Libyan Journal of Scientific Research
المجلة الليبية للبحث العلمي
E-ISSN: 0000-0000
Volume 1, Issue 1, 2026
Page No: 1-13
Website: https://ljsr.ly/index.php

Flexible Data Centres as Grid-Interactive Energy Resources: A Multi-Objective Optimisation Framework for Demand Response, Renewable Integration, and Energy Storage

Mohamed Abubaker[1*], Ali Hesri[2]
[1]Libyan Center for Sustainable Development Research, Al-Khums, Libya
[2]Electrical and Electronic Engineering Department, Faculty of Engineering, Fezzan University, Murzuq, Libya

مراكز البيانات المرنة بوصفها موارد طاقة تفاعلية مع الشبكة: إطار تحسين متعدد الأهداف للاستجابة للطلب، ودمج الطاقة المتجددة، وتخزين الطاقة

محمد أبوبكر[1*]، علي هسري
[1]المركز الليبي لأبحاث التنمية المستدامة، الخمس، ليبيا
[2]قسم الهندسة الكهربائية والإلكترونية، كلية الهندسة، جامعة فزان، مرزق، ليبيا

*Corresponding author: mabubaker@lsd.ly

Received: January 02, 2026 Revised: February 25, 2026 Accepted: February 30, 2026
Published: March 14, 2026

Abstract:
The rapid expansion of cloud computing and artificial intelligence is increasing electricity demand from data centres while intensifying local grid congestion and connection delays. Nevertheless, data centres also possess controllable resources, including delay-tolerant computing workloads, battery energy storage, onsite renewable generation and flexible cooling, that can enable them to operate as grid-interactive energy resources. This paper develops a multi-objective linear optimisation framework for coordinating these resources. The formulation simultaneously minimises electricity expenditure, operational carbon emissions and peak grid demand while preserving computational service requirements, battery operating limits and power balance. A transparent 24-hour illustrative case study is used to demonstrate the framework. Four operating strategies are compared: conventional operation, workload shifting, battery-only flexibility and coordinated optimisation. Under the stated synthetic assumptions, the coordinated strategy reduces peak grid demand, daily energy expenditure and operational emissions relative to the conventional case while maintaining full use of available onsite renewable generation. The analysis also shows that workload scheduling and storage are complementary: scheduling relocates delay-tolerant demand toward low-price and low-carbon periods, whereas storage provides fast temporal balancing and peak support. The proposed formulation offers a reproducible foundation for larger studies involving multiple data-centre sites, uncertainty, ancillary-service markets, thermal dynamics and carbon-aware computing. The results should be interpreted as an internally consistent methodological demonstration rather than measurements from a commercial facility.

Keywords: Data Centres; Demand Response; Energy Storage; Renewable Energy; Workload Scheduling; Multi-Objective Optimisation; Grid Flexibility; Carbon-Aware Computing.

الملخص:
يؤدي التوسع المتسارع في الحوسبة السحابية والذكاء الاصطناعي إلى زيادة الطلب على الكهرباء في مراكز البيانات، بالتزامن مع تفاقم ازدحام الشبكات الكهربائية المحلية وتأخر عمليات الربط بها. ومع ذلك، تمتلك مراكز البيانات مجموعة من الموارد القابلة للتحكم، تشمل أحمال الحوسبة القابلة للتأجيل، وأنظمة تخزين الطاقة بالبطاريات، وتوليد الطاقة المتجددة في الموقع، وأنظمة التبريد المرنة، مما يتيح لها العمل بوصفها موارد طاقة تفاعلية مع الشبكة. تطوّر هذه الدراسة إطار

تحسين خطي متعدد الأهداف لتنسيق تشغيل هذه الموارد. ويهدف النموذج المقترح بصورة متزامنة إلى تقليل تكاليف الكهرباء، والانبعاثات الكربونية التشغيلية، والطلب الأقصى على الشبكة، مع الحفاظ على متطلبات خدمات الحوسبة، والحدود التشغيلية للبطارية، وتوازن القدرة. استُخدمت دراسة حالة توضيحية وشفافة تمتد على مدار 24 ساعة لعرض فاعلية الإطار المقترح. وقورنت أربع استراتيجيات تشغيلية، هي: التشغيل التقليدي، وترحيل أحمال الحوسبة زمنياً، والمرونة المعتمدة على البطارية فقط، والتحسين المنسق. وفي ظل الافتراضات الاصطناعية المحددة، أسهمت استراتيجية التحسين المنسق في خفض ذروة الطلب على الشبكة، وتكاليف الطاقة اليومية، والانبعاثات التشغيلية مقارنةً بحالة التشغيل التقليدي، مع المحافظة على الاستفادة الكاملة من الطاقة المتجددة المتاحة في الموقع .كما توضح النتائج أن جدولة أحمال الحوسبة وتخزين الطاقة يمثلان موردين متكاملين؛ إذ تعمل الجدولة على نقل الطلب القابل للتأجيل إلى الفترات التي تنخفض فيها أسعار الكهرباء والانبعاثات الكربونية، في حين يوفر التخزين موازنة زمنية سريعة ودعماً خلال فترات ذروة الطلب. ويوفر الإطار المقترح أساساً قابلاً لإعادة التطبيق في دراسات أوسع تشمل مواقع متعددة لمراكز البيانات، وحالات عدم اليقين، وأسواق الخدمات المساندة، والديناميكيات الحرارية، والحوسبة المراعية للانبعاثات الكربونية. وينبغي تفسير النتائج بوصفها عرضاً منهجياً متسقاً داخلياً، وليس قياسات فعلية مأخوذة من منشأة تجارية.

الكلمات المفتاحية :مراكز البيانات؛ الاستجابة للطلب؛ تخزين الطاقة؛ الطاقة المتجددة؛ جدولة أحمال الحوسبة؛ التحسين متعدد الأهداف؛ مرونة الشبكة؛ الحوسبة المراعية للانبعاثات الكربونية.

1. Introduction

Data centres accounted for approximately 1.5% of global electricity consumption in 2024, corresponding to nearly 415 terawatt-hours (TWh). The United States represented the largest share of this consumption, at approximately 45%, followed by China and Europe at 25% and 15%, respectively. Since 2017, global electricity consumption by data centres has increased at an average annual rate of approximately 12%, exceeding the growth rate of total electricity demand by more than fourfold [1-3]. AI-oriented data centres can exhibit electricity requirements comparable to those of highly energy-intensive industrial facilities, including aluminium smelters. Unlike many conventional industrial loads, however, data-centre capacity is highly concentrated geographically. In the United States, nearly half of the installed data-centre capacity is located within only five regional clusters. Consequently, although data centres currently represent a relatively modest proportion of global electricity consumption, their spatial concentration can impose substantial pressure on local electricity markets, generation resources, and transmission and distribution networks.

Global electricity demand from data centres is projected to more than double by 2030, reaching approximately 945 TWh, slightly exceeding Japan’s current annual electricity consumption. The rapid expansion of AI workloads is expected to be the principal driver of this increase, alongside the continued growth of cloud computing and other digital services. The United States is projected to account for the largest proportion of the additional demand, followed by China. Within the United States, data centres could contribute nearly half of the total growth in electricity demand through 2030 [4-6]. By the end of the decade, electricity consumed by American data centres may surpass the combined electricity requirements of aluminium, steel, cement, chemical production, and other energy-intensive manufacturing activities. Although projections become increasingly uncertain beyond 2030 because of variations in AI adoption, computing efficiency, hardware development, and energy-infrastructure availability, the base-case scenario indicates that global data-centre electricity consumption could reach approximately 1,200 TWh by 2035 [7-10].

Data centres contain several sources of flexibility. Delay-tolerant AI training, batch analytics and maintenance tasks can be shifted across hours. Geographically distributed operators can migrate selected workloads among regions [11-13]. Uninterruptible power-supply batteries can be dispatched beyond their backup role when reliability reserves are preserved [14-19]. Cooling can be pre-cooled or supported by thermal storage, and onsite photovoltaics (PV) can reduce grid imports [20-24]. Studies and industry programmes increasingly recognise these facilities as potential flexible loads rather than passive consumers [25,26].

Existing energy-efficiency research has largely focused on reducing facility overhead through improved PUE, server virtualisation and advanced cooling [27]. These measures remain indispensable, but efficiency and flexibility should not be treated as substitutes. Efficiency lowers the total energy required to deliver computing services, whereas flexibility changes the timing or location of the remaining demand. A highly efficient facility can still worsen a local evening peak, while a flexible facility with modest storage may alleviate congestion without reducing its daily computing output [28,29]. An integrated planning framework must therefore evaluate energy, power, time and location simultaneously.

However, flexibility decisions involve competing objectives. Aggressive price arbitrage may increase emissions when cheap electricity is carbon intensive; carbon minimisation can increase operating cost; peak reduction can require battery cycling or workload delays; and excessive workload interruption can reduce expensive server utilisation. A suitable control framework must represent these trade-offs explicitly. This paper contributes: (i) an integrated model of grid imports, flexible workloads, onsite PV, battery storage and renewable curtailment; (ii) a normalised multi-objective formulation covering cost, operational emissions and peak demand; (iii) operational constraints that maintain computing energy requirements and cyclic battery state of charge; and (iv) a reproducible illustrative case study comparing four flexibility strategies. The framework is intentionally linear and transparent so that it can be extended to stochastic, mixed-integer or multi-site applications.

2. Grid-Interactive Data Centres and Energy-Flexibility Resources

2.1 Electrical demand and operating characteristics

A data centre's facility demand consists primarily of information-technology (IT) equipment and supporting infrastructure. Servers often dominate total consumption, while cooling commonly represents 10–30% depending on climate, technology and power usage effectiveness (PUE) [30-33]. Network equipment, lighting and auxiliary systems contribute smaller shares. IT load can be divided into inflexible, latency-sensitive services and flexible jobs whose start time or location can change within a service-level window [34,35]. Moreover, reliability requirements shape all flexibility decisions. Grid-interactive operation must not consume the battery reserve needed for ride-through, violate redundancy policies, or delay critical applications [36-43]. Consequently, flexibility should be defined as bounded modulation around a secure operating envelope, not unrestricted load shedding.

2.2 Demand response through computational workload management

Temporal workload shifting schedules delay-tolerant jobs during periods of low price, low marginal emissions, abundant renewable generation or reduced network stress. AI training and batch inference are promising because some jobs can be queued or checkpointed. Spatial workload shifting transfers eligible tasks between data centres, creating a virtual interconnector between power systems [3]. Practical limits include data sovereignty, network latency, software architecture, server availability and contractual service levels [44].

Workload flexibility is heterogeneous. Interactive search, financial transactions and real-time control applications have strict response-time requirements and should normally remain inflexible. By contrast, model training, data preprocessing, backups, software builds and some batch-inference tasks can tolerate delays ranging from minutes to hours [45]. Treating all computation as one homogeneous flexible block would overstate the resource available to the grid. Operational models could classify jobs by release time, deadline, power requirement, checkpointing capability and migration cost. The aggregate energy constraint adopted here is a transparent first approximation; a production implementation would use queue-level constraints and explicit penalties for delay or interruption [46].

Spatial migration can expand the flexibility envelope beyond the physical resources of one facility. A geographically diversified operator could route suitable tasks to a region with available renewable generation or lower congestion. Nevertheless, spatial shifting is not automatically carbon beneficial. The destination may have a cleaner average mix but a high marginal emission rate at the relevant hour, and additional data transfer consumes network energy [47]. Moreover, regional transfers may be limited by latency, bandwidth, jurisdictional rules and the availability of equivalent accelerators. Above-mentioned considerations support a multi-site extension in which migration is a bounded decision variable rather than a cost-free action.

2.3 Battery and thermal energy storage

Most large data centres already employ energy storage for uninterruptible power supply. Grid-interactive control can reserve a mandatory state of charge for reliability while using the remaining capacity for peak shaving, renewable-energy absorption and price or carbon arbitrage. Battery degradation, round-trip losses and market rules must be considered in deployment studies. Thermal storage offers an additional pathway by shifting cooling electricity; its detailed thermodynamic representation is outside the present model but can be incorporated as another storage state [48,49].

The economic value of battery flexibility depends on more than the energy-price spread. Avoided demand charges, accelerated grid connection, ancillary-service revenue and deferred network reinforcement may be more important than arbitrage alone. Conversely, cycling can accelerate degradation and reduce the lifetime of assets installed primarily for reliability. A complete investment model should therefore distinguish power-related value from energy-related value and include a degradation term linked to throughput, depth of discharge, temperature and calendar ageing [50,51].

The present operational model omits this cost to preserve linear transparency, so its battery schedules represent technical potential rather than a complete business case.

Cooling flexibility may be particularly attractive because thermal inertia can shift electrical demand without interrupting computation. Pre-cooling, chilled-water tanks and ice storage can move compressor demand away from grid peaks, but the available duration depends on climate, cooling architecture, server inlet limits and thermal-storage size. Higher temperature setpoints may reduce cooling electricity but can also affect equipment reliability [52-55]. Future co-optimisation should connect server utilisation, heat generation, cooling coefficient of performance and indoor thermal states, thereby preventing schedules that are electrically feasible but thermally unrealistic.

2.4 Onsite renewable generation and grid coordination

Onsite PV directly offsets imported electricity, but its variability can create curtailment when generation exceeds coincident demand or export is restricted. Coordinating workloads and batteries increases self-consumption. At the system level, the value of flexibility depends on timely grid signals. Prices, marginal emission factors, congestion alerts or explicit flexibility requests must arrive early enough for workload and storage scheduling. Clear baselines, telemetry and settlement rules are therefore prerequisites for participation in demand-response programmes [56-58].

Renewable matching should be evaluated at an appropriate temporal resolution. Annual procurement can balance total consumption and contracted clean generation while leaving hourly demand dependent on carbon-intensive or constrained supply. Hourly or sub-hourly matching provides a stronger operational signal but requires better forecasts, metering and scheduling. For data centres, the most credible strategy combines additional renewable procurement with temporal demand flexibility, storage and transparent treatment of residual grid emissions [59-62].

Table 1. Grid-interactive flexibility options and principal constraints.

Resource	Grid service	Key constraint
Temporal workload shifting	Peak reduction; renewable matching; cost and carbon response	Deadlines, checkpointing, server utilisation and service-level agreements
Spatial workload migration	Regional congestion relief and carbon-aware routing	Latency, bandwidth, data sovereignty and capacity at destination
Battery storage	Fast response, peak shaving, backup and renewable absorption	Reserve requirement, degradation, efficiency and interconnection rules
Cooling/thermal storage	Short-duration load shifting	Temperature limits, thermal inertia and climate dependence
Onsite renewables	Reduced imports and operational emissions	Variability, land/roof area and curtailment/export limits

3. Proposed Multi-Objective Optimisation Framework

3.1 System boundary and assumptions

The scheduling horizon is divided into T hourly intervals. The facility has an inflexible base load, a schedulable workload, onsite PV, a battery and a grid connection. Export is excluded in the illustrative case, although it can be added. Forecasts of price, grid carbon intensity, inflexible demand and renewable generation are treated as known. The battery model uses constant charging and discharging efficiencies. Workload completion is represented by a daily energy requirement and hourly bounds.

The selected boundary is operational rather than life-cycle based. It accounts for emissions associated with imported electricity during use but excludes embodied emissions from servers, batteries, buildings and PV modules. This distinction is necessary because an operational schedule cannot reverse manufacturing impacts already incurred. Nevertheless, investment studies should add embodied emissions and equipment replacement, especially when comparing larger batteries with software-based workload flexibility. Water consumption and waste heat are also outside the objective set, although both may be important in water-stressed regions or urban heat-recovery projects.

3.2 Decision variables and objective functions

For each interval t, the decision variables are grid import P^g_t, battery charging P^c_t and discharging P^d_t, battery state of charge E_t, scheduled flexible computing demand P^f_t and renewable curtailment P^u_t. An auxiliary variable P^{peak} represents the maximum grid import.

The first objective minimises electricity expenditure:

$$J_1 = \Sigma_t \ c_t \ P^g_t \ \Delta t, \tag{1}$$

where c_t is the time-dependent electricity price. The second objective minimises operational emissions:

$$J_2 = \Sigma_t \ \gamma_t \ P^g_t \ \Delta t, \tag{2}$$

where γ_t is the grid carbon intensity. The third objective minimises connection stress:

$$J_3 = P^{peak}. \tag{3}$$

The normalised weighted-sum objective is:

$$\min J = w_1(J_1/J_{1,ref}) + w_2(J_2/J_{2,ref}) + w_3(J_3/J_{3,ref}), \tag{4}$$

where non-negative weights sum to one and the reference quantities place the objectives on comparable scales. A Pareto-front approach may be preferred when decision makers do not wish to specify weights in advance.

Normalisation is essential because cost, emissions and power are expressed in different units and numerical ranges. Without reference values, the objective with the largest magnitude could dominate even when its assigned weight is small. Reference values may be obtained from conventional operation, independent single-objective optima or policy targets. The interpretation of the weights must also be explicit: they express decision preferences rather than physical constants. Utilities may assign greater importance to peak reduction during capacity scarcity, while operators in regions with strong carbon policies may prioritise emissions. Reporting several weight combinations or an epsilon-constraint Pareto frontier makes the trade-off visible and reduces the risk of presenting one subjective schedule as universally optimal.

Price and carbon signals can either reinforce or oppose one another. Charging a battery during inexpensive coal-dominated hours and discharging during a cleaner but expensive period may reduce cost while increasing emissions. Conversely, strict carbon minimisation may schedule jobs during high-renewable hours even when prices do not fully reflect that environmental value. Including both signals prevents the common assumption that low price is automatically equivalent to low carbon. A fourth objective or constraint could represent workload delay, battery degradation or renewable curtailment when those effects are important to the decision maker.

3.3 Operating constraints

Hourly power balance is imposed by:

$$P^g_t + (P^r_t - P^u_t) + P^d_t = P^b_t + P^f_t + P^c_t, \tag{5}$$

where P^r_t is available renewable generation and P^b_t is inflexible demand. Battery dynamics are:

$$E_t = E_{t-1} + \eta^c P^c_t \ \Delta t - (P^d_t \ \Delta t/\eta^d), \tag{6}$$

Subject to minimum and maximum state of charge, charge/discharge power limits and a cyclic terminal condition $E_T = E_0$. In practical applications, binary variables can prevent simultaneous charging and discharging; the positive efficiency losses and objective coefficients avoid that behaviour in the present linear case.

The cyclic state-of-charge condition prevents the optimiser from obtaining an artificial end-of-horizon benefit by emptying the battery during the final interval. For rolling operation, a terminal reserve or terminal-value function may be more appropriate. Reliability can be represented by imposing a minimum reserve above the normal lower state-of-charge limit, with the reserve varying according to outage risk and required ride-through duration. Network connection limits, ramping limits and minimum notification times can be added without changing the overall framework. Flexible workload requirements are represented by:

$$\Sigma_t \ P^f_t \ \Delta t = W, \quad P^{f,min}_t \le P^f_t \le P^{f,max}_t, \tag{7}$$

where W is the energy required to complete all schedulable jobs. Deadline-aware queues can replace this aggregate constraint. Finally:

$$P^g_t \le P^{peak}, \quad 0 \le P^u_t \le P^r_t. \tag{8}$$

3.4 Solution procedure and implementation

The formulation is a linear programme and can be solved using standard optimisation packages. Each scheduling cycle follows five steps: update forecasts and workload queues; set technical and service constraints; solve the selected weighted or epsilon-constraint problem; dispatch feasible near-term actions; and repeat with updated measurements. Model-predictive control is recommended for operational deployment because it corrects forecast errors and state deviations.

A practical architecture separates planning from real-time protection. The optimisation layer can issue target profiles at hourly or fifteen-minute resolution, while local controllers maintain voltage, frequency, thermal and UPS constraints at faster timescales. If a workload forecast changes or a grid event occurs, model-predictive control re-solves the remaining horizon using the measured state of charge and updated job queue. This hierarchy allows economic optimisation without replacing safety-critical facility controls. It also creates an auditable record of requested and delivered flexibility for settlement with a utility or market operator.

4. Case Study, Scenarios, and Performance Evaluation

4.1 Illustrative data and assumptions

A synthetic 24-hour case demonstrates the model without implying access to confidential commercial data. The nominal facility demand combines a 6.0–9.3 MW inflexible profile and 24 MWh of schedulable computing energy. Hourly schedulable demand may vary from 0.25 to 1.75 MW while preserving total work. The onsite PV system produces 31.25 MWh during the day. The battery has 8 MWh usable capacity, a 3 MW power rating and 95% charge and discharge efficiencies. Prices range from USD 49/MWh overnight to USD 145/MWh during the evening peak. Grid carbon intensity ranges from 0.30 to 0.60 tCO_2/MWh. Export is unavailable, and the battery ends the day at its initial state of charge.

Table 2. Main parameters of the illustrative case study.

Parameter	Value	Interpretation
Scheduling horizon	24 hourly intervals	Day-ahead operational schedule
Inflexible demand	6.0–9.3 MW	Synthetic IT and supporting load
Flexible workload	24 MWh; 0.25–1.75 MW	Daily work conserved within hourly bounds
Onsite PV	31.25 MWh/day; 4.7 MW peak	No export; curtailment allowed
Battery	8 MWh / 3 MW	$\eta c = \eta d = 0.95$; cyclic SOC
Electricity price	USD 49–145/MWh	Synthetic time-of-use profile
Grid carbon intensity	0.30–0.60 tCO_2/MWh	Synthetic marginal/operational signal

4.2 Evaluated operating scenarios

Four strategies are compared. S0 fixes the flexible load at 1 MW and does not use storage. S1 enables temporal workload shifting. S2 fixes workload but enables the battery. S3 coordinates workload and battery decisions. S1 and S2 use normalised weights of 0.65, 0.20 and 0.15 for cost, emissions and peak demand, respectively; S3 uses 0.55, 0.25 and 0.20 to place greater emphasis on system and environmental performance. These weights are illustrative policy preferences, not universal values.

The scenario design isolates the contribution of each flexibility resource. Comparing S1 with S0 measures the value of software-enabled scheduling without storage losses. Comparing S2 with S0 shows the effect of electrical storage when computing demand remains fixed. Comparing S3 with both intermediate cases tests complementarity: if the coordinated solution performs better than either resource alone, the improvement results from joint timing decisions rather than simple addition of capacities. All scenarios complete the same 24 MWh of flexible computing work and use the same demand, PV, price and carbon profiles, ensuring a consistent comparison.

4.3 Results

The coordinated strategy reduces peak grid demand by 15.0%, daily electricity expenditure by 6.7% and operational emissions by 3.7% relative to conventional operation. Grid energy changes from 175.85 MWh to 177.18 MWh. The battery may increase total imported energy slightly because of round-trip losses, but it can still reduce cost, emissions and connection stress by changing when energy is imported. This distinction is important: flexibility improves the temporal quality of demand even when it does not minimise total consumption. Table 3 highlights optimization results under the stated synthetic assumptions.

Table 3. Optimisation results under the stated synthetic assumptions.

Scenario	Grid energy (MWh)	Peak (MW)	Cost (USD thousand)	Emissions (tCO_2)	PV used (MWh)
S0 Conventional	175.85	10.30	14.25	80.17	31.25
S1 Workload shifting	175.85	9.55	13.93	79.12	31.25
S2 Battery only	177.49	8.78	13.55	78.20	31.25
S3 Coordinated	177.18	8.75	13.29	77.20	31.25

Workload shifting alone reduces the peak from 10.30 MW to 9.55 MW, the daily cost from USD 14.25 thousand to USD 13.93 thousand and emissions from 80.17 to 79.12 tCO_2. Because no storage conversion is involved, daily grid energy remains unchanged. This outcome illustrates the operational value of software-defined flexibility: a modest redistribution of delay-tolerant computing can improve several metrics without additional electrical infrastructure. Its limitation is the hourly scheduling bound; once flexible jobs reach their minimum during stressed periods or maximum during preferred periods, further improvement is unavailable.

The battery-only case produces a larger peak reduction, reaching 8.78 MW, and lowers cost to USD 13.55 thousand and emissions to 78.20 tCO_2. Total grid energy rises to 177.49 MWh because the energy recovered during discharge is lower than the energy absorbed during charging. This is not a contradiction. The battery creates value by moving imports from expensive, carbon-intensive and high-demand hours to more favourable intervals. However, the additional energy and unmodelled degradation show why storage should not be evaluated solely through peak reduction.

Coordinated operation achieves the lowest peak, cost and emissions among the four scenarios: 8.75 MW, USD 13.29 thousand and 77.20 tCO_2. The incremental peak improvement over the battery-only case is small because the 8.75 MW level is close to the binding combination of load, PV and flexibility limits. The stronger improvement in cost and emissions indicates that joint scheduling gives the optimiser more freedom to decide whether a marginal adjustment should be provided by workload movement or battery dispatch. In other words, coordination improves resource allocation even when the physical peak cannot be reduced much further.

All scenarios use the full 31.25 MWh of available PV because onsite generation never exceeds the facility's instantaneous ability to absorb it under the stated profile. Therefore, this case does not demonstrate a curtailment reduction. In a facility with a larger PV array, lower midday demand or restricted export, workload shifting, and battery charging would be expected to increase renewable self-consumption. Retaining the curtailment variable in the formulation is still useful because it enables that behaviour in other datasets without changing the model structure. Figure 1 represents conventional and coordinated grid-import profiles with onsite PV. Figure 2 depicts comparison of peak grid demand and daily electricity expenditure.

4.4 Discussion and sensitivity considerations

4.4.1 Complementarity of computing and electrical flexibility

Workload shifting, and storage provide different but complementary capabilities. Shifting moves computational energy without conversion losses and is therefore attractive when jobs are delay tolerant. Batteries respond rapidly and can support latency-sensitive facilities, but incur efficiency losses and degradation. Coordinated control uses workload flexibility first where service constraints permit and reserves storage for sharper ramps, evening peaks and hours with unfavourable price or carbon signals. The complementarity is also temporal. Workload schedulers normally require advance knowledge of job deadlines and accelerator availability, while batteries can respond almost immediately to a grid signal. A facility could therefore use day-ahead optimisation to schedule flexible jobs and reserve part of the battery for forecast errors or unexpected demand-response events. This layered approach is more robust than relying entirely on either advance workload migration or real-time battery dispatch.

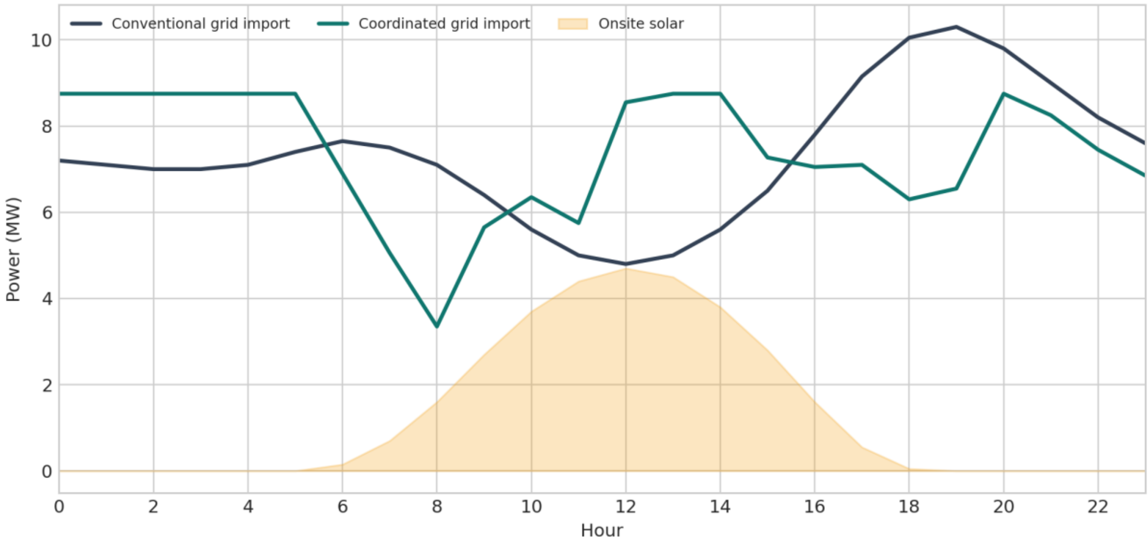

Figure 1. Conventional and coordinated grid-import profiles with onsite PV.

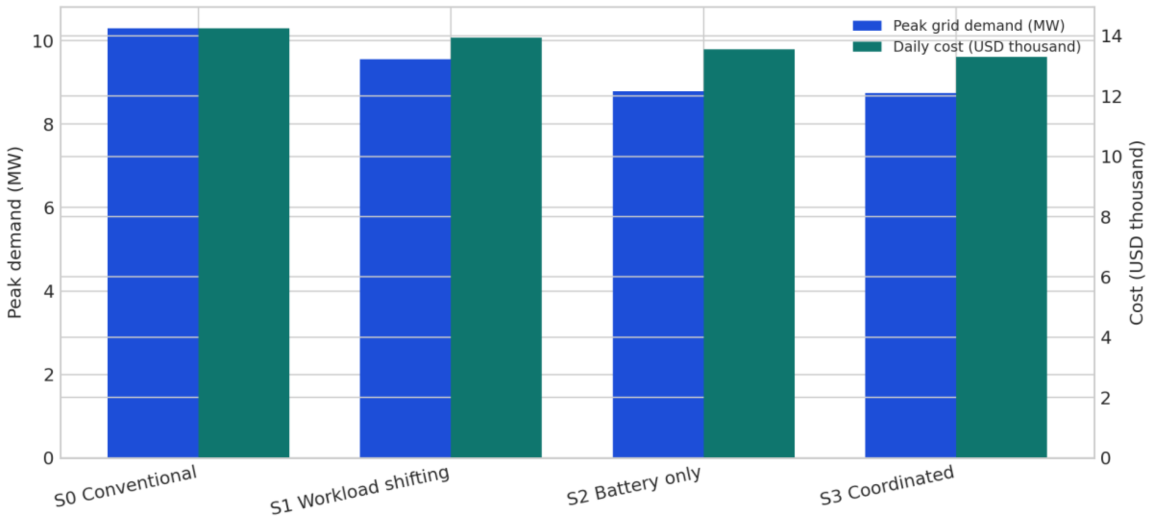

Figure 2. Comparison of peak grid demand and daily electricity expenditure.

4.4.2 Objective conflicts and Pareto interpretation

The numerical outcome is sensitive to the correlation between price, grid carbon intensity and renewable output. When price and carbon signals align, a single schedule can improve both objectives. When they diverge, the Pareto frontier becomes important. Battery size also exhibits diminishing returns once it is sufficient to cover the duration of critical peaks. Likewise, widening workload windows improves flexibility only until server-capacity, deadline or network constraints become binding.

Peak reduction has a different economic meaning from energy-cost minimisation. A lower peak may avoid a connection upgrade, reduce a demand charge or provide capacity during a small number of system-stress hours, even if daily energy cost changes little. This helps explain why a utility may value the coordinated schedule more highly than an energy-only tariff reveals. Conversely, a data-centre operator may be unwilling to reduce accelerator utilisation during expensive but commercially important computing periods. A negotiated flexibility product should therefore compensate the opportunity cost of curtailed or delayed computation, not merely the avoided electricity consumption.

4.4.3 Sensitivity to battery, workload and renewable parameters

Increasing battery energy capacity extends the duration over which demand can be shifted, while increasing its power rating strengthens short peak shaving. The preferred design depends on whether the binding problem is a brief network peak or a sustained renewable mismatch. Efficiency improvements reduce the grid-energy penalty of storage, whereas degradation cost discourages unnecessary cycling. Sensitivity analysis should vary these parameters jointly because a high-power battery with insufficient energy capacity and a high-energy battery with insufficient converter power provide different services.

The flexible-workload bounds represent both technical and commercial constraints. Expanding the upper bound permits more work during favourable hours, but may require spare server capacity; lowering the minimum permits stronger demand response, but may underutilise expensive accelerators. Deadline length, checkpoint overhead and migration bandwidth can be translated into these bounds or represented directly in a mixed-integer job-scheduling model. As flexibility increases, the marginal benefit will eventually decline because price, carbon or peak conditions are already flattened.

Renewable capacity changes the role of both resources. With small PV penetration, nearly all generation is self-consumed and the principal value of flexibility is price, carbon and peak management. At higher penetration, midday curtailment can emerge, making workload concentration and battery charging more valuable. If export is permitted, the optimiser must compare self-consumption with export revenue and grid constraints. Locational marginal prices and nodal emission factors would improve this decision in congested networks.

4.4.4 Uncertainty, rebound effects and validation

The case study is deliberately deterministic. Real facilities face uncertain renewable output, workload arrivals, energy prices and grid events. Robust or stochastic optimisation can incorporate forecast distributions. Additional extensions should include battery degradation cost, cooling thermodynamics, backup-reserve constraints, demand charges, ancillary-service revenues, network nodal limits and geographically distributed data centres.

Forecast error can reduce both savings and reliability. Underestimated workload arrivals may leave insufficient computing capacity after jobs have been shifted, while overestimated PV can cause the battery to charge from the grid unexpectedly. Chance constraints can limit the probability of violating a peak or service deadline, and distributionally robust methods can protect against poorly characterised uncertainty. Scenario-based evaluation should report not only expected performance but also worst-case cost, emissions, unmet workload and reserve shortfall.

Efficiency gains may also create rebound effects. Cheaper and more available computing can stimulate additional AI demand, offsetting part of the energy saved per task. The operational optimisation presented here holds required computing work constant and therefore does not capture long-term demand growth. Strategic assessments should distinguish short-run scheduling benefits from economy-wide changes in computing demand, generation investment and network expansion.

Validation should proceed in stages. First, historical demand, workload, PV, price and carbon data can be replayed offline to compare optimised and observed operation. Second, a hardware- or controller-in-the-loop test can verify communication delays and reserve logic. Third, a limited field pilot can deliver non-critical flexibility while monitoring service-level compliance, battery health and thermal conditions. Only after these steps should savings be generalised to a commercial fleet.

5. Results Implications, Policy Recommendations, and Conclusions

5.1 Implications for data-centre operators

Operators should classify workloads by latency, deadline, interruptibility, location and data-governance constraints. This creates a verifiable flexibility envelope. Energy-management systems should then co-optimise computing, cooling and electrical assets while maintaining explicit reliability reserves. Carbon-aware scheduling should use marginal or operational signals where available rather than annual average emissions alone. Battery dispatch policies should incorporate degradation and preserve emergency autonomy.

Organisational integration is as important as the optimiser. Computing teams typically prioritise service availability and accelerator utilisation, whereas facilities teams focus on electrical reliability and cooling. A grid-interactive programme requires shared performance indicators and escalation rules so that energy actions never bypass service or safety responsibilities. The business case should include avoided connection delay, demand charges, market revenue and sustainability value alongside software development, metering, storage degradation and cybersecurity costs.

5.2 Implications for utilities and regulators

Utilities can use flexible interconnection agreements to connect facilities more rapidly while retaining the right to request limited reductions during critical hours. Such agreements require transparent event frequency, notification time, compensation, performance measurement and non-performance provisions. Regulators should prevent the transfer of avoidable network costs to smaller customers, reward verifiable flexibility and permit qualified behind-the-meter assets to participate in relevant markets. Locational incentives can guide new facilities toward areas with renewable resources, available network capacity and opportunities for heat reuse.

Interoperability standards are also needed. Grid operators, data-centre energy systems and workload orchestrators must exchange secure, machine-readable signals without exposing sensitive customer or operational data. Cybersecurity is essential because increased connectivity expands the attack surface of both digital and electricity infrastructure.

Utilities should distinguish dependable flexibility from theoretical technical potential. A resource should be credited according to tested availability, response time, duration, notification requirements and recovery effects. Baselines must reflect normal operation so that operators are not rewarded for demand reductions that would have occurred anyway. Flexible interconnection contracts can specify a maximum number and duration of events, protecting both grid reliability and the data centre's commercial planning.

Policy should also address distributional effects. Large new loads may require substations, transmission upgrades and additional firm capacity. If cost allocation is weak, existing customers may bear part of these expenses. Transparent connection charges, long-term procurement commitments and performance-based flexibility incentives can align private investment with system needs. At the same time, excessive restrictions could push computing investment toward onsite fossil generation or regions with weaker environmental standards. Regulation should therefore combine timely connection pathways with credible flexibility, efficiency and clean-energy obligations.

5.3 Limitations

This study provides a methodological demonstration rather than a validated commercial deployment. Its 24-hour inputs are synthetic, the network is represented by a single connection point, and uncertainty, cooling dynamics, battery degradation and detailed computing performance are simplified. Therefore, the numerical results demonstrate internal model behaviour and should not be generalised as sector-wide savings. Validation should use measured facility demand, workload traces, local marginal emissions, tariff data and utility constraints.

The weighted-sum formulation may not reveal non-convex portions of a Pareto frontier when binary operating decisions are introduced. Hourly resolution may also hide short-duration demand spikes, UPS behaviour and ancillary-service dynamics. Carbon-intensity data can be uncertain and conceptually different depending on whether average, marginal or consequential accounting is used. These limitations do not invalidate the framework, but they affect how results should be interpreted and which extensions are necessary for a specific application.

5.4 Conclusion

Data centres need not remain passive high-density loads. When computational scheduling, energy storage and onsite renewables are coordinated, they can reduce peak grid demand, improve renewable utilisation and respond to economic and carbon signals while completing required computing work. The proposed multi-objective formulation provides a transparent way to balance these goals and quantify trade-offs. The illustrative results confirm that workload flexibility and storage are complementary and that minimising total electricity alone is not an adequate measure of grid value. Future research should extend the framework to uncertainty-aware, multi-site and network-constrained operation and validate it with real data. With suitable technical standards and market rules, flexible data centres can become active participants in secure, affordable and lower-emissions electricity systems.

Author Contributions: Authors have contributed significantly to the development and completion of this article.

Funding: This article received no external funding.

Data Availability Statement: Not applicable.

Acknowledgments: The author would like to express their sincere gratitude to Libyan Center for Sustainable Development Research, AL-khums, Libya for their invaluable support and resources throughout the course of this research.

Conflicts of Interest: The author(s) declare no conflict of interest.

References

[1] X. Chen, X. Wang, A. Colacelli, M. Lee, and L. Xie, "Electricity demand and grid impacts of AI data centers: Challenges and prospects," *arXiv [eess.SY]*, 2025.

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

[3] Y. Van Fan, C. Wilson, G. Kamiya, and A. Mastrucci, "Data centre energy demand projections within shared socioeconomic pathways," *Energy Clim. Chang.*, vol. 7, no. 100253, p. 100253, 2026.

[4] R. Jha, R. Jha, and M. Islam, "Forecasting US data center CO2 emissions using AI models: emissions reduction strategies and policy recommendations," *Front. Sustain.*, vol. 5, 2025.

[5] H. Wojtaszek, “Energy transition 2024–2025: New demand vectors, technology oversupply, and shrinking Net-Zero 2050 premium,” *Energies*, vol. 18, no. 16, p. 4441, 2025.

[6] A. Ashraf and M. Sagheer, “Renewable energy capacity and technological innovations: A review of global trends and future directions,” *Environ. Prog. Sustain. Energy*, vol. 44, no. 6, 2025.

[7] J. Bolaños-Zuñiga and A. J. Lamadrid, “Environmental and economic implications of artificial intelligence data centers in the United States,” *arXiv [econ.GN]*, 2026.

[8] P. Devarakota, N. Tsesmetzis, F. O. Alpak, A. Gala, and D. Hohl, “AI and the net-zero journey: Energy demand, emissions, and the potential for transition,” *arXiv [cs.AI]*, 2025.

[9] T. Lu, X. Lü, D. Clements-Croome, Y. Yuan, and B. Martinkauppi, “Developing an active heat recovery system to boost data centre waste heat utilization: a novel and cost-effective solution for building decarbonisation,” *Energy (Oxf.)*, vol. 341, no. 139324, p. 139324, 2025.

[10] S. Hlabisa, “The ecology of artificial intelligence: energy, water, materials, and land limits of digital systems,” *Carbon Neutral Syst.*, vol. 1, no. 1, 2025.

[11] C. Crozier and M. Liska, “The potential of data center energy demand to provide grid flexibility,” *Curr. Sustain./Renew. Energy Rep.*, vol. 12, no. 1, 2025.

[12] M. T. Takci, J. Day, and M. Qadrdan, “Characterisation and quantification of data centre flexibility for power system support,” *arXiv [eess.SY]*, 2025.

[13] R. Buyya, S. Ilager, and P. Arroba, “Energy-efficiency and sustainability in new generation cloud computing: A vision and directions for integrated management of data centre resources and workloads,” *Softw. Pract. Exp.*, vol. 54, no. 1, pp. 24–38, 2024.

[14] Y. F. Nassar *et al.*, “Carbon footprint and energy life cycle assessment of wind energy industry in Libya,” *Energy Convers. Manag.*, vol. 300, no. 117846, p. 117846, 2024.

[15] M. Khaleel *et al.*, “Evolution of emissions: The role of clean energy in sustainable development,” *Chall. Sustain.*, vol. 12, no. 2, pp. 122–135, 2024.

[16] H. J. El-Khozondar *et al.*, “Analysis of a grid/wind/hydrogen hybrid energy system for producing electricity, vehicle fuel and freshwater for Gaza Strip-Palestine,” *Fuel (Lond.)*, vol. 428, no. 140358, p. 140358, 2027.

[17] M. Khaleel *et al.*, “The impact of SMES integration on the power grid: Current topologies and nonlinear control strategies,” in *New Technologies, Development and Application VII*, Cham: Springer Nature Switzerland, 2024, pp. 108–121.

[18] M. Elmnifi, M. Khaleel, Z. Yusupov, Y. F. Nassar, and H. J. El-Khozondar, “Economic and environmental benefits of recycling electrical and electronic waste products in north African countries,” in *Environmental Science and Engineering*, Cham: Springer Nature Switzerland, 2025, pp. 55–93.

[19] E. Mohamed, A. Souli, A. Beladel, and M. Khaleel, “Simulation and analysis of lightning strikes in electrical systems by MATLAB/SIMULINK and ATP/EMTP,” *ITEGAM- J. Eng. Technol. Ind. Appl. (ITEGAM-JETIA)*, vol. 10, no. 47, 2024.

[20] A. Ghayth *et al.*, “Evaluating the effect of electric vehicles charging stations on residential distribution grid,” *ijees*, vol. 2, no. 3, pp. 1–22, 2024.

[21] M. Khaleel, N. El-Naily, H. Alzargi, M. Amer, T. Ghandoori, and A. Abulifa, “Recent progress in synchronization approaches to mitigation voltage sag using HESS D-FACTS,” in *2022 International Conference on Emerging Trends in Engineering and Medical Sciences (ICETEMS)*, 2022.

[22] M. Khaleel, Z. Yusupov, and W. A. Shah, “Artificial intelligence utilization for PQ concerns in power grid: Cybersecurity, optimization and policy implications,” in *2026 IEEE 5th International Maghreb Meeting of the Conference on Sciences and Techniques of Automatic Control and Computer Engineering (MI-STA)*, 2026, pp. 1–6.

[23] M. Khaleel, Z. Yusupov, O. M. Bshina, M. M. Dali, and W. A. Shah, “Toward renewable energy deployment in North African countries: Potential resources, investments, transition scenario, opportunities, challenges, and policy,” *Unconventional Resources*, vol. 12, no. 100363, p. 100363, 2026.

[24] Y. F. Nassar *et al.*, “Thermoelectrical analysis of a new hybrid PV-thermal flat plate solar collector,” in *2023 8th International Engineering Conference on Renewable Energy & Sustainability (ieCRES)*, 2023.

[25] Y. F. Nassar *et al.*, “Design of reliable standalone utility-scale pumped hydroelectric storage powered by PV/Wind hybrid renewable system,” *Energy Convers. Manag.*, vol. 322, no. 119173, p. 119173, 2024.

[26] M. Khaleel, Z. Yusupov, A. Ahmed, A. Alsharif, Y. Nassar, and H. El-Khozondar, "Towards sustainable renewable energy," *Appl. Sol. Energy*, vol. 59, no. 4, pp. 557–567, 2023.

[27] A. Katal, S. Dahiya, and T. Choudhury, "Energy efficiency in cloud computing data center: a survey on hardware technologies," *Cluster Comput.*, vol. 25, no. 1, pp. 675–705, 2022.

[28] G. Fieni, R. Rouvoy, and L. Seinturier, "xPUE: Extending power usage effectiveness metrics for cloud infrastructures," *IEEE Trans. Sustain. Comput.*, pp. 1–13, 2025.

[29] N. Nandhini, V. Manikandan, and S. Elango, "An interpretable generalized additive neural networks for electricity theft detection in smart cities using balanced data and intelligent grid management," *Energy Build.*, vol. 346, no. 116123, p. 116123, 2025.

[30] A. Katal, S. Dahiya, and T. Choudhury, "Energy efficiency in cloud computing data centers: a survey on software technologies," *Cluster Comput.*, vol. 26, no. 3, pp. 1845–1875, 2023.

[31] X. Wang *et al.*, "A review on data centre cooling system using heat pipe technology," *Sustain. Comput. Inform. Syst.*, vol. 35, no. 100774, p. 100774, 2022.

[32] D. Mytton and M. Ashtine, "Sources of data center energy estimates: A comprehensive review," *Joule*, vol. 6, no. 9, pp. 2032–2056, 2022.

[33] S. Long, Y. Li, J. Huang, Z. Li, and Y. Li, "A review of energy efficiency evaluation technologies in cloud data centers," *Energy Build.*, vol. 260, no. 111848, p. 111848, 2022.

[34] H.-A. Ounifi, A. Gherbi, and N. Kara, "Deep machine learning-based power usage effectiveness prediction for sustainable cloud infrastructures," *Sustain. Energy Technol. Assessments*, vol. 52, no. 101967, p. 101967, 2022.

[35] Y. Sun, Y. Wang, G. Jiang, B. Cheng, and H. Zhou, "Deep learning-based power usage effectiveness optimization for IoT-enabled data center," *Peer Peer Netw. Appl.*, vol. 17, no. 3, pp. 1702–1719, 2024.

[36] M. Khaleel, Z. Yusupov, M. Elmnifi, T. Elmenfy, Z. Rajab, and M. Elbar, "Assessing the financial impact and mitigation methods for voltage sag in power grid," *ijees*, vol. 1, no. 3, pp. 10–26, 2023.

[37] A. Hesri, M. Khaleel, A. A. Ahmed, and Y. Nassar, "Grid integration of EVs: Deploying measures, standards and interoperability and policy," *ijees*, vol. 3, no. 3, pp. 1–18, 2025.

[38] A. A. Ahmed, M. Mohamed, Y. Nassar, and A. Alsharif, "Recent trends in optimization objectives for power system operation improvement using FACTS," *ijees*, vol. 3, no. 4, pp. 33–46, 2025.

[39] A. Alkhair and G. A. M. Ghalboun, "Environmental impacts on energy storage systems integrated with renewables," *ijees*, vol. 4, no. 1, pp. 15–31, 2026.

[40] I. Abdaldaem, B. Alskekh, and Z. Mady, "Advancing Power Quality in distribution grids through AI: Opportunities, challenges, optimization, and policy pathways," *ijees*, vol. 3, no. 3, pp. 48–61, 2025.

[41] O. S. Abualoyoun and A. A. Amar, "Advancing the global integration of solar and wind power: Current status and challenges," *ijees*, vol. 3, no. 2, pp. 73–88, 2025.

[42] M. Khaleel, Z. Yusupov, B. Alfalh, M. T. Guneser, Y. Nassar, and H. El-Khozondar, "Impact of smart grid technologies on sustainable urban development: DOI: 10.5281/zenodo.11577746," *ijees*, vol. 2, no. 2, pp. 62–82, 2024.

[43] O. Hassan, Z. Yusupov, and M. M. Abubakr, "Hybrid AI and optimization algorithms for performance enhancement of grid-connected solar PV systems," *ijees*, vol. 4, no. 1, pp. 63–84, 2026.

[44] U. Assad *et al.*, "Smart grid, demand response and optimization: A critical review of computational methods," *Energies*, vol. 15, no. 6, p. 2003, 2022.

[45] V. Arumugham *et al.*, "An artificial-intelligence-based Renewable Energy prediction program for demand-side management in smart Grids," *Sustainability*, vol. 15, no. 6, p. 5453, 2023.

[46] B. Dey, S. Dutta, and F. P. Garcia Marquez, "Intelligent demand side management for exhaustive techno-economic analysis of microgrid system," *Sustainability*, vol. 15, no. 3, p. 1795, 2023.

[47] J. Xie, A. Ajagekar, and F. You, "Multi-Agent attention-based deep reinforcement learning for demand response in grid-responsive buildings," *Appl. Energy*, vol. 342, no. 121162, p. 121162, 2023.

[48] P. Xie, L. Jin, G. Qiao, C. Lin, C. Barreneche, and Y. Ding, "Thermal energy storage for electric vehicles at low temperatures: Concepts, systems, devices and materials," *Renew. Sustain. Energy Rev.*, vol. 160, no. 112263, p. 112263, 2022.

[49] M. Brandt, J. Woods, and P. C. Tabares-Velasco, "An analytical method for identifying synergies between behind-the-meter battery and thermal energy storage," *J. Energy Storage*, vol. 50, no. 104216, p. 104216, 2022.

[50] C. Zhao, J. Yan, X. Tian, X. Xue, and Y. Zhao, “Progress in thermal energy storage technologies for achieving carbon neutrality,” *Carb Neutrality*, vol. 2, no. 1, 2023.

[51] H. Tang and S. Wang, “Life-cycle economic analysis of thermal energy storage, new and second-life batteries in buildings for providing multiple flexibility services in electricity markets,” *Energy (Oxf.)*, vol. 264, no. 126270, p. 126270, 2023.

[52] M. Khaleel *et al.*, “Battery technologies in electrical power Systems: Pioneering secure energy transitions,” *J. Power Sources*, vol. 653, no. 237709, p. 237709, 2025.

[53] M. Amer, A. Elmojarrush, Y. F. Nassar, and M. Khaleel, “Critical materials for EV batteries: Challenges, opportunities, and policymakers,” *ijees*, vol. 3, no. 1, pp. 119–133, 2025.

[54] M. Khaleel, Z. Yusupov, N. Yasser, and H. J. El-Khozondar, “Enhancing Microgrid performance through hybrid energy storage system integration: ANFIS and GA approaches,” *ijees*, vol. 1, no. 2, pp. 38–48, 2023.

[55] I. Imbayah, O. A. Eseid, M. M. Beiek, K. Akter, A. Alsharif, and A. Ali Ahmed, “Recent developments in EV charging infrastructure: Opportunities and IoE framework challenges,” *ijees*, vol. 1, no. 4, pp. 47–63, 2023.

[56] K.-X. Zhao, D.-X. Chen, T.-T. Wang, Y. Zhou, and T. Ma, “Institutional coordination expansion in power sector reshapes renewable energy system layout and trans-administrative region power grid,” *Appl. Energy*, vol. 402, no. 126857, p. 126857, 2025.

[57] L. Richter, V. Lenz, M. Dotzauer, and J. Seifert, “Coordinated energy systems in decentralized districts: Evaluating the cellular approach for improved grid stability and renewable integration,” *Smart Energy*, vol. 20, no. 100215, p. 100215, 2025.

[58] M. Alanazi, “A sustainable approach to renewable integration in distribution grids through coordinated control of flexibility technologies,” *Clean. Eng. Technol.*, vol. 28, no. 101053, p. 101053, 2025.

[59] M. Khaleel, I. Imbayah, Y. Nassar, and H. J. El-Khozondar, “Renewable energy transition pathways and net-zero strategies,” *ijees*, vol. 3, no. 4, pp. 01–16, 2025.

[60] O. S. Abualoyoun and A. A. Amar, “Toward achieving net-zero in the evolution of the global energy sector,” *ijees*, vol. 3, no. 1, pp. 143–156, 2025.

[61] M. Khaleel and M. Elbar, “Exploring the rapid growth of solar photovoltaics in the European Union,” *ijees*, vol. 2, no. 1, pp. 61–68, 2024.

[62] E. Salim, “Exploring the contribution of wind energy to the electrical load at wadi alshatti university,” *ijees*, vol. 2, no. 2, pp. 46–61, 2024.