2025. [Online]. Available: https://transparency.entsoe.eu/
[67] "ERA5 hourly data on single levels from 1940 to present." Accessed: Oct. 29, 2025. [Online]. Available: https://cds.climate.copernicus.eu/datasets/reanalysis-era5-single-levels?tab=overview
[68] "Climate change indicators for the EU." Accessed: Oct. 29, 2025. [Online]. Available: https://ec.europa.eu/eurostat/cache/infographs/climate-change/
[69] "Google Maps Platform," Google for Developers. Accessed: Oct. 29, 2025. [Online]. Available: https://developers.google.com/maps