Energy Reports 14 (2025) 5061–5071

Contents lists available at ScienceDirect

Energy Reports

journal homepage: www.elsevier.com/locate/egyr

Research paper

Multi-agent systems for optimising smart energy clusters: A case study on cost and emission reduction in industrial seaport facilities

Amin Amin [a,*], Ali Ghoroghi [a], Ioan Petri [a], Andrei Hodorog [a], Thomas Genest [b], Yacine Rezgui [a]

[a] School of Engineering, Cardiff University, Cardiff, CF24 3AA, United Kingdom
[b] Independent Researcher, Paris, France

ARTICLE INFO

Keywords:
Multi-agent
Energy optimisation
Building cluster
Peer-to-peer
Decarbonisation
Demand response

ABSTRACT

This study investigates the use of multi-agent systems (MAS) to optimise energy flow in industrial seaport clusters, with an emphasis on reducing operational costs, minimising carbon emissions, and increasing renewable energy utilisation. The proposed framework combines decentralised control via MAS, dynamic load scheduling, and smart grid interaction across three operational strategies: individual building-level optimisation, cluster-level coordination, and peer-to-peer energy trading. Simulation results from a UK-based pilot site indicate that cluster-level optimisation achieves an average 60% cost saving and 30% emission reduction. The peer-to-peer strategy enables up to 90% renewable self-consumption, reducing grid dependence by 30–35%. In contrast, the individual strategy remains sensitive to demand fluctuations, with grid dependency rising to 70% under high-load conditions due to limited energy sharing and battery saturation.

These findings demonstrate the scalability and adaptability of MAS-based energy frameworks in industrial contexts. By critically evaluating performance under variable operational scenarios, this study offers practical insights for sustainable energy management and industrial decarbonisation pathways.

1. Introduction

Demand-side energy management in building clusters presents a promising pathway for enhancing both economic performance and environmental sustainability, while accelerating the transition to low-carbon industrial systems (Kaspar et al., 2022; Amin et al., 2022). The increasing deployment of renewable energy sources (RES) requires intelligent energy management solutions to mitigate production intermittency, ensure operational stability, and optimise energy use and associated costs (Amin et al., 2026). Recent advances in artificial intelligence (AI) and cyber–physical systems have enabled the emergence of decentralised energy management paradigms, where interconnected actors coordinate energy flows, dynamically adjust loads, and engage in peer-to-peer (P2P) trading to increase renewable self-consumption and reduce grid dependency (Huang et al., 2020).

Industrial seaport facilities — especially those supporting fisheries and food processing — face unique energy and environmental challenges. These facilities are often located in coastal or remote areas, characterised by ageing infrastructure, limited grid access, and elevated electricity tariffs (Food and Agriculture Organization of the United Nations, 2020; Alqarni et al., 2023). In response, the concept of the "smart port" has emerged, underpinned by the digitalisation of energy infrastructure and operational processes. Yang et al. (2018) describe smart ports as cyber–physical ecosystems integrating sensor networks, data platforms, and energy control mechanisms to enable intelligent decision-making. Buiza et al. (2015) further highlight that smart ports prioritise environmental efficiency, operational reliability, and economic competitiveness.

This study builds upon the theoretical foundations of multi-agent systems (MAS) and distributed optimisation. MAS enables autonomous coordination among energy entities — such as buildings, storage units, and local generators — by facilitating real-time decision-making and inter-agent communication, which is suitable for complex energy systems with heterogeneous actors, temporal variability, and spatial distribution of resources (Du et al., 2023; Bregar, 2020).

Despite extensive literature on microgrids and industrial energy management, energy-intensive fishing facilities have received limited attention despite their high energy volatility, reliance on fossil fuels, and insufficient integration of renewable resources. Existing approaches often lack scalability, adaptability, and holistic coordination across the

* Corresponding author.
E-mail addresses: AminA8@cardiff.ac.uk (A. Amin), ghoroghia@cardiff.ac.uk (A. Ghoroghi), petrii@cardiff.ac.uk (I. Petri), hodoroga@cardiff.ac.uk (A. Hodorog), genest.thomas.ba@gmail.com (T. Genest), rezguiy@cardiff.ac.uk (Y. Rezgui).

https://doi.org/10.1016/j.egyr.2025.11.003
Received 12 March 2025; Received in revised form 30 June 2025; Accepted 5 November 2025
Available online 28 November 2025

Nomenclature

Acronyms

CO_2	Carbon dioxide
DR	Demand response
EC	Electricity costs
GD	Grid dependency
KPI	Key perfomance indicator
MAS	Multi-agent system
P2P	Peer-to-peer
PV	Photovoltaic
RES	Renewable energy sources
UK	United Kingdom

Units

%	Percentage
£	British Pound Sterling
kWh	Kilowatt-hour
kWp	Kilowatt-peak
MWh	Megawatt-hour
ton	Weight in ton

Variables

A, b	Coefficients for energy flow equation
Aeq, beq	Coefficients for energy balance equation
ci(x)	Non-linear constraint for energy contradiction
L	Number of days
SoC	State of charge
X	Decision three-dimensional vector

cluster level, especially in integrating dynamic P2P trading with multi-layer optimisation (Alzahrani et al., 2021; Iris and Lam, 2021). This paper addresses these gaps by proposing a layered MAS-based optimisation framework tailored to the operational profiles of industrial seaport facilities. The framework integrates three complementary strategies:

- Dynamic load scheduling to optimise energy consumption and production across buildings and time periods.
- Cluster-level coordination to optimise collective energy operations and reduce costs and grid reliance.
- P2P energy trading to facilitate decentralised energy exchange among buildings and maximise on-site renewable self-consumption.

The core contributions of this research are as follows:

1. Articulation of literature gaps in MAS-based energy management specific to industrial seaports and fisheries, contributing to emerging discourse on sustainable port development.
2. Design of a multi-layered MAS framework integrating local control, central coordination, and cooperative trading mechanisms to enhance operational flexibility and system resilience.
3. Evaluation of decentralised P2P energy trading in a real-world industrial cluster, demonstrating its impact on renewable self-consumption and cost savings under different operational conditions.
4. Critical assessment of technical and environmental performance metrics — including cost, emissions, and grid dependence — under various optimisation strategies using validated simulation data.

The paper is structured as follows. Section 3 presents the methodology, pilot site characterisation, and optimisation formulation. Section 4 discusses the simulation results and scenario-specific insights. Conclusions and future research directions are summarised in Section 5.

2. Related work

Recent research highlights a wide range of strategies for energy optimisation and decarbonisation in industrial settings, including smart ports and seafood processing facilities. This section categorises relevant contributions across four primary domains: energy management in seaports, renewable energy integration, peer-to-peer (P2P) energy trading, and multi-agent systems (MAS) for energy optimisation.

Energy management in seaports: Recent research has explored smart energy systems in seaports, particularly focusing on fish processing and cold-chain logistics. Alzahrani et al. (2020) developed a smart microgrid model for the Milford Haven fishing port, incorporating PV panels, battery storage, and a MILP-based optimiser to manage electricity cost and emissions. The study achieved 15.4% electricity cost reductions and 13.7% carbon emissions savings through the integrated scheduling of deferrable loads and storage operation. However, the model relied on static scheduling and assumed full demand forecast accuracy, which may limit its responsiveness under real-time operational fluctuations.

In a broader seaport energy context, Iris and Lam (2021) introduced an integrated energy system combining demand response (DR), energy arbitrage, and multi-vector control in a smart port framework. The model considered various load types (e.g., HVAC, refrigeration, lighting) and optimised power flow from renewable sources and the grid. Their results indicated 15.2% cost savings and a 4.3% improvement in energy autonomy. Notably, their framework lacked explicit agent-based coordination or peer-to-peer mechanisms, instead depending on centralised optimisation for load shifting and battery dispatch.

Although these studies highlight the importance of flexible demand modelling, especially for energy-intensive assets such as like cold rooms and ice-making units, they underscore the growing emphasis on integrating renewable generation and storage within port environments. Further research is required to investigate impacts of fixed price tariffs, decentralised intelligence, and stochastic variability modelling on seaport energy environments.

Renewable energy integration: Several studies explore the integration of renewable energy sources (RES) in port environments. Ramos et al. (2014) assessed the potential of tidal energy to power the port of Ribadeo in Spain, using high-resolution simulation of tidal currents. The methodology included monthly energy output estimation based on flow velocities and turbine efficiency, indicating a yearly energy potential of 94.9 MWh and a carbon reduction of up to 85% compared to conventional diesel generation. The results suggest feasibility for partial port electrification using tidal currents under consistent hydrodynamic conditions.

Similarly, Arena et al. (2018) investigated wave energy converters (WECs) to supply EV charging stations in ports. Their work combined energy production forecasting with EV charging profiles and concluded that 2–3 WEC units could meet 45%–60% of energy demand, although production was seasonally variable. The study highlights the importance of aligning marine energy production with real-time demand profiles.

Buonomano et al. (2023) proposed a dynamic simulation model for the Port of Naples, integrating photovoltaic (PV), solar thermal, and hydrogen energy systems. The framework included multi-domain energy modelling and optimisation via genetic algorithms. The system achieved a renewable self-consumption rate of 84.3% and cut carbon emissions by 59.7%. This approach combined multiple RES and storage systems, offering higher flexibility and energy independence.

Collectively, integrating multiple energy vectors (e.g. tidal, wave, and hybrid solar systems) can enhance energy sustainability in port facilities, with typical outcomes showing up to 85% renewable integration and 60% emission reductions through applying simulation-based potential assessments or real-world optimisation frameworks.

P2P energy trading: P2P models have emerged as viable mechanisms to enhance distributed energy coordination and self-sufficiency in local energy communities. Alzahrani et al. (2021) proposed a P2P energy-sharing algorithm using a dynamic pricing mechanism within smart port microgrids, emphasising agent-based coordination under seasonal demand variations. Their two-stage model, tested at a fishing port, achieved a 21.6% reduction in energy costs and a 9.4% improvement in renewable self-consumption.

May and Huang (2023) explored a MAS framework using reinforcement learning for decentralised P2P trading. The proposed framework incoprated dynamic price signals led to improved system efficiency and user satisfaction. The findings noted the influence of prosumer behaviour on market stability, suggesting behavioural modelling as critical for long-term performance.

Tushar et al. (2020) provided a comprehensive review of market mechanisms in P2P energy systems. They categorised designs into community-based and fully decentralised architectures, identifying trade-offs in computational complexity, trust mechanisms, and scalability. Dynamic pricing was identified as a critical enabler for grid responsiveness and fair compensation. Key limitations include transaction costs, lack of standardised protocols, and absence of legal frameworks for decentralised markets. Otherwise, the integration of distributed ledger technologies (DLT) may enhance traceability and automate settlements through smart agreements (AlSkaif et al., 2021).

Overall, P2P energy trading literature highlights the need for scalable, agent-coordinated, and incentive-compatible models to support energy flexibility in microgrids and industrial settings.

MAS in energy optimisation: Recent advancements have significantly expanded the role of MAS in complex energy environments, especially within industrial and building-scale microgrids. MAS-based frameworks have demonstrated strong potential in handling decentralised decision-making, dynamic energy coordination, and stochastic conditions without requiring centralised control.

Zhu et al. (2022) proposed a MAS deep reinforcement learning (DRL) model for multi-energy management in industrial parks. The system leveraged decentralised execution and centralised training using a soft actor–critic algorithm and attention-based mechanisms to reduce coordination overhead. Their simulation, based on real data, reported improved learning stability and efficiency while minimising long-term energy costs under fluctuating demand conditions. Similarly, Shen et al. (2022) designed a multi-agent DRL framework for building energy systems using duelling deep Q-networks and a value-decomposition approach. Their results demonstrated a 43% reduction in unutilised renewable energy and 8% cost savings compared to rule-based methods, indicating MAS adaptability to supply–demand mismatches and resource constraints.

Incorporating blockchain into MAS energy control has gained attention for enhancing transparency and trust in decentralised environments. AlSkaif et al. (2021) introduced a blockchain-enabled MAS platform for residential P2P trading that supports contract enforcement and transaction traceability. Their findings emphasised reduced transaction delays and improved user trust, which are essential for decentralised energy ecosystems. Further, integrating federated learning (FL) has emerged as a promising approach to MAS-based optimisation. Yang et al. (2025) applied FL to decentralised energy control across buildings, allowing agents to collaboratively train control policies without sharing raw data. This method improved resilience under fluctuating loads and protected data privacy, a critical consideration in competitive industrial environments.

Communication constraints in MAS systems have also been studied. Domyshev et al. (2024) presented a MAS control model addressing both power and communication limitations in microgrids. Their simulation showed that distributed negotiation outperforms centralised scheduling in terms of adaptability and network resilience. Wang et al. (2024) proposed a distributed online optimisation approach using MAS to manage integrated energy systems. Their method dynamically updates local agent strategies under time-varying prices and demand, proving effective in reducing both energy cost and operational complexity in real-time settings.

Hybrid optimisation methods have also been explored in literature. Liu and Tang (2024) combined stochastic programming with generalised Nash bargaining in a MAS framework to manage energy distribution among microgrids. Their simulations demonstrated improved fairness and cost-effectiveness, particularly under uncertain renewable generation scenarios.

Finally, Petri et al. (2019) implemented a MAS-based control framework tailored for energy optimisation in fish-processing industries. Their study modelled building-level agents to simulate various operational scenarios, highlighting the value of agent-based modelling in industrial load optimisation and emission reduction.

Despite these advancements, two critical challenges remain in the MAS literature:

- **Limited integration of hybrid renewables**: Most MAS frameworks rely exclusively on solar PV. While sufficient in some contexts, this limits applicability in regions with high potential for tidal, wind, or wave energy, particularly coastal ports. Altough the feasibility of such sources have demonstrated in literature (Arena et al., 2018; Ramos et al., 2014), few MAS models incorporate multi-resource forecasting, control, or optimisation logic. Without hybrid integration, MAS energy management remains constrained in its environmental and economic adaptability.
- **Absence of lifecycle contextualisation**: Operational efficiency alone is insufficient to claim sustainability. Embedded carbon in PV panels, battery manufacture, and disposal must be accounted for to determine net environmental benefit Petri et al. (2025). However, few MAS frameworks embed lifecycle assessment (LCA) metrics into their optimisation objectives, resulting in incomplete sustainability evaluation. Without such metrics, decision-makers may overestimate the benefits of decentralised energy systems, especially over long time horizons.

Overall, MAS represents a frontier technology in industrial energy optimisation. The combination of decentralised autonomy, adaptive learning, and agent-level intelligence provides a flexible architecture for managing increasingly complex energy systems (Adewoyin et al., 2025). As industrial sectors seek scalable solutions to meet decarbonisation targets, MAS — enhanced with federated learning, blockchain trading, and multi-renewable control — offers a viable pathway, particularly when incorporating lifecycle metrics and real-world implementation barriers (Celik et al., 2021; Alzahrani et al., 2021; Iris and Lam, 2021).

To address these gaps, this study investigates the following hypotheses:

- A layered MAS-based framework can effectively reduce energy costs and emissions in industrial seaport environments (Yang et al., 2018; Alzahrani et al., 2021).
- Cluster-level optimisation outperforms individual building-level control in terms of cost efficiency and renewable energy utilisation (Buonomano et al., 2023).
- P2P energy trading enhances renewable self-consumption and grid independence beyond conventional control strategies (Shen et al., 2022; Ramos et al., 2014).

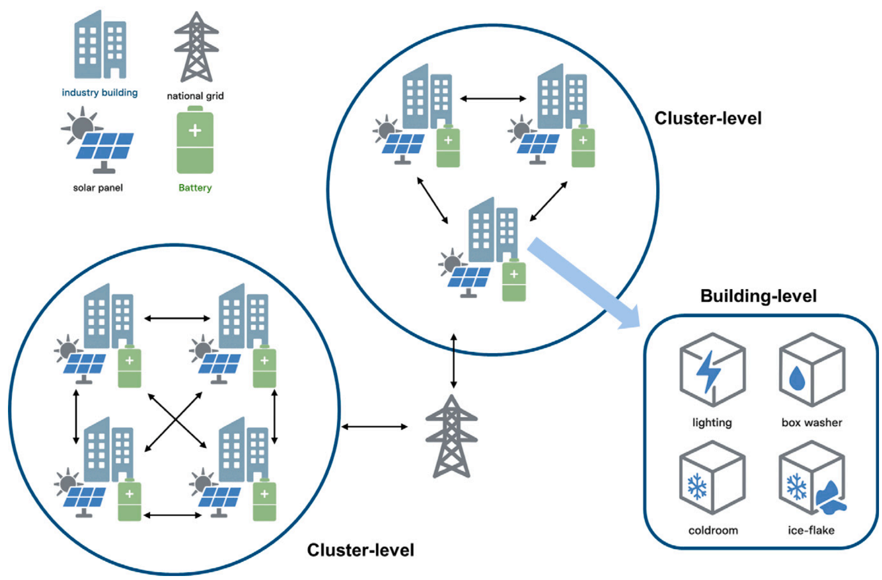

Fig. 1. Layered multi-agent framework for smart energy management across cluster and building levels in industrial seaport facilities.

3. Methodology

Our methodology implements a layered multi-agent structure to explore energy optimisation across multiple operational levels in industrial energy clusters. The proposed framework integrates dynamic load scheduling, decentralised control, and real-time data exchange to minimise operational costs and environmental impact while maximising renewable energy use and grid independence. It leverages a deep reinforcement learning (DRL) scheme that adapts to changing demand, intermittent generation, and dynamic electricity tariffs.

The architecture comprises two hierarchical agent layers, as illustrated in Fig. 1: (i) **building-level agents** independently manage energy production, consumption, and storage within each facility; and (ii) **cluster-level agents** coordinate energy transactions and distributed storage systems across the entire site, facilitating centralised decision-making for enhanced efficiency and resilience.

To capture variations in consumption and generation patterns, the model applies temporal segmentation, dividing the day into four periods (early morning, daytime, afternoon, evening). A week-ahead optimisation horizon is adopted to balance computational cost and predictive accuracy, enabling the evaluation of strategic energy decisions under different operational conditions.

3.1. Pilot and energy modelling

The framework is applied to a case study at Milford Haven port in Wales, UK—one of the country's largest energy hubs, as shown in Fig. 2. Five energy-intensive buildings are modelled (Fig. 2):

- J-Shed: multi-tenanted facility
- Packaway: equipped with a 50 kWp PV system
- K-Shed: includes a 50 kWp solar PV and a monthly consumption of approximately 4000 kWh
- M-Shed: multi-purpose industrial unit
- F-Shed: dedicated to fish processing activities

Each building integrates one or more energy sources — solar PV, battery storage, and the national grid — and supports four operational appliances: cold room, lighting, ice flake machine, and box washer. The model is built on the following key assumptions:

- Energy demand: The model incorporates historical demand patterns and adaptive forecasting to account for future variations, enhanced by real-time monitoring. Real-time monitoring enhances this approach, allowing the system to respond dynamically to both predictable cycles and unexpected fluctuations.
- Renewable energy production: Solar energy generation is modelled using historical irradiance data specific to the pilot location, with adjustments for seasonal changes. The capacity is determined by available rooftop space, ensuring that local renewable resources are fully utilised.
- Battery storage: Each building is equipped with a fixed-capacity battery system based on average energy storage systems in industrial facilities, considering average daily energy use in the pilot site. This configuration provides sufficient capacity for effective load management without incurring excessive costs.
- Charging/discharging rates: The power rates for battery charging and discharging are set based on the operational specifications of industrial battery technologies. The model optimises charge/discharge cycles through a cost-minimisation algorithm that prioritises renewable energy utilisation during peak tariff periods, directly impacting energy cost and grid dependency.

Model validation includes:

- Calibration against measured demand with ±5% accuracy.
- Scenario simulation under normal, high-demand, and high-generation conditions.
- Key performance indicator (KPI) monitoring for energy cost, carbon emissions, and grid reliance.
- Sensitivity analysis on system size (varying PV and battery capacities), demand volatility (using stochastically scaled demand patterns), weather uncertainty (via altered solar irradiance profiles), and dynamic pricing shifts (modifying tariff inputs).

3.2. Optimisation formulation

The optimisation focuses on controlling key energy vectors: (i) energy consumption by buildings; (ii) energy imported from the power grid; (iii) energy exported to the grid; (iv) battery charging/discharging rates; and (v) energy traded among buildings in peer-to-peer (P2P) scenarios. These variables were selected for their direct influence on the system's energy balance and their ability to accommodate dynamic

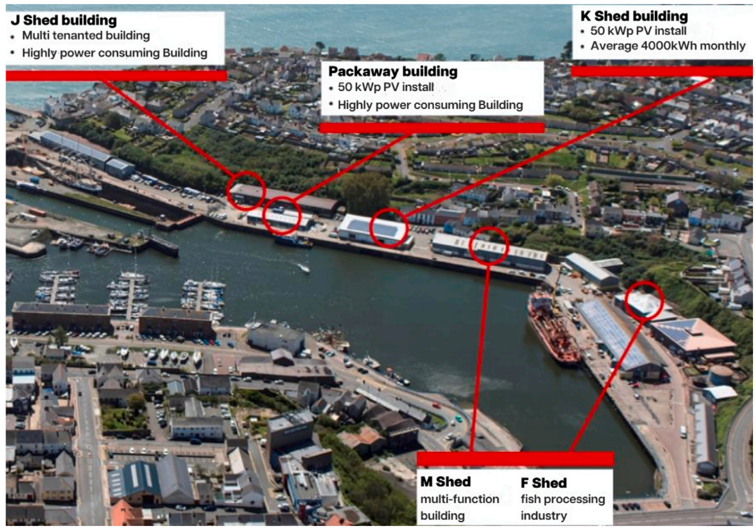

Fig. 2. Major industrial buildings and their energy attributions located in Milford Haven port, Wales.

operational goals, including: minimising electricity cost, reducing carbon emissions, and enhancing self-consumption of renewables. Each variable is linked to physical constraints and cost structures, forming the basis for multi-objective optimisation.

P2P energy trading adopts a fixed mid-tariff benchmark, based on average industrial rates at the pilot site (approximately £0.14/kWh during the study period). This simplification reduces model complexity while enabling economically balanced transactions. However, it limits its responsiveness to dynamic market signals such as time-of-use or marginal pricing. Future enhancements could adopt decentralised price finding, such as agent-based bidding, blockchain smart contracts, to reflect temporal energy value (Tushar et al., 2020).

The optimisation problem is formulated to minimise the net daily energy cost across a given scheduling horizon. The objective function is defined as:

$$f(x) = \sum_{i=1}^{L} [x_{L+i} - 0.9 \cdot x_{2L+i}] \cdot \text{dailyPrice}_i \tag{1}$$

Where:

- x is a decision variable vector of length 3L, representing: x_i: daily energy demand, x_{L+i}: energy purchased from the grid, and x_{2L+i}: energy sold to the grid.
- dailyPrice_i represents the unit electricity price (in £/kWh) on day i.
- The coefficient 0.9 reflects a standard feed-in-tariff assumption, where buildings sell electricity back to the grid at 90% of the purchase price.

The model is subject to four critical constraint sets:

- **Energy balance** ensures that total demand is met by the sum of local generation, battery discharge, and grid imports:

$$A_{\text{eq}} \cdot x = b_{\text{eq}} \tag{2}$$

where A_{eq} maps energy flows, and b_{eq} denotes daily forecasted demand.

- **Energy flow limitations** prevent total energy flows (e.g., battery or grid imports) from exceeding system capacity:

$$A \cdot x \leq b \tag{3}$$

- **Non-negativity** enforces all energy variables cannot be non-negative:

$$x > 0 \tag{4}$$

- **Contradiction energy** prevents economically illogical behaviour (e.g., buying and selling to the grid simultaneously).

$$c_i(x) = x_{L+i} \cdot x_{2L+i} = 0 \quad \text{for } i = 1, \dots, L \tag{5}$$

The framework is implemented in MATLAB, using `fmincon` function for constrained nonlinear multivariable optimisation, chosen for its robust object-oriented programming capabilities and extensive library of built-in mathematical tools.

3.3. Scenario development

The analysis incorporates temporal segmentation which divides the day into four distinct periods: early morning (00:00-06:00), daytime (06:00-12:00), afternoon (12:00-18:00), and evening (18:00-00:00), allowing to accurately capture the typical daily routine and activities of a fish processing facility. The analysis focuses on three distinct scenarios: (i) long-term optimisation of individual buildings, (ii) cluster optimisation, and (iii) P2P optimisation. In addition, each scenario is performed under normal conditions, high operational conditions (doubled demand), and high penetration of renewable energy (doubled production).

3.3.1. Scenario 1: Long-term optimisation of the individual building

The individual building optimisation algorithm is designed to allow each building to independently manage its energy use, storage, and scheduling, as described in Algorithm 1. It operates in a daily loop across the simulation horizon, using local forecasts and tariff data to minimise cost and grid dependency.

Each day, the algorithm:

1. Forecasts building energy demand based on historical usage and operating schedules.
2. Estimates available solar generation using irradiance profiles and system capacity.
3. Retrieves the applicable electricity tariff, which may vary by time of day.

4. Schedules appliances according to flexibility, priority, and operational constraints.
5. Optimises battery charging and discharging, ensuring physical limits and satae of charge (SoC) bounds are respected.
6. Solves a cost-minimisation problem to balance solar, grid, and battery energy sources.
7. Logs energy flows and cost indicators to evaluate performance over time.

This decentralised model allows each agent to act autonomously, but its performance may be suboptimal when excess solar energy cannot be shared, especially during high demand.

Algorithm 1 Individual Building Optimisation

```
procedure OptimiseBuilding(building, timeHorizon)
    for each day in timeHorizon do
        predictDemand(building, day)
        estimateSolarProduction(building, day)
        optimiseEnergySchedule(building, day)
        for each appliance in building do
            scheduleApplianceUsage(appliance, day)
        end for
        updateBatteryState(building)
        logDailyData(building)
    end for
end procedure
```

3.3.2. Scenario 2: Cluster optimisation

The cluster optimisation strategy treats the set of industrial buildings as a single, integrated energy system with centralised coordination, as coded in Algorithm 2. This allows for smarter resource distribution across the site. The algorithm follows this structure:

1. Aggregate forecast data: Total site demand and total site solar production are computed.
2. Determine global battery status: The shared battery's charge/discharge status is updated centrally.
3. Optimise the site-wide energy schedule, deciding how much energy to draw from the grid, store, use, or export.
4. Distribute optimised energy allocations to each building based on their needs and constraints.
5. Schedule appliances and local operations while respecting the centralised plan.
6. Calculate costs, emissions, and grid interaction metrics.

This strategy leverages economies of scale but assumes full coordination and communication, which may be challenging in practice.

Algorithm 2 Cluster Optimisation

```
procedure OptimiseCluster(buildings, timeHorizon)
    initialiseGlobalBuilding(buildings)
    for each day in timeHorizon do
        aggregateDemand(buildings)
        aggregateSolarProduction(buildings)
        optimiseClusterEnergySchedule(day)
        for each building in buildings do
            allocateEnergy(building, day)
            for each appliance in building do
                scheduleApplianceUsage(appliance, day)
            end for
        end for
        updateGlobalBatteryState()
        distributeCostsAndRevenues(buildings)
        logClusterData()
    end for
end procedure
```

3.3.3. Scenario 3: P2P optimisation

In this scenario, we focus on decentralised energy trading within the energy cluster, enabling buildings to directly exchange energy. The P2P trading rates are set using a static mid-range tariff value, derived from averaged industrial electricity pricing data specific to the Welsh context. This simplification allows the model to emphasise agent-level interaction dynamics while maintaining economic feasibility, as described in Algorithm 3. The P2P optimisation algorithm enables decentralised trading among buildings, where each agent decides when and how much energy to trade with others based on local conditions.

The steps are:

1. Forecast individual building demand and solar output.
2. Optimise each building's local schedule and battery usage.
3. Identify energy surplus and deficit buildings during each time slot.
4. Evaluate possible trades between surplus and deficit agents using a static mid-tariff baseline.
5. Execute energy trades if both parties benefit and physical constraints allow it.
6. Update battery status and grid interactions after trading.
7. Record KPIs for each agent, including cost, emissions, and grid reliance.
8. Record KPIs for each agent, including cost, emissions, and grid reliance.

This method adds flexibility and autonomy but introduces coordination complexity and potential inefficiencies when surplus energy is not matched effectively.

Algorithm 3 Peer-to-Peer Optimisation

```
procedure OptimisePeerToPeer(buildings, timeHorizon)
    for each day in timeHorizon do
        for each building in buildings do
            predictDemand(building)
            estimateSolarProduction(building)
            optimiseIndividualSchedule(building)
        end for
        for each timeslot in day do
            for each building A in buildings do
                for each building B in buildings where B ≠ A do
                    if canTrade(A, B, timeslot) then
                        price = calcTradePrice(A, B, timeslot)
                        execEnergyTrade(A, B, price, timeslot)
                    end if
                end for
            end for
        end for
        for each building in buildings do
            finaliseGridTransactions(building)
            updateBatteryState(building)
            calculateEmissions(building)
            logBuildingData(building)
        end for
        calculateClusterEmissions()
        logClusterData()
    end for
end procedure
```

4. Results and discussion

This section presents a comprehensive evaluation of the three proposed optimisation strategies — individual, cluster, and peer-to-peer (P2P) — under varying operational conditions. Simulations were carried out over a four-week horizon at the Milford Haven pilot site. Each strategy was assessed under three distinct operational states: normal, high demand, and high renewable generation. The performance

Table 1
Monthly electricity costs, carbon emissions, and grid dependence for the pilot seaport in Wales under baseline and three optimised scenarios, with corresponding potential savings across different operating conditions.

Use-case	KPI	Scenario				Savings [%]		
		Baseline	SC1	SC2	SC3	SC1	SC2	SC3
Normal	EC [£]	3785.7 (±75)	3043.7 (±64)	1559.7 (±32)	1813.3 (±37)	19.6	58.8	52.1
	CO_2 [ton]	3.33 (±0.06)	2.84 (±0.05)	2.37 (±0.04)	2.17 (±0.03)	14.7	28.7	34.7
	GD [MWh]	14.7 (±0.3)	11.1 (±0.2)	5.8 (±0.1)	6.9 (±0.1)	24.6	60.2	53.3
High Demand	EC [£]	7481.5 (±130)	6881.5 (±120)	5981.5 (±105)	6481.5 (±112)	8.3	18.1	13.4
	CO_2 [ton]	6.66 (±0.12)	6.11 (±0.10)	5.31 (±0.09)	5.76 (±0.10)	8.4	18.2	13.5
	GD [MWh]	29.5 (±0.6)	27.1 (±0.5)	23.5 (±0.4)	25.5 (±0.5)	8.4	18.2	13.5
High Production	EC [£]	−2163.7 (±50)	−3118.0 (±61)	−3674.0 (±70)	−3351.0 (±66)	−44.1	−69.8	−54.9
	CO_2 [ton]	0.97 (±0.02)	0.15 (±0.01)	0.12 (±0.01)	0.11 (±0.01)	84.6	87.3	89.2
	GD [MWh]	−9.1 (±0.2)	−12.8 (±0.3)	−15.1 (±0.4)	−13.7 (±0.3)	−42.3	−67.1	−52.7

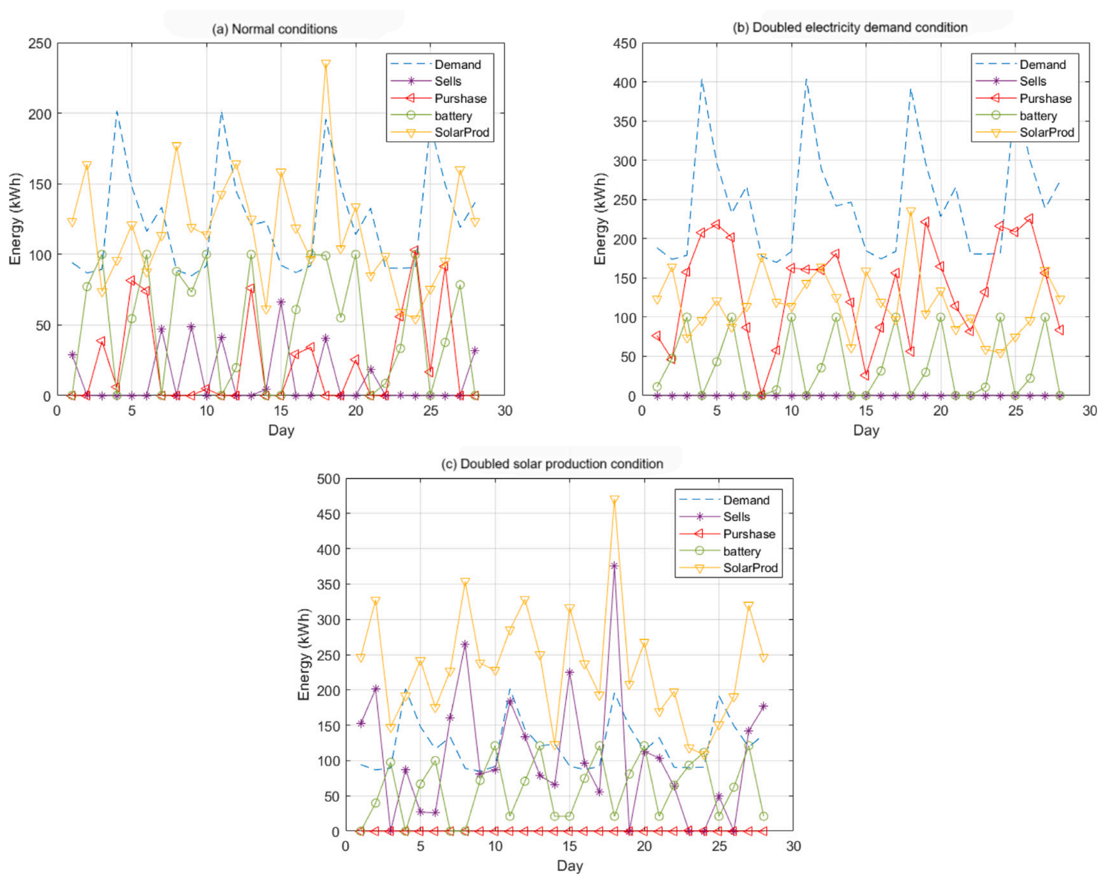

Fig. 3. Daily electricity profiles for individual building optimisation under normal, high-demand, and high-production scenarios, showing demand, grid interactions (sell and purchase), battery usage, and solar generation.

was measured across three critical KPIs: electricity cost (EC), carbon emissions (CO_2), and grid dependency (GD). Table 1 summarises the monthly results, including error deviations to reflect system sensitivity to input variability (e.g., weather fluctuations and demand profiles). These statistical measures enhance the reliability of the analysis and allow a more rigorous comparison across strategies.

4.1. Scenario 1: Individual building optimisation

The individual optimisation strategy enables each building to independently schedule its energy usage, without inter-building coordination or shared energy resources. As shown in Fig. 3, daily energy profiles vary considerably depending on the operational condition.

Under normal conditions, the buildings demonstrate consistent daily consumption and modest solar generation. The system achieves a 19.6% reduction in energy costs and a 14.7% cut in emissions, with standard deviations of ±2.0% and ±1.5% respectively. Approximately 85% of solar generation is self-consumed, with the remainder unused due to the absence of sharing mechanisms. Grid reliance remains above 40%, revealing the inherent limitations of isolated decision-making.

Under high demand, system rigidity becomes more noticeable. Despite intensified battery cycling and optimised scheduling, the cost and emission reductions decrease to around 8.3% and 8.4% respectively. The low variability (±1.1%) indicates limited adaptivity, as buildings reach storage and flexibility limits. The battery rapidly saturates, and excess demand must be met through grid imports. Without coordinated

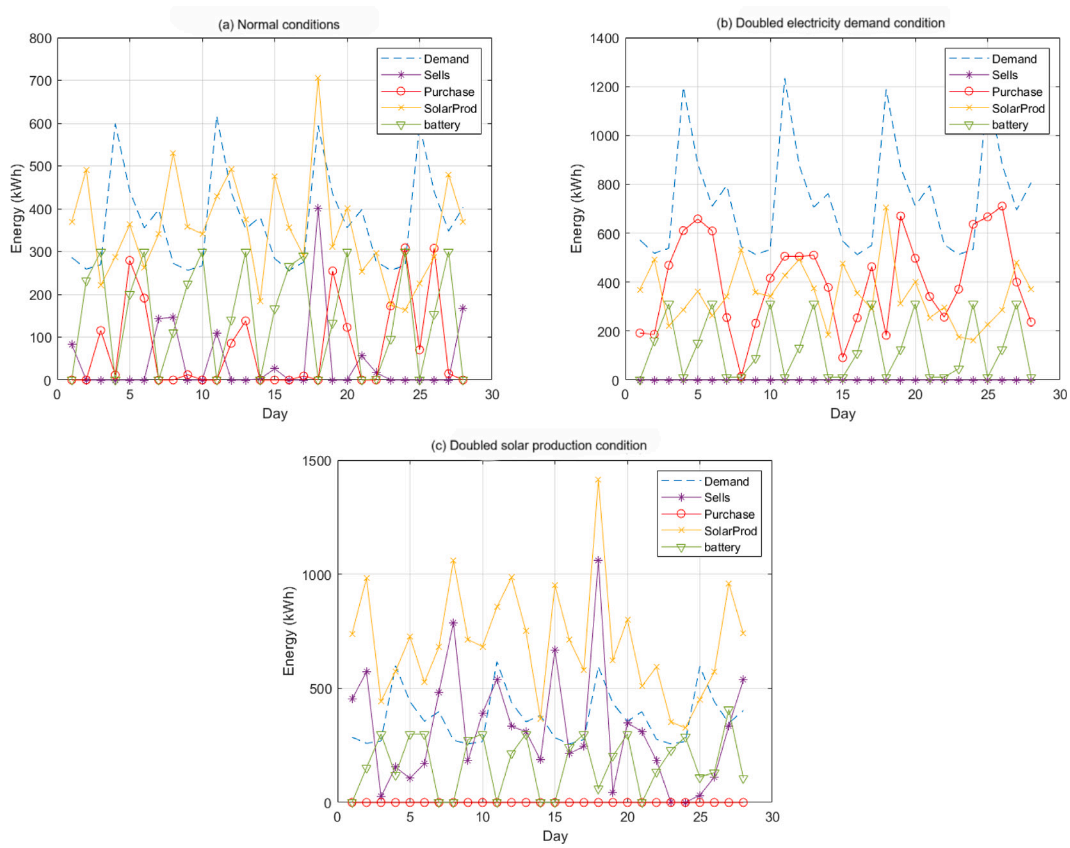

Fig. 4. Cluster-level electricity profiles under normal, high-demand, and high-production scenarios, showing aggregated demand, grid exchanges (sell and purchase), battery operation, and solar generation.

Table 2
Summary comparison of the three strategies across operational scenarios.

Metric	Individual	Cluster	P2P
Cost saving (%)	16.0	**45.0**	39.0
Emissions reduction (%)	35.9	48.1	**45.8**
Grid import reduction (%)	25.1	**62.5**	53.2
Coordination complexity	Low	High	Medium
Sensitivity to demand spikes	High	Low	Medium

load balancing or shared energy assets, the strategy is constrained in peak demand scenarios.

By contrast, high renewable availability yields net financial gains, with buildings becoming net exporters. The system achieves an 84.6% emissions reduction and an effective reversal in grid dependency (–9.1 MWh), facilitated by high battery utilisation and surplus generation. Nevertheless, some solar energy remains underutilised due to the lack of demand-matching. This underscores the limitation of independent optimisation: while effective under abundant supply, it underperforms in collaborative potential and dynamic system-wide efficiency.

4.2. Scenario 2: Cluster optimisation

The cluster optimisation strategy treats the entire facility as a coordinated energy system, as illustrated in Fig. 4. Under normal conditions, this approach achieves a 58.8% cost saving and 28.7% emissions reduction—around triple the performance of the individual strategy. The standard deviations (±2.5% cost, ±1.8% emissions) indicate robust system stability. This efficient performance is mainly driven by the pooled battery system, which redistributes energy across buildings and eliminates solar curtailment.

Even under high demand, the system maintains 18.1% cost and 18.2% emissions reductions, outperforming other strategies in this stress scenario. However, performance plateaus as shared battery saturation limits further optimisation. These diminishing returns reveal that while central coordination is more resilient, it cannot fully offset physical system constraints such as fixed storage capacity.

In high-production scenarios, cluster optimisation produces net profits and 87.3% emissions reduction. Solar utilisation exceeds 95%, enabled by optimised distribution and load shifting. However, while statistically consistent, these predicted results are based on perfect communication and ideal storage efficiency assumptions. Real-world constraints — latency, interoperability, fault tolerance — may reduce practical gains.

4.3. Scenario 3: P2P optimisation

The P2P strategy introduces a decentralised trading mechanism where buildings exchange surplus energy directly, as shown in Fig. 5. Under normal conditions, it yields a 52.1% cost saving and 34.7% emissions reduction. Although slightly less efficient than the cluster approach, its decentralised nature offers greater autonomy and flexibility. However, wider standard deviations (±3.1% cost, ±2.0% emissions) indicate more sensitivity to matching delays and agent behaviour.

Under high demand, P2P outperforms the individual strategy but trails behind the cluster approach, with 13.4% cost and 13.5% emissions reduction. The battery system is heavily utilised, but occasional mismatches in energy trading lead to underutilised renewables. Static mid-tariff pricing, used in the current model, fails to incentivise dynamic trade responses and may constrain future scalability. Incorporating real-time pricing or auction models could mitigate this limitation.

In high-production conditions, average building profits reach £670/month, and carbon emissions decrease by 89.2%. However,

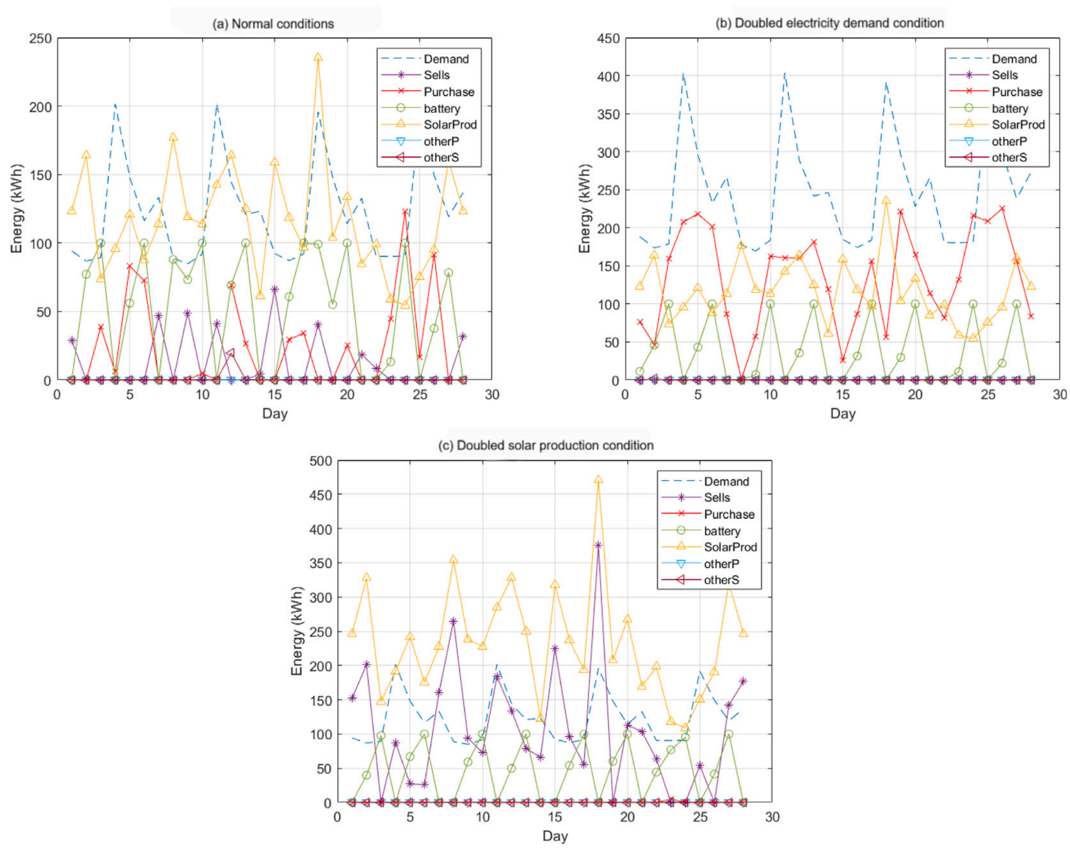

Fig. 5. P2P optimisation electricity profiles under normal, high-demand, and high-production scenarios, showing demand, grid exchanges, battery use, solar output, and peer-to-peer transactions.

this performance highly depends on sufficient surplus–deficit alignment. Variability in agent participation or forecasting errors could reduce real-world efficiency. Moreover, transaction costs, regulatory tariffs, and infrastructure interoperability may affect implementation feasibility.

4.4. Comparative insights

Table 2 summarises the trade-offs between all strategies. While cluster optimisation provides the highest average performance and stability, it requires robust central coordination and full interoperability. It has limited responsiveness to individual demand spikes. P2P trading balances efficiency and autonomy but is exposed to negotiation inefficiencies. Independent building optimisation, while simplest to implement, underperforms in all scenarios except those with significant renewable surplus.

These findings suggest a hybrid framework — combining cluster-wide control for baseline optimisation and P2P trading for local balancing — may yield the most resilient and scalable solution. Moreover, integrating real-time pricing and stakeholder engagement mechanisms (e.g., fair cost allocation, demand response incentives) can further enhance adaptability and fairness in future implementations.

5. Conclusion

This study has demonstrated the effectiveness of Multi-Agent Systems (MAS) in addressing the complex energy management challenges of industrial seaport facilities, where energy demand is highly variable, infrastructure constraints are pronounced, and renewable integration is often limited. The proposed MAS-based optimisation framework — applied and validated using real-world data from Milford Haven port in the UK — delivers significant improvements in operational efficiency, environmental performance, and energy autonomy.

Through the simulation of three optimisation strategies — individual building control, centralised cluster coordination, and decentralised peer-to-peer (P2P) energy trading — the study provides a comprehensive evaluation of trade-offs across cost savings, carbon emissions, and grid dependency. Key findings include:

- **Individual building optimisation** yielded moderate gains, achieving 19.6% (±2.0%) cost reduction and 14.7% (±1.5%) emissions reduction. While suitable for isolated building management, this strategy struggles under high-demand scenarios due to limited flexibility and the absence of shared resources.
- **Cluster-level optimisation** demonstrated the highest performance across all KPIs, with average savings of 58.8% (±2.5%) in electricity costs and 30% (±1.8%) in carbon emissions. Coordinated energy scheduling and shared battery systems enabled near-optimal utilisation of renewables and substantial grid independence.
- **Peer-to-peer optimisation** provided a decentralised alternative with promising outcomes — achieving 52.1% (±3.1%) cost reduction and 34.7% (±2.0%) emissions reduction — primarily through direct energy exchange between buildings. Although performance was slightly lower than cluster control, P2P offered enhanced flexibility and autonomy, making it suitable for contexts with heterogeneous ownership or decentralised governance structures.

These findings underscore the scalability, resilience, and adaptability of MAS-based control frameworks for energy-intensive sectors. While the pilot focused on a seaport context, the proposed methodology is readily transferable to other industrial domains such as food processing, aquaculture, and manufacturing. It provides a promising pathway towards cost-effective, decentralised, and low-carbon energy systems.

Nonetheless, several limitations necessitate critical consideration:

- Limited validation scope: The current framework was evaluated using data from a single industrial site, which may restrict generalisability. Performance under different regulatory regimes, weather profiles, or industrial sectors remains to be tested.
- Restricted renewable portfolio: The study focused exclusively on solar PV systems. Future research should integrate other renewable energy sources — including wind, tidal, and biomass — as well as long-duration storage technologies such as hydrogen or thermal energy storage.
- Static optimisation parameters: PV size, battery capacity, and temporal segmentation were pre-defined. Dynamic reconfiguration of these parameters in response to real-time conditions could unlock greater system flexibility.
- Simplified P2P market structure: The current P2P trading model does not account for transaction expenses, market liquidity limits, verification latency, or regulatory conditions, such as grid codes, feed-in tariffs, and local trading permits, which can significantly impact trading viability and outweigh economic gains in small-scale P2P setups. Additionally, the absence of mechanisms for dispute resolution and billing validation poses legal and operational risks.

In light of these limitations, several future research directions are proposed:

- Cross-sector scalability assessment: Apply the MAS framework to diverse industrial environments to evaluate its generalisability, adaptability, and cost–benefit potential across sectors and regions.
- Integration of multiple RES and hybrid storage: Extend the model to include additional renewable sources and complementary storage systems. Hybrid configurations may improve energy resilience and enable higher penetration of renewables in constrained environments.
- P2P market design and stakeholder participation: Future work should incorporate dynamic pricing, transaction validation, and regulatory mechanisms. Involving industrial stakeholders in the design-phase participatory modelling would ensure alignment with real-world economic and institutional constraints.
- Adaptive control and learning-based optimisation: Introduce adaptive algorithms — such as reinforcement learning or evolutionary optimisation — to enable real-time response to fluctuating supply, demand, and market signals.
- Lifecycle impact assessment: Incorporate life cycle analysis (LCA) to assess long-term environmental and economic sustainability. While short-term gains are evident, the embodied carbon and financial cost of PV panels and battery systems must be accounted for in a full-system decarbonisation assessment.

In conclusion, this study advances the application of intelligent energy systems in industrial clusters and provides a robust foundation for future work on distributed energy management, collaborative optimisation, and industrial decarbonisation. The MAS framework offers not only technical efficiency but also a governance model for future-ready, low-carbon industrial ecosystems.

CRediT authorship contribution statement

Amin Amin: Writing – review & editing, Visualization, Validation, Supervision. **Ali Ghoroghi:** Writing – review & editing, Validation, Supervision. **Ioan Petri:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Funding acquisition, Conceptualization. **Andrei Hodorog:** Conceptualization, Data curation, Investigation, Software, Visualization, Writing – original draft. **Thomas Genest:** Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Data curation, Conceptualization. **Yacine Rezgui:** Writing – review & editing, Supervision, Resources.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The work in this study was carried out under the IAA-EPSRC grant: "Decarbonising seaports using digital twins: A CUSP platform feasibility study". This project also uses data and information from the EU Interreg piSCES project.

References

Adewoyin, Musa Adekunle, Adediwin, Olugbenga, Audu, Audu Joseph, 2025. Artificial intelligence and sustainable energy development: A review of applications, challenges, and future directions. Int. J. Multidiscip. Res. Growth Eval. 6 (2), 196–203.

Alqarni, Z., Rezgui, Y., Petri, I., Ghoroghi, A., 2023. Factors and strategies affecting indoor air quality in educational buildings. In: 2023 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC). Edinburgh, UK, http://dx.doi.org/10.1109/ICE/ITMC58018.2023.10332421.

AlSkaif, Tarek, Crespo-Vazquez, Jose L, Sekuloski, Milos, Van Leeuwen, Gijs, Catalão, João PS, 2021. Blockchain-based fully peer-to-peer energy trading strategies for residential energy systems. IEEE Trans. Ind. Informatics 18 (1), 231–241.

Alzahrani, A., Petri, I., Rezgui, Y., 2020. Modelling and implementing smart microgrids for the fish-processing industry. In: 2019 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC). IEEE, pp. 1–8. http://dx.doi.org/10.1109/ICE.2019.8792575.

Alzahrani, A., Petri, I., Rezgui, Y., 2021. Optimal control-based price strategies for smart fishery ports micro-grids. In: 27th International Conference on Engineering, Technology and Innovation (ICE/ITMC). Cardiff, UK, http://dx.doi.org/10.1109/ICE/ITMC52061.2021.9570267.

Amin, Amin, Kem, Oudom, Gallegos, Pablo, Chervet, Philipp, Ksontini, Feirouz, Mourshed, Monjur, 2022. Demand response in buildings: Unlocking energy flexibility through district-level electro-thermal simulation. Appl. Energy (ISSN: 0306-2619) 305, 117836. http://dx.doi.org/10.1016/j.apenergy.2021.117836, URL https://www.sciencedirect.com/science/article/pii/S0306261921011612.

Amin, Amin, Petri, Ioan, Minett, Simon, 2026. Local energy community with a power-to-gas system: a feasibility case study and performance analysis. Renewable and Sustainable Energy Reviews (ISSN: 1364-0321) 225, 116147. http://dx.doi.org/10.1016/j.rser.2025.116147, URL https://www.sciencedirect.com/science/article/pii/S1364032125008202.

Arena, F., Barbaro, G., Romolo, A., Malara, G., Fiamma, V., 2018. From green energy to green logistics: A pilot study in an Italian port area. Transp. Res. Procedia 30, 111–118. http://dx.doi.org/10.1016/j.trpro.2018.09.013.

Bregar, Andrej, 2020. Implementation of a multi-agent multi-criteria negotiation protocol for self-sustainable smart grids. J. Decis. Syst. 29 (sup1), 87–97. http://dx.doi.org/10.1080/12460125.2020.1733701.

Buiza, G, Cepolina, S, Dobrijevic, A, del Mar Cerbán, M, Djordjevic, O, González, C, 2015. Current situation of the mediterranean container ports regarding the operational, energy and environment areas. In: International Conference on Industrial Engineering and Systems Management. IESM, IEEE, pp. 530–536. http://dx.doi.org/10.1109/IESM.2015.7380209.

Buonomano, Annamaria, Del Papa, Gianluca, Giuzio, Giovanni Francesco, Palombo, Adolfo, Russo, Giuseppe, 2023. Future pathways for decarbonization and energy efficiency of ports: Modelling and optimization as sustainable energy hubs. J. Clean. Prod. (ISSN: 0959-6526) 420, 138389. http://dx.doi.org/10.1016/j.jclepro.2023.138389, URL https://www.sciencedirect.com/science/article/pii/S0959652623025477.

Celik, Y., Petri, I., Rezgui, Y., 2021. Leveraging BIM and blockchain for digital twins. In: 27th ICE/IEEE International Technology Management and Innovation Conference. Cardiff, UK, http://dx.doi.org/10.1109/ICE/ITMC52061.2021.9570246.

Domyshev, Alexander, Sidorov, Denis, Liu, Fang, 2024. Microgrids multi-agent control based on power distribution and communication constraints. In: 2024 43rd Chinese Control Conference. CCC, IEEE, pp. 5812–5816.

Du, Yuquan, Li, Chen, Wang, Tingsong, Xu, Ya, 2023. Special issue on "smart port and shipping operations" in maritime policy & management. Marit. Policy & Manag. 50 (4), 413–414. http://dx.doi.org/10.1080/03088839.2023.2196754.

Food and Agriculture Organization of the United Nations, 2020. The State of World Fisheries and Aquaculture 2020. FAO, URL https://www.fao.org/interactive/state-of-fisheries-aquaculture/2020/en/.

Huang, Pei, Lovati, Marco, Zhang, Xingxing, Bales, Chris, 2020. A coordinated control to improve performance for a building cluster with energy storage, electric vehicles, and energy sharing considered. Appl. Energy 268, 114983.

Iris, Çağatay, Lam, Jasmine Siu Lee, 2021. Optimal energy management and operations planning in seaports with smart grid while harnessing renewable energy under uncertainty. Omega (ISSN: 0305-0483) 103, 102445. http://dx.doi.org/10.1016/j.omega.2021.102445, URL https://www.sciencedirect.com/science/article/pii/S0305048321000542.

Kaspar, Kathryn, Ouf, Mohamed, Eicker, Ursula, 2022. A critical review of control schemes for demand-side energy management of building clusters. Energy Build. 257, 111731.

Liu, Hualong, Tang, Wenyuan, 2024. Optimal operation of multi-energy microgrids: A hybrid methodology using stochastic programming and generalized Nash bargaining. In: 2024 IEEE Energy Conversion Congress and Exposition. ECCE, IEEE, pp. 1045–1052.

May, Ross, Huang, Pei, 2023. A multi-agent reinforcement learning approach for investigating and optimising peer-to-peer prosumer energy markets. Appl. Energy (ISSN: 0306-2619) 334, 120705. http://dx.doi.org/10.1016/j.apenergy.2023.120705, URL https://www.sciencedirect.com/science/article/pii/S0306261923000697.

Petri, Ioan, Amin, Amin, Ghoroghi, Ali, Hodorog, Andrei, Rezgui, Yacine, 2025. Digital twins for dynamic life cycle assessment in the built environment. Science of The Total Environment (ISSN: 0048-9697) 993, 179930. http://dx.doi.org/10.1016/j.scitotenv.2025.179930, URL https://www.sciencedirect.com/science/article/pii/S0048969725015700.

Petri, Ioan, Yama, Aida, Rezgui, Yacine, 2019. Agent-based appliance scheduling for energy management in industry 4.0. In: GECON 2019 - 16th International Conference on the Economics of Grids, Clouds, Systems, and Services. Springer, pp. 199–207. http://dx.doi.org/10.1007/978-3-030-36027-6_17.

Ramos, V, Carballo, R, Álvarez, M, Sánchez, M, Iglesias, G, 2014. A port towards energy self-sufficiency using tidal stream power. Energy 71, 432–444. http://dx.doi.org/10.1016/j.energy.2014.04.098.

Shen, Rendong, Zhong, Shengyuan, Wen, Xin, An, Qingsong, Zheng, Ruifan, Li, Yang, Zhao, Jun, 2022. Multi-agent deep reinforcement learning optimization framework for building energy system with renewable energy. Appl. Energy (ISSN: 0306-2619) 312, 118724. http://dx.doi.org/10.1016/j.apenergy.2022.118724, URL https://www.sciencedirect.com/science/article/pii/S0306261922001829.

Tushar, Wayes, Saha, Tapan Kumar, Yuen, Chau, Smith, David, Poor, H Vincent, 2020. Peer-to-peer trading in electricity networks: An overview. IEEE Trans. Smart Grid 11 (4), 3185–3200.

Wang, Guofeng, Liu, Yongqi, Zhang, Youbing, Yan, Jun, Xie, Shuzong, 2024. Distributed online optimization for integrated energy systems: A multi-agent system consensus approach. Internat. J. Adapt. Control Signal Process. 38 (10), 3401–3421.

Yang, Ting, Xu, Zheming, Ji, Shijie, Liu, Guoliang, Li, Xinhong, Kong, Haibo, 2025. Cooperative optimal dispatch of multi-microgrids for low carbon economy based on personalized federated reinforcement learning. Appl. Energy 378, 124641.

Yang, Y, Zhong, M, Yao, H, Yu, F, Fu, X, Postolache, O, 2018. Internet of things for smart ports: Technologies and challenges. IEEE Instrum. Meas. Mag. 21 (1), 34–43. http://dx.doi.org/10.1109/MIM.2018.8278808.

Zhu, Dafeng, Yang, Bo, Liu, Yuxiang, Wang, Zhaojian, Ma, Kai, Guan, Xinping, 2022. Energy management based on multi-agent deep reinforcement learning for a multi-energy industrial park. Appl. Energy (ISSN: 0306-2619) 311, 118636. http://dx.doi.org/10.1016/j.apenergy.2022.118636, URL https://www.sciencedirect.com/science/article/pii/S0306261922001064.