SWISSI INSTITUTE FOR AI

Multi-Agent AI for ESG-Tracked Energy Production and Trading on a Decentralised DAG-Based Ledger

Walter Kurz[1], Wojtek Stricker[1]

December 2025

Abstract

This paper proposes an architectural blueprint for a multi-agent AI system that supports ESG-tracked energy trading on a decentralised DAG-based ledger. The goal is to integrate verifiable environmental, social and governance attributes into the core logic of energy transactions, rather than treating ESG as an external reporting layer. The envisioned system links Industrial Internet of Things (IIoT) devices of energy producers and grid operators to smart contracts that are stored in a DAG structure with fine-grained timestamping and geolocation. Specialised AI agents are able to read, write and execute these smart contracts, enriching them with real time production data, grid conditions and ESG metrics associated with different energy types. On top of this contract and data layer, coordinating agents manage energy trading and allocation. Game theoretic mechanisms based on Nash equilibrium guide bidding, pricing and matching routines under both economic and ESG constraints. Selected agents can act as automated market makers that provide liquidity and price discovery for ESG-differentiated energy products.

The paper focuses on the entire chain from IIoT-enabled energy producers, through the DAG-based contract and data layer, up to AI-driven trading and allocation. Its main contribution is an AI-centric reference architecture that embeds verifiable ESG tracking and automated compliance into this end-to-end flow while enabling efficient coordination of energy flows in industrial energy systems.

Keywords: artificial intelligence, multi-agent systems, energy trading, ESG tracking, DAG-based ledger, smart contracts, Industrial IoT, automated market making, Nash equilibrium

1 Introduction

Decarbonisation targets, volatile energy prices and regulatory pressure on environmental, social and governance (ESG) reporting increase the demand for verifiable, machine-readable information on how energy is produced, transported and traded [1, 2]. Existing ESG disclosures often rely on heterogeneous data sources and manual aggregation, which leads to fragmented and weakly auditable information [1, 3]. Industrial energy systems already employ extensive IoT and IIoT infrastructures for monitoring and control, yet these data streams are only partially integrated with ESG reporting and trading processes [4, 5]. Blockchain platforms and smart contracts are increasingly investigated for energy applications, including peer-to-peer trading, certificate management and automated settlement [6, 7]. ESG-focused architectures combine these technologies with sensor data to support tamper-resistant, token-incentivised reporting and traceability [3, 1]. Most existing work, however, treats data acquisition, market interaction and decision support as largely separate layers, rather than analysing how artificial intelligence can coordinate operational and trading decisions across the full energy value chain. The growing penetration of renewable resources such as wind, solar photovoltaics and battery storage intensifies variability and forecast uncertainty, which complicates operational planning and market participation [8]. Electricity is traded across several short-term markets with distinct time horizons and rules, typically including day-ahead, intraday and imbalance stages that jointly determine the value of flexibility and forecasting quality [9, 10, 11].

In the Day-Ahead Market (DAM), producers submit hourly bids for the following day based on expected generation and demand, and prices result from market-clearing mechanisms that match aggregate supply and demand [9, 10]. A simplified expression for the clearing price in hour t is

$$\text{Price}_{DA,t} = \frac{\sum_i \text{Bid}_i(t) \cdot \text{Quantity}_i(t)}{\sum_i \text{Quantity}_i(t)}, \tag{1}$$

where i indexes market participants, $\text{Bid}_i(t)$ is the price offer of participant i, and $\text{Quantity}_i(t)$ the corresponding quantity. Forecast errors in generation or consumption directly affect bidding decisions and thus financial outcomes. In the Intraday Market (IDM), participants adjust their positions close to real time in response to deviations from Day-Ahead schedules, outages or stochastic renewable output [11, 8]. The imbalance in hour t can be written as

$$\text{Imbalance}_t = \text{ActualGeneration}_t - \text{ScheduledGeneration}_t, \tag{2}$$

where $\text{ActualGeneration}_t$ denotes measured production and $\text{ScheduledGeneration}_t$ the quantity committed in the DAM. Positive imbalances represent surplus energy that can be sold intraday, while negative imbalances require additional procurement at potentially unfavourable prices, which increases the economic value of flexible assets such as batteries and controllable loads [8].

For battery storage systems, revenue generation often relies on arbitrage opportunities between market segments, where storage operators buy energy during low-price periods and sell during high-price periods [12, 13]. The profit from charging and discharging a battery during a given period t can be written as

$$\text{Profit}_t = (\text{Price}_{sell,t} - \text{Price}_{buy,t}) \cdot \text{EnergyTransferred}_t, \tag{3}$$

where $\text{Price}_{sell,t}$ and $\text{Price}_{buy,t}$ denote the selling and buying prices at time t, and $\text{EnergyTransferred}_t$ is the amount of energy moved through the storage system.

Beyond economic objectives, producers are increasingly required to comply with environmental, social, and governance (ESG) criteria, frequently operationalised as composite scores with indicator-specific weights [2, 1]. An overall ESG score can be defined as

$$\text{ESG\_Score} = w_E E + w_S S + w_G G, \tag{4}$$

where E, S, and G represent environmental, social, and governance indicators, and w_E, w_S, and w_G are their respective weights. In digital trading architectures, smart contracts and AI agents can implement minimum thresholds $E_{min}, S_{min}, G_{min}$ as hard constraints on operational and portfolio decisions [3, 1]:

$$E \geq E_{min}, \quad S \geq S_{min}, \quad G \geq G_{min}. \tag{5}$$

Many renewable energy producers still access short-term electricity markets through third-party marketers, which introduces commission costs, information delays and portfolio-level trading strategies that often ignore asset-specific flexibility. ESG-related data, certificates and regulatory compliance steps are handled in parallel, largely manual workflows that remain decoupled from bidding and dispatch decisions. These structural frictions motivate decentralised, producer-centric architectures in which AI agents represent individual assets, interact directly with day-ahead and intraday markets, and execute trades through smart contracts on a DAG-based ledger that embeds ESG constraints into the transaction logic.

2 Research Gap

Existing work on digital energy trading tends to specialise in three directions. First, there is a large body of research on market design and pricing mechanisms for wholesale and joint energy–reserve or energy–ancillary service markets, often formulated as equilibrium or bilevel optimisation problems and used to study strategic bidding by generators, retailers and flexible resources [14, 15, 16, 17, 18]. Second, multi-agent control and optimisation methods have been applied to microgrids and integrated energy systems, where agents represent distributed generators, storage units and loads and coordinate real-time scheduling based on local measurements and IIoT data streams [19, 20, 21, 22]. Third, distributed ledgers and smart contracts have been explored as infrastructures for recording energy transactions, peer-to-peer trading and renewable energy certificates; recent surveys and systematisations analyse smart-contract design patterns, performance and open issues in energy-sector applications, including DAG-style ledgers and block-free architectures for high-frequency trading [23, 24, 25, 7, 26, 27]. ESG-related work builds mainly on this last strand: blockchain and smart contracts are proposed as tamper-resistant layers for ESG reporting, certification and audit trails, often in combination with IoT-based data collection, but typically without tight coupling to operational control of physical assets or market-clearing algorithms [28, 29, 30, 31, 32]. Across these literatures, multi-agent AI, IIoT data flows, smart contracts and distributed ledgers are usually treated as loosely coupled components rather than as a single, end-to-end architecture, and ESG objectives appear as reporting outputs instead of first-class constraints in agent policies and trading logic. Existing reviews highlight fragmentation and the predominance of proof-of-concept platforms rather than deployable reference architectures, which leaves a gap for designs that show how ESG-tracked energy, automated compliance checks and strategic interaction between agents can be combined in a technically implementable system, in particular when using DAG-style ledgers as the transaction backbone [23, 24, 25, 7, 30].

3 Contribution to the Field

This paper proposes an AI-centric reference architecture for ESG-tracked energy production and trading on a decentralised DAG-based ledger. The architecture spans IIoT-enabled producers and grid operators, which provide measurements and asset states, a DAG-based contract and data layer that stores timestamped and geolocated smart contracts, and AI agents that coordinate trading, allocation and liquidity provision. Within this pipeline, the paper specifies agent roles with explicit responsibilities and interaction patterns: agents that read, write and execute smart contracts; agents that aggregate and transform IIoT data into ESG-relevant indicators; and agents that participate in trading routines and pricing schemes subject to ESG and system-level constraints. Strategic interaction between these agents is organised through game-theoretic coordination mechanisms based on Nash-equilibrium concepts, so that economic incentives and ESG objectives can be encoded directly in agent policies. The contribution lies in making the links between multi-agent AI design, contract structures and ledger primitives explicit, and in providing an extensible architectural blueprint that can guide implementation and evaluation of ESG-tracked energy trading systems in industrial settings.

4 Research

Building on the gap identified above, we propose a two-level approach that distinguishes between the market game played by stakeholder-controlled participants and the multi-agent AI systems that implement their decision policies on a decentralised DAG-based ledger. At the market level, short-term electricity trading is represented as a non-cooperative game between producer portfolios, grid operators, aggregators, automated market makers and large consumers, with Nash equilibrium as the benchmark for strategic interaction. At the implementation level, we specify an AI-centric reference architecture in which IIoT-enabled assets feed data into a DAG-based smart-contract layer, and stakeholder policies are executed by agents under ESG and regulatory constraints. We instantiate this architecture in detail for a producer-side decision system, since producers combine direct access to physical assets with strong exposure to price volatility and ESG requirements. In this instantiation, a multi-agent AI system coordinates specialised agents for forecasting, bidding, storage control and ESG compliance on the basis of shared data from the DAG ledger and IIoT infrastructure. We also outline a complementary design for an automated market maker and order-matching service, which appears as a single participant in the market game while internally using multi-agent AI to manage liquidity provision, price updates and congestion-aware matching. The remainder of this section first formalises the market game and the associated ESG-aware objective and constraints, then presents the reference architecture, followed by the producer-side multi-agent system and the automated market maker design.

4.1 End-to-end physical and value flow

We consider a discrete time horizon $\mathcal{T}$ and a set of nodes $\mathcal{N}_{\text{phys}}$ in the physical system (for example grid zones or connection points). For each node $n \in \mathcal{N}_{\text{phys}}$ and time $t \in \mathcal{T}$, let $G_{n,t}$ denote the electrical energy generated by local assets, $L_{n,t}$ the electrical load, $B^{\text{ch}}_{n,t}$ and $B^{\text{dis}}_{n,t}$ storage charging and discharging, $C_{n,t}$ curtailed generation, $F_{n\to m,t}$ the energy flow from node n to node m, and $\ell_{n,t}$ network losses.

A simplified nodal balance for each (n, t) is

$$G_{n,t} - C_{n,t} + B^{\text{dis}}_{n,t} - B^{\text{ch}}_{n,t} + \sum_{m\neq n} F_{m\to n,t} - \sum_{m\neq n} F_{n\to m,t} = L_{n,t} + \ell_{n,t}. \tag{6}$$

Collecting all physical variables in a vector x gives the compact form

$$A_{\text{phys}} x = b_{\text{phys}}, \qquad 0 \le x \le \bar{x}, \tag{7}$$

where A_{phys} and b_{phys} encode network constraints and $\bar{x}$ denotes capacity limits.

On the market side, let $\mathcal{M}$ be the set of trading venues (for example day-ahead, intraday, balancing). For each market $m \in \mathcal{M}$, time $t \in \mathcal{T}$ and node or zone z, let $P_{m,z,t}$ be the market price and $Q_{p,m,z,t}$ the net quantity traded by participant p (positive for net sales, negative for net purchases). The financial payoff of participant p from energy trading is

$$\Pi^{\text{energy}}_p(s, x) = \sum_{m\in\mathcal{M}} \sum_{z} \sum_{t\in\mathcal{T}} P_{m,z,t}(s, x)\, Q_{p,m,z,t}(s, x), \tag{8}$$

where the dependence on (s, x) reflects that prices and cleared quantities result from the joint strategies s and the physical state x.

Let $\text{ESG}(x)$ denote a vector of ESG indicators computed from IIoT measurements and contract metadata on the DAG-based ledger, and let $\underline{e}$ be minimum ESG thresholds. The end-to-end chain is subject to

$$\text{ESG}(x) \ge \underline{e}, \tag{9}$$

which links physical operation, trading outcomes and ESG performance.

4.2 Stakeholders, participants and motives

Let $\mathcal{K}$ denote the set of stakeholder groups (for example producers, grid operators, aggregators, large consumers, regulators, ESG auditors and infrastructure providers). Each stakeholder group $k \in \mathcal{K}$ controls one or more market participants collected in a set $\mathcal{P}$. For each participant $p \in \mathcal{P}$, we define a payoff function of the form

$$u_p(s, x) = \Pi^{\text{energy}}_p(s, x) - \text{Cost}_p(s, x) - \Gamma_p\big(\text{ESG}(x)\big), \tag{10}$$

where Cost_p covers operational and contractual costs, and Γ_p encodes ESG-related penalties or incentives, for example through higher financing costs, fees or rewards linked to ESG performance.

At an aggregate level, a system-level objective can be written as a weighted combination of stakeholder utilities and ESG penalties

$$W(s,x) = \sum_{k\in\mathcal{K}} \omega_k \, U_k(s,x) - \lambda_E \, E(s,x) - \lambda_S \, S(s,x) - \lambda_G \, G(s,x), \tag{11}$$

where U_k is the utility of stakeholder group k, E, S and G measure deviations from target levels for environmental, social and governance metrics, and $\omega_k, \lambda_E, \lambda_S, \lambda_G \geq 0$ are weighting parameters.

4.3 Market game and Nash equilibrium

The strategic interaction between market participants is modelled as a non-cooperative game on top of the DAG-based ledger. Let S_p be the strategy set of participant $p \in \mathcal{P}$ and $S = \prod_{p\in\mathcal{P}} S_p$ the joint strategy space. A strategy profile is $s = (s_p)_{p\in\mathcal{P}} \in S$, with s_{-p} denoting the strategies of all participants except p. Each participant seeks to maximise its payoff $u_p(s,x)$ as defined in (10), subject to the physical and ESG constraints (7) and (9).

A strategy profile $s^* \in S$ is a Nash equilibrium [33] if no participant can unilaterally improve its payoff,

$$u_p\big(s_p^*, s_{-p}^*, x^*\big) \;\geq\; u_p\big(s_p, s_{-p}^*, x^*\big) \quad \forall s_p \in S_p,\ \forall p \in \mathcal{P}, \tag{12}$$

where x^* is the physical state induced by s^* through (7) and (9). In the proposed architecture, each participant p is implemented by one or more AI agents that operate on the DAG-based ledger and smart-contract layer and that learn policies intended to approximate such equilibrium strategies while respecting ESG and regulatory constraints.

4.4 AI Agent roles and ontology

In the AI architecture, each market participant $p \in \mathcal{P}$ is implemented by one or more AI-based decision services. We refer to these services as AI agents and use a common ontology of roles across participants:

- *Data/state agents* ingest IIoT measurements, on-chain transaction data and external signals, and maintain a consistent internal state that other agents can query.
- *Forecasting agents* map historical and real-time data to probabilistic forecasts of generation, demand, prices and network conditions over the relevant horizons.
- *Decision agents* (for example bidding, scheduling and storage-control agents) transform forecasts, risk preferences and ESG constraints into actions such as market orders, dispatch schedules or curtailment decisions.
- *ESG and compliance agents* compute ESG indicators from IIoT and contract metadata and enforce regulatory and contractual constraints when other agents propose actions.

These roles appear in different combinations for different participants. In the following, we focus first on the producer-side decision system and then on automated market making. In both cases, agents read from and write to the DAG-based ledger: the ledger provides a shared, tamper-evident source of truth for trades, contract states and ESG metadata that all agent roles can access.

4.5 Producer-side decision system

From the viewpoint of the market game, an IIoT-enabled energy supplier or producer aggregator is represented by a single market participant $p^{\text{prod}} \in \mathcal{P}$ that controls a portfolio of physical assets. Let $\mathcal{A}^{\text{prod}}$ be the set of generation and storage assets operated by p^{prod}, each connected to a node $n(a) \in \mathcal{N}_{\text{phys}}$. For each asset $a \in \mathcal{A}^{\text{prod}}$ and time $t \in \mathcal{T}$, let $G_{a,t}$ denote electrical output, $C_{a,t}$ curtailed output, and $B^{\text{ch}}_{a,t}$ and $B^{\text{dis}}_{a,t}$ storage charging and discharging decisions where applicable. These variables contribute to the nodal balances in (6) and to the compact network constraints (7). The producer also submits orders to the set of short-term markets $\mathcal{M}$, with $Q^{\text{prod}}_{m,z,t}$ denoting its net traded quantity in market m at zone z and time t.

The producer's economic motive is to obtain revenue from selling energy and flexibility while covering operating costs and respecting contractual and ESG requirements. Its trading payoff can be written as

$$\Pi^{\text{energy}}_{p^{\text{prod}}}(s,x) = \sum_{m\in\mathcal{M}} \sum_{z} \sum_{t\in\mathcal{T}} P_{m,z,t}(s,x)\, Q^{\text{prod}}_{m,z,t}(s,x), \tag{13}$$

and its total payoff takes the form

$$u_{p^{\text{prod}}}(s,x) = \Pi^{\text{energy}}_{p^{\text{prod}}}(s,x) - \text{Cost}_{p^{\text{prod}}}(x) - \Gamma_{p^{\text{prod}}}\big(\text{ESG}(x)\big), \tag{14}$$

where $\text{Cost}_{p^{\text{prod}}}$ covers fuel, maintenance, start-up, balancing and contract-related costs, and $\Gamma_{p^{\text{prod}}}$ captures ESG-related penalties or rewards, for example through certificate revenues or financing advantages.

At the level of the producer-side decision system, the induced optimisation problem can be expressed as

$$\max_{s_{p^{\text{prod}}},\, x_{p^{\text{prod}}}} \quad u_{p^{\text{prod}}}\big(s_{p^{\text{prod}}}, s_{-p^{\text{prod}}}, x\big) \tag{15}$$

subject to:

$$A_{\text{phys}}x = b_{\text{phys}}, \quad 0 \le x \le \bar{x} \tag{16}$$

$$\text{ESG}(x) \ge \underline{e} \tag{17}$$

$$\text{SC\_consistency}\big(s_{p^{\text{prod}}}, x\big) = 0 \tag{18}$$

$$Q^{\text{prod}}_{m,z,t}(s,x) = \sum_{a\in\mathcal{A}^{\text{prod}}} \phi_{a,m,z,t}\big(G_{a,t}, B^{\text{ch}}_{a,t}, B^{\text{dis}}_{a,t}, C_{a,t}\big) \quad \forall m,z,t, \tag{19}$$

where $s_{p^{\text{prod}}}$ denotes the strategy components controlled by the producer, $x_{p^{\text{prod}}}$ the subset of physical variables related to its assets, and $s_{-p^{\text{prod}}}$ the strategies of all other participants. Constraint (16) enforces physical feasibility, (17) enforces ESG thresholds, and (18) requires consistency between producer strategies, smart-contract logic and ledger updates. The mapping (19) links asset-level operational decisions to market orders via allocation functions $\phi_{a,m,z,t}(\cdot)$ that implement the producer's internal hedging and portfolio allocation rules. In the proposed architecture, $s_{p^{\text{prod}}}$ and $x_{p^{\text{prod}}}$ are implemented by a multi-agent AI system whose AI agents (for forecasting, bidding, storage control and ESG compliance) operate on IIoT data and DAG-based smart contracts to approximate solutions of (15)–(19).

To connect this payoff to the day-ahead and intraday formulations, we decompose the trading term into contributions from the Day-Ahead Market (DAM) and the Intraday Market (IDM). Let $Q^{\text{prod}}_{\text{DA},z,t}$ and $Q^{\text{prod}}_{\text{ID},z,t}$ denote the net quantities that the producer clears in the DAM and IDM, respectively, at zone z and time t, and let $\text{Price}_{\text{DA},z,t}$ and $\text{Price}_{\text{ID},z,t}$ be the corresponding clearing prices (with $\text{Price}_{\text{DA},z,t}$ consistent with (1)). The trading payoff in (13) can then be written as

$$\Pi^{\text{energy}}_{p^{\text{prod}}}(s,x) = \sum_{z} \sum_{t\in\mathcal{T}} \text{Price}_{\text{DA},z,t}(s,x)\, Q^{\text{prod}}_{\text{DA},z,t}(s,x) + \sum_{z} \sum_{t\in\mathcal{T}} \text{Price}_{\text{ID},z,t}(s,x)\, Q^{\text{prod}}_{\text{ID},z,t}(s,x). \tag{20}$$

Imbalances between scheduled and actual generation, as defined in (2), determine how much of the position is adjusted in the IDM and through balancing mechanisms, and therefore how strongly the producer is exposed to intraday and imbalance prices.

Battery storage units in $\mathcal{A}^{\text{prod}}$ provide an additional degree of freedom. For a storage asset at time t, charging and discharging decisions $(B^{\text{ch}}_{a,t}, B^{\text{dis}}_{a,t})$ imply an energy transfer $\text{EnergyTransferred}_t$, and the arbitrage profit between two price levels $(\text{Price}_{\text{buy},t}, \text{Price}_{\text{sell},t})$ follows the local relation in (3). In the portfolio-level payoff (20), this appears as the contribution of storage-driven trades to $Q^{\text{prod}}_{\text{DA},z,t}$ and $Q^{\text{prod}}_{\text{ID},z,t}$, together with associated efficiency losses and degradation costs in $\text{Cost}_{p^{\text{prod}}}(x)$. The storage-control component of the producer-side decision system uses these spreads between DAM and IDM prices, and between low and high intraday prices, to schedule charging and discharging such that the overall payoff (14) is improved while respecting the physical and ESG constraints.

4.6 Multi-agent AI system for producer-side decision making

From the perspective of the producer, the optimisation problem (15)–(19) involves heterogeneous tasks: forecasting generation and prices, scheduling assets, forming bids in multiple markets, managing storage and enforcing ESG and contractual constraints. A monolithic control system becomes difficult to design, maintain and verify under these requirements.

We propose that the producer implements its strategy components $s_{p^{\text{prod}}}$ through a multi-agent AI system in which specialised AI agents act on shared data from IIoT infrastructure and the DAG-based ledger. The purpose of this system is to approximate high-quality solutions of (15) in real time, adapt to changing market and system conditions, and provide an auditable link between decisions, on-chain transactions and ESG indicators. The producer-side multi-agent system is structured into specialised roles:

- *Data and state aggregation agent.* This agent ingests IIoT measurements from assets in $\mathcal{A}^{\text{prod}}$ (for example power output, state-of-charge, availability) together with on-chain information from the DAG ledger, such as executed trades, contract states and recent price paths. It maintains a consistent internal state that other agents can query.

- *Forecasting agent.* This agent combines weather data, historical IIoT measurements and market information to generate probabilistic forecasts of generation, load and prices for the relevant markets and zones, aligned with the DAM and IDM horizons. Its outputs feed into bidding and storage decisions.

- *Bidding and scheduling agent.* This agent maps forecasts, risk preferences and ESG constraints into DAM and IDM orders $Q^{\text{prod}}_{\text{DA},z,t}$ and $Q^{\text{prod}}_{\text{ID},z,t}$, as well as internal production and curtailment schedules for assets in $\mathcal{A}^{\text{prod}}$. It implements producer-specific hedging and portfolio allocation rules consistent with (19).

- *Storage-control agent.* For assets with storage capabilities, this agent schedules charging and discharging actions $(B^{\text{ch}}_{a,t}, B^{\text{dis}}_{a,t})$ to exploit price spreads between low and high price periods as captured in (20) and (3), subject to technical limits and ESG constraints.

- *ESG-compliance agent.* This agent computes ESG indicators from IIoT data and contract metadata on the DAG, evaluates them against thresholds $\underline{e}$, and enforces these bounds in the producer's decisions by constraining admissible actions and adjusting cost and penalty terms in (14).

All agents read and write to the DAG-based ledger through smart contracts. The ledger serves as a shared, tamper-evident source of truth for executed trades, contract states and ESG-relevant metadata, and as a confirmation layer for actions taken by the producer-side agents. In the following, we abstract from algorithmic details and summarise the structure of the producer-side multi-agent system in terms of agent roles and data flows, as shown in Fig. 1.

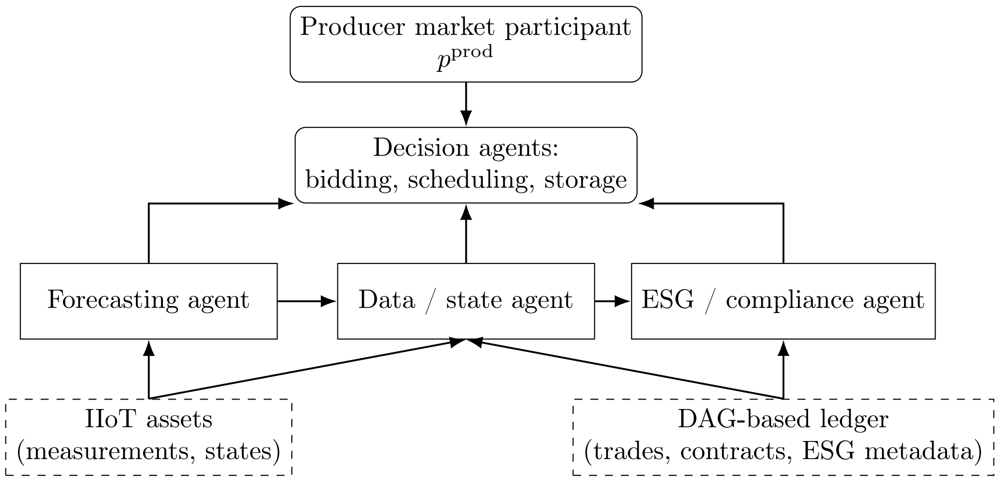

Figure 1: Ontology of the producer-side multi-agent AI system. The producer market participant p^{prod} is implemented by AI agents in distinct roles that operate on shared data from IIoT assets and the DAG-based ledger.

Trading and scheduling decisions proposed by the bidding and storage-control agents are submitted as smart-contract transactions, while the data aggregation and ESG-compliance agents continuously monitor on-chain events to update their internal state.

In practice, the producer-side multi-agent system relies on established methods from stochastic optimisation and reinforcement learning. Scenario generation and uncertainty quantification for renewable production and prices can be handled via Monte Carlo simulation [34, 35]. Given stochastic processes for prices P_t and generation G_t, the forecasting agent draws N scenarios

$$\big\{(P_t^{(n)}, G_t^{(n)})_{t\in\mathcal{T}}\big\}_{n=1}^{N}, \tag{21}$$

and approximates expectations in (14) by the sample average

$$\mathbb{E}\big[u_{p^{\mathrm{prod}}}(s,x)\big] \;\approx\; \frac{1}{N}\sum_{n=1}^{N} u_{p^{\mathrm{prod}}}\big(s, x^{(n)}\big). \tag{22}$$

This leads to a sample average approximation of the producer problem

$$\max_{s_{p^{\mathrm{prod}}}} \quad \frac{1}{N}\sum_{n=1}^{N} u_{p^{\mathrm{prod}}}\big(s_{p^{\mathrm{prod}}}, s_{-p^{\mathrm{prod}}}, x^{(n)}\big) \tag{23}$$

subject to:

$$A_{\mathrm{phys}} x^{(n)} = b_{\mathrm{phys}}, \quad 0 \le x^{(n)} \le \bar{x}, \quad \mathrm{ESG}\big(x^{(n)}\big) \ge \underline{e}, \quad n = 1, \dots, N, \tag{24}$$

$$\mathrm{SC\_consistency}\big(s_{p^{\mathrm{prod}}}, x^{(n)}\big) = 0, \quad n = 1, \dots, N, \tag{25}$$

with optional risk measures (for example CVaR) included by replacing the sample average with an appropriate risk functional.

For sequential decision making and adaptation to other strategic participants, the AI agents can use reinforcement learning and multi agent reinforcement learning techniques [36, 37, 38]. In a value based RL formulation, an agent with state x_t (including prices, asset states and on chain information) and action a_t (for example adjusting bids or storage schedules) learns an action value function

$$Q(x_t, a_t) = \mathbb{E}\left[\sum_{\tau=t}^{T} \gamma^{\tau-t} r_\tau \,\middle|\, x_t, a_t\right], \tag{26}$$

where r_τ is a reward derived from $u_{p^{\mathrm{prod}}}$ and $\gamma \in (0,1]$ is a discount factor. A standard temporal difference update with learning rate α is

$$Q_{t+1}(x_t, a_t) = Q_t(x_t, a_t) + \alpha\Big(r_t + \gamma \max_{a'} Q_t(x_{t+1}, a') - Q_t(x_t, a_t)\Big). \tag{27}$$

Policy gradient and actor critic methods parameterise a policy $\pi_\theta(a \mid x)$ and update parameters θ along estimated performance gradients [36]. In the multi agent setting, several producer side agents and external participants co evolve their policies; practical algorithms approximate equilibrium or no regret behaviour rather than solving (12) in closed form [37, 38]. The DAG based ledger provides the shared, time stamped data stream of actions, outcomes and ESG indicators on which these learning procedures condition.

4.7 Automated market makers in the proposed market design

Automated market makers (AMMs) and constant-function market makers (CFMMs) provide a programmable alternative to traditional limit-order books. In both cases, liquidity is pooled in a smart contract and prices are quoted as a deterministic function of current pool balances rather than via explicit order matching [39, 40, 41]. An AMM or CFMM smart contract holds reserves of one or more assets and enforces a pricing rule through an invariant of the form

$$F(R) = k, \tag{28}$$

where R is the vector of pool reserves, F is a predefined function and k is a constant. Trades adjust the reserves while keeping (28) satisfied, and the marginal price of one asset in terms of another is determined by the gradient of F at the current reserve vector [40, 41]. This mechanism can be implemented entirely as on-ledger logic and executed without an order book.

A widely used AMM design is the constant-product model, where the pool holds reserves X and Y of two assets (for example an energy-linked token and a settlement token) and maintains the invariant

$$XY = k, \tag{29}$$

for some constant $k > 0$ [42, 43]. A trade that adds ΔX units of asset X to the pool and removes ΔY units of asset Y must satisfy

$$(X + \Delta X)(Y - \Delta Y) = k, \tag{30}$$

so that the invariant (29) remains unchanged. The marginal price of X in terms of Y is given by the reserve ratio

$$P_{X/Y} = \frac{Y}{X}, \tag{31}$$

which adjusts endogenously as trades change the pool composition. Constant-product AMMs are a special case of the CFMM framework in (28) with $F(R) = XY$, and serve as the baseline mechanism on which more advanced CFMM designs, such as constant-sum, hybrid and concentrated-liquidity models, build [42, 41].

Beyond the constant-product case, CFMMs specify more general invariants of the form (28) with $R = (R_1, \dots, R_d)$ and F a smooth function [40, 41]. For an infinitesimal trade that increases reserve R_i and decreases reserve R_j while preserving (28), the marginal price of asset i in terms of asset j follows from the total differential

$$\sum_{\ell=1}^{d} \frac{\partial F}{\partial R_\ell}(R)\, \mathrm{d}R_\ell = 0 \quad \Rightarrow \quad P_{i/j}(R) = \frac{\mathrm{d}R_j}{\mathrm{d}R_i} = -\frac{\partial F/\partial R_i}{\partial F/\partial R_j}. \tag{32}$$

Constant-sum, hybrid and stable-swap designs correspond to particular choices of F that stabilise prices for tightly correlated assets, while concentrated-liquidity CFMMs, such as Uniswap v3, define F piecewise over price intervals and restrict non-zero liquidity to selected ranges [42, 41]. This increases capital efficiency and allows fine-grained shaping of price response at the expense of more complex inventory, fee and risk management.

From the perspective of the multi-agent AI system implementing the AMM/CFMM participant p^{amm}, the choice of pricing rule introduces additional decision variables and state dimensions. The data/state agent must track not only pool reserves R and recent order flow but also CFMM parameters, such as segment boundaries and weights in a concentrated-liquidity design. The pricing and liquidity-management agent treats these parameters as controllable variables θ^{cfmm} that influence the invariant $F(\cdot;\theta^{\mathrm{cfmm}})$ and hence the price function (32). Its task is to adapt θ^{cfmm} and fee levels $\varphi_{m,z,t}$ based on forecasts of order flow and volatility in order to improve its payoff $u_{p^{\mathrm{amm}}}(s,x)$ while respecting inventory and ESG constraints. In a learning-based implementation, this agent's action space consists of discrete or continuous updates to θ^{cfmm} and $\varphi_{m,z,t}$, and its state includes the reserve vector R, the derived price sensitivities in (32), and ESG-related indicators, all supplied by the data/state and ESG agents.

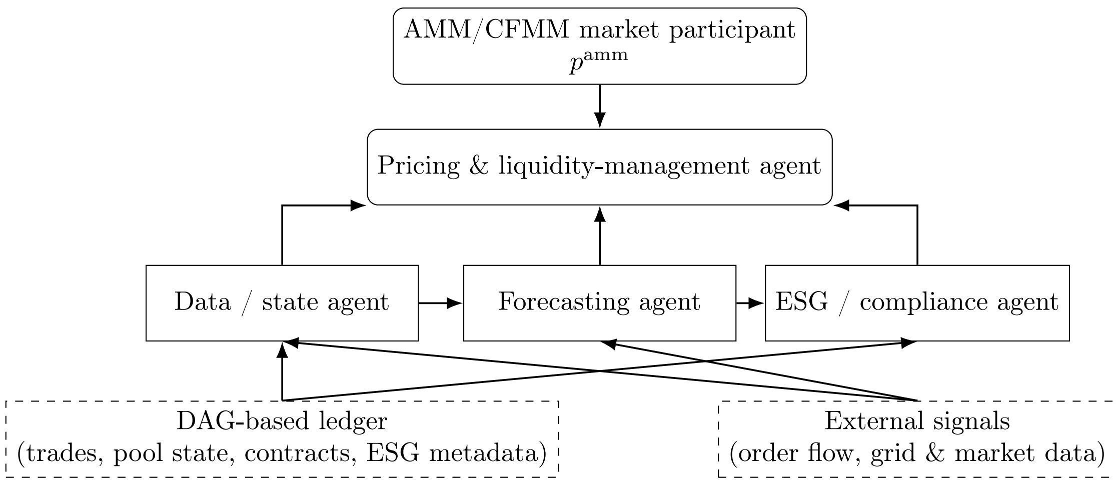

Figure 2: Ontology of the AMM/CFMM-side multi-agent AI system. The AMM/CFMM participant p^{amm} is implemented by software agents in distinct roles. The DAG-based ledger acts as the primary source of truth for trades, pool state, contracts and ESG metadata, complemented by external system and market signals.

After specifying the AMM/CFMM pricing rules and their implications for the multi-agent AI system, we summarise the internal structure of the AMM/CFMM participant in terms of agent roles and data flows, as shown in Fig. 2.

4.8 Implications for energy trading practice

The proposed architecture treats energy and related ESG attributes as tokenised products that are traded through automated market makers on a decentralised DAG-based ledger. Physical production and consumption remain in the power system, while token flows represent contractual claims on energy delivery and ESG characteristics such as origin, emissions intensity or certification status. Producers, suppliers and other participants interact through specialised AI agents that observe IIoT measurements, forecasts and market data, and that submit orders or liquidity updates to AMM and CFMM pools.

On the producer side, decision agents coordinate unit commitment, storage operation and procurement or sale of energy tokens. Their policies take into account physical constraints, portfolio positions and ESG-related objectives or obligations. On the market side, AMM and CFMM pools provide continuous pricing and liquidity for differentiated energy products. The ledger records trades, pool states and ESG attestations in an immutable structure that can be queried by compliance agents and external stakeholders. The overall trading arrangement therefore links operational decisions in industrial energy systems, decentralised price formation for ESG-tagged energy and verifiable tracking of environmental attributes.

This trading arrangement changes the role and incentives of market participants. Producers face a richer product space, since they can sell energy in forms that differ by delivery profile, location and ESG attributes, and they can choose how much liquidity to provide to AMM pools rather than only submitting discrete orders to centralised auctions or brokers. Algorithmic decision agents internalise a wider set of constraints and objectives, which can lead to more systematic exploitation of flexibility, for example through coordinated use of storage and demand-side management.

Consumers and retailers can express preferences not just over price and time of delivery, but also over ESG properties of the purchased energy. Their agents can trade continuously in small increments, rebalance positions as new information arrives and adjust portfolios in response to updated ESG constraints or disclosure requirements. Liquidity provision through AMM pools gives them a transparent rule for price formation, and makes the cost of specific ESG preferences explicit in the marginal price of the corresponding tokens.

For system operators and regulators, the architecture provides more granular data on physical and contractual positions. The ledger and its ESG tracking layer produce machine-readable histories of production, transfer and consumption that can be audited by compliance agents. At the same time, this creates new requirements for validation of on-ledger data against physical measurements and for supervision of algorithmic trading strategies implemented by decision agents. Governance arrangements become more important, since platform rules, pool parameters and admissible agent behaviours shape both market outcomes and the credibility of ESG claims.

Conventional electricity trading in many jurisdictions relies on a combination of centralised day-ahead auctions, continuous intraday order book trading and bilateral over-the-counter contracts, cleared through recognised counterparties. Products are typically standardised by delivery period and location, while ESG attributes are handled through separate certificate schemes. Trading decisions are made by human traders supported by analytics tools, and contract data are stored in proprietary systems, with only aggregated information visible to the wider market.

In contrast, the proposed architecture integrates energy and ESG attributes into a single token space and uses AMM and CFMM pools for price formation and liquidity. Price adjustment is continuous and endogenous to pool flows, rather than the result of periodic auctions or order book matching. Market access is mediated by AI agents that operate on shared state and contract data, which reduces manual intervention and enables faster reactions to system changes. Settlement and position management occur directly on the ledger, which reduces reconciliation overhead but introduces new forms of operational and cyber risk.

The architecture does not eliminate existing exchanges and bilateral contracts. Instead, it can be viewed as an additional layer that tokenises positions and ESG attributes and that provides continuous trading and hedging possibilities between established market time frames. Producers and consumers can still use day-ahead and intraday markets as reference and can calibrate AMM pool parameters to align with external price signals. The main differences are the degree of automation in decision making, the integration of ESG tracking into the core trading process and the use of a decentralised ledger and AMM mechanism for contract management and price discovery.

5 Discussion and outlook

5.1 Alignment between decentralised incentives and system-level objectives

The proposed architecture combines individual participant payoffs, the system-level objective (11), the market game (12) and AMM/CFMM pricing rules into a single interacting system. A central question is whether Nash equilibria of the market game are aligned with the weighted welfare function $W(s, x)$ that aggregates stakeholder utilities and ESG penalties. In the current formulation, each participant $p \in \mathcal{P}$ maximises its own payoff (10), while $W(s, x)$ is a conceptual benchmark for regulators, platform operators and governance bodies.

Alignment depends on several design choices. The weighting parameters ω_k and $\lambda_E, \lambda_S, \lambda_G$ in (11) encode a policy stance on the trade-off between economic outcomes and ESG performance. At the same time, ESG penalties and rewards enter directly into participant payoffs through Γ_p, which influences the set of best responses and thus the Nash equilibria (12). If penalties are too weak or poorly targeted, equilibria may remain economically efficient in a narrow sense while still violating desired ESG outcomes. If penalties are too strong or mis-specified, equilibria may shift towards strategies that satisfy formal ESG metrics but create hidden risks for reliability or long-term investment.

The multi-agent AI layer provides an instrument to steer this alignment. Decision and ESG agents implement the payoff structure in operational policies rather than in static contracts. This permits regulators or platform operators to update ESG-related terms in Γ_p and corresponding constraints without redesigning the entire market. It also creates a need for robust governance, since changes in policies or ESG weightings propagate through learned strategies and may induce transitions between qualitatively different equilibria. The architecture shifts part of the regulatory challenge from designing once-off rules to managing an evolving set of algorithmic incentives.

5.2 Interplay between AMM design, physical constraints and ESG tracking

The AMM and CFMM layer, expressed through the invariant $F(R) = k$ in (28) and its special cases such as (29) and (30), provides a rule-based mapping from order flow to prices. This mapping interacts with both physical constraints and ESG tracking. In a traditional market, the link between physical scarcity and prices is mediated by centralised auctions, order books and operator decisions. Here, the link is partly delegated to an on-ledger pricing function that reacts mechanically to changes in pool reserves.

This delegation has concrete implications. Pool design determines how price sensitivity varies with inventory in the pool and thus shapes the risk of large price swings for illiquid ESG-differentiated products. Constant-product pools penalise trades that move reserves towards extremes, which provides an implicit liquidity risk premium for large orders. For energy products with strong physical coupling across locations and time periods, such penalties can amplify local scarcity signals if pool composition reflects network conditions. If pool composition drifts away from physically motivated ratios, prices may signal artificial scarcity or abundance, which then feeds back into producer and consumer agent decisions.

The DAG-based ledger closes a loop between physical events, ESG tracking and pricing. IIoT data and ESG attestations flow into state agents, update on-ledger records and change the effective eligibility of tokens for ESG-sensitive pools. Pool reserves thus depend on both trading strategies and the validity of ESG tags. In turn, prices influence incentives to invest in low-emission assets, flexibility and data quality. This creates a coupled system where data integrity, physical operations and token liquidity must be considered together. AMM design becomes an instrument not only for managing financial liquidity but also for stabilising the interaction between physical and ESG constraints.

5.3 Overall formulation of the architecture

From a design perspective, the architecture can be summarised by an optimisation problem that captures the intended relationship between decentralised strategies, physical feasibility, ledger dynamics and AMM pricing. Let $s = (s_p)_{p \in \mathcal{P}}$ denote the joint strategy profile, and let $x(s)$, $L(s)$ and $R(L(s))$ denote the induced physical state, ledger state and AMM reserve vectors. A compact formulation of the design goal is

$$\begin{aligned} \max_{s \in S} \quad & W(s, x(s)) \\ \text{s.t.} \quad & x(s) \in \mathcal{X}_{\text{phys}}, \\ & L(s) \in \mathcal{X}_{\text{ledger}}, \\ & F(R(L(s))) = k, \\ & s \text{ satisfies the Nash conditions (12).} \end{aligned} \tag{33}$$

The first two constraints express physical and ledger-level feasibility, including network limits, device constraints and basic consistency of token flows. The third constraint enforces the AMM or CFMM invariant that defines how prices adjust to changes in reserves. The final condition expresses the requirement that the implemented strategies form a Nash equilibrium under the given market rules and payoffs.

In practice, no single decision-maker solves (33) directly. The problem instead serves as a conceptual reference for mechanism design. Platform operators choose the function F, pool parameters, admissible strategy sets and ESG weightings with the aim that the resulting Nash equilibria approximate the maximisers of W. The multi-agent AI layer then learns and adapts strategies within these constraints. This formulation highlights that architecture and governance choices are inseparable from equilibrium outcomes, and that redesigning AMM rules or ESG penalties amounts to changing the feasible set and objective in (33).

5.4 Limitations and directions for further work

The conceptual framework abstracts from several layers of complexity that would matter in deployment. First, the description of strategies s_p and payoffs in (10) and (12) hides rich temporal structure. Real participants operate under multi-period contracts, ramping constraints, evolving ESG regulation and learning dynamics of AI agents. A more detailed model would represent the architecture as a stochastic dynamic game in which policies map observable states to actions and where the ledger and AMM states enter explicitly into the state vector.

Second, the model treats ESG indicators E, S and G as well-defined functions of states and actions. In practice, measurement and verification of ESG performance rely on heterogeneous data sources, protocols and certifiers. Inconsistent or low-quality data can undermine the effectiveness of penalties and rewards. The architecture assumes that data and attestations that reach the ledger are trustworthy enough to serve as inputs for decision agents. Research on detection of manipulation, cross-checking with independent measurements and design of ESG metrics that are robust to strategic behaviour becomes central in such a setting.

Third, the AMM specification in (28) and its constant-product special case are stylised representations of liquidity provision. Energy markets involve correlated risks across time, space and ESG attributes that may require richer pool structures, such as multi-dimensional reserves, time-coupled invariants or hybrid designs that combine order books and CFMMs. Studying how such designs propagate into participant risk exposures, equilibrium prices and system-level indicators is a natural next step.

Finally, the multi-agent AI layer introduces its own class of risks. Decision agents optimise with respect to modelled payoffs and constraints, which may omit rare events, structural breaks or complex failure modes in power systems and markets. Safety considerations then require constraints on admissible policies, monitoring of agent behaviour and mechanisms for human intervention in exceptional situations. Simulation environments that couple detailed power system models, market rules, AMM dynamics and learning agents would provide a test bed for such safety concepts and for the calibration of ESG-related incentives.

The outlook of this architecture is shaped by these limitations and opportunities. The combination of on-ledger AMM-based trading, ESG tracking and multi-agent AI decision-making opens a design space in which financial, physical and sustainability objectives can be treated jointly. At the same time, this design space transfers part of the traditional work of human traders, regulators and system operators into the choice of algorithms, incentives and governance for autonomous agents. Future work will need to explore this space systematically, with a focus on concrete deployment settings, regulatory constraints and empirical evidence from pilot implementations.

References

1. Caldarelli G. Integration of Blockchain in Accounting and ESG Reporting: A Systematic Review from an Oracle-Based Perspective. Journal of Risk and Financial Management 2025; 18:491. DOI: 10.3390/jrfm18090491
2. Suta A and Tóth Á. Systematic Review on Blockchain Research for Sustainability Accounting Applying Methodology Coding and Text Mining. Cleaner Engineering and Technology 2023; 14:100648. DOI: 10.1016/j.clet.2023.100648
3. Wu W, Fu Y, Wang Z, and Liu X. Consortium Blockchain-enabled Smart ESG Reporting Platform with Token-based Incentives for Corporate Crowdsensing. Computers & Industrial Engineering 2022; 172:108456. DOI: 10.1016/j.cie.2022.108456
4. Ahmad T and Zhang D. Using the Internet of Things in Smart Energy Systems and Networks. Sustainable Cities and Society 2021; 68:102783. DOI: 10.1016/j.scs.2021.102783
5. Sithole KD, Lehloka MC, and Monchusi BB. Employing Internet of Things (IoT) Devices for Monitoring and Controlling Energy Management Systems: A Review. Journal of Electrical Systems 2024; 20:753–60. DOI: 10.52783/jes.7273
6. Borkovcová A, Černá M, and Sokolová M. Blockchain in the Energy Sector—Systematic Review. Sustainability 2022; 14:14793. DOI: 10.3390/su142214793
7. Egunjobi OO, Gomes A, Egwim CN, and Morais H. A Systematic Review of Blockchain for Energy Applications. e-Prime: Advances in Electrical Engineering, Electronics and Energy 2024; 9:100751. DOI: 10.1016/j.prime.2024.100751
8. Bessa RJ, Moreira C, Silva B, and Matos MA. Handling Renewable Energy Variability and Uncertainty in Power Systems Operation. Wiley Interdisciplinary Reviews: Energy and Environment 2014; 3:156–78. DOI: 10.1002/wene.76
9. Shahidehpour M, Yamin H, and Li Z. Market Operations in Electric Power Systems: Forecasting, Scheduling, and Risk Management. New York: Wiley-IEEE Press, 2002
10. Tanrisever F, Shahmanzari M, and Büke B. European Electricity Day-Ahead Markets: A Review of Models and Solution Methods. SSRN Electronic Journal 2020. DOI: 10.2139/ssrn.3517267
11. Gao X, Yang C, Zhu Z, Wu Y, Chen Y, and Mu Z. A Mini-review on Trading Mechanisms of Emerging Joint Energy Markets with High Shares of Renewables. Frontiers in Energy Research 2024; 12:1391813. DOI: 10.3389/fenrg.2024.1391813
12. He G, Chernyakhovskiy I, and Rômeo D. Estimating the Economic Value of Behind-the-Meter Battery Storage for Power System Operators. Applied Energy 2016; 182:322–30. DOI: 10.1016/j.apenergy.2016.08.189
13. Schill WP, Zerrahn A, and Kunz F. Prosumage of Solar Electricity: Pros, Cons, and the System Perspective. Economics of Energy & Environmental Policy 2017; 6:7–32. DOI: 10.5547/2160-5890.6.1.wsch
14. Gao X, Yang C, Zhu Z, Wu Y, Chen Y, and Mu Z. A mini-review on trading mechanisms of emerging joint energy markets with high shares of renewables. Frontiers in Energy Research 2024; 12:1391813. DOI: 10.3389/fenrg.2024.1391813

15. González P, Villar J, and Campos FA. Joint Energy and Reserve Markets: Current Implementations and Modeling Trends. Electric Power Systems Research 2014; 109:101–11. DOI: 10.1016/j.epsr.2013.12.013

16. Gan D and Litvinov E. Energy and Reserve Market Designs with Explicit Consideration to Lost Opportunity Costs. IEEE Transactions on Power Systems 2003; 18:53–9. DOI: 10.1109/TPWRS.2002.807061

17. Shivaie M, Kiani-Moghaddam M, and Weinsier PD. Bilateral Bidding Strategy in Joint Day-Ahead Energy and Reserve Electricity Markets Considering Techno-Economic-Environmental Measures. Energy and Environment 2022; 33:696–727. DOI: 10.1177/0958305X211014875

18. Block CA, Collins J, Ketter W, and Weinhardt C. A Multi-Agent Energy Trading Competition. ERIM Report Series ERS-2009-054-LIS. Erasmus Research Institute of Management, 2009. Available from: https://ssrn.com/abstract=1522683

19. El Hafiane D, El Magri A, Chakir HE, Lajouad R, and Boudoudouh S. A multi-agent system approach for real-time energy management and control in hybrid low-voltage microgrids. Results in Engineering 2024; 24:103035. DOI: 10.1016/j.rineng.2024.103035

20. Raju L and Raj V. Energy Management System Using Multi Agent Systems with IoT and Machine Learning. *2023 International Conference on Circuit Power and Computing Technologies (ICCPCT)*. 2023. DOI: 10.1109/ICCPCT58313.2023.10245738

21. Gazafroudi AS, De Paz JF, Prieto-Castrillo F, Villarrubia G, Talari S, Shafie-Khah M, and Catalão JPS. A Review of Multi-Agent Based Energy Management Systems. *Ambient Intelligence – Software and Applications, 8th International Symposium on Ambient Intelligence (ISAmI 2017)*. Vol. 615. Advances in Intelligent Systems and Computing. Springer, 2017 :203–9. DOI: 10.1007/978-3-319-61118-1_25

22. Ghazimirsaeid SS, Jonban MS, Mudiyanselage MW, Marzband M, Romeral Martinez JL, and Abusorrah A. Multi-Agent-Based Energy Management of Multiple Grid-Connected Green Buildings. Journal of Building Engineering 2023; 74:106866. DOI: 10.1016/j.jobe.2023.106866

23. Vionis P and Kotsilieris T. The Potential of Blockchain Technology and Smart Contracts in the Energy Sector: A Review. Applied Sciences 2024; 14:253. DOI: 10.3390/app14010253

24. Honari K, Rouhani S, Falak NE, Liu Y, Li Y, Liang H, Dick S, and Miller J. Smart Contract Design in Distributed Energy Systems: A Systematic Review. Energies 2023; 16:4797. DOI: 10.3390/en16124797

25. Kirli D, Couraud B, Robu V, and Kiprakis A. Smart Contracts in Energy Systems: A Systematic Review of Fundamental Approaches and Implementations. Renewable and Sustainable Energy Reviews 2022; 158:112013. DOI: 10.1016/j.rser.2021.112013

26. Santos LM, Gomes
bibinitperiod, and Rupino Cunha P. Energy Trading Using Blockchain: Smart Contracts Functionalities – A Systematic Review. Energy Strategy Reviews 2025; 61:101825. DOI: 10.1016/j.esr.2025.101825

27. Park J, Chitchyan R, Angelopoulou A, and Murkin J. A Block-Free Distributed Ledger for P2P Energy Trading: Case with IOTA? *Advanced Information Systems Engineering Workshops*. Vol. 11483. Lecture Notes in Computer Science. Springer, 2019 :111–25. DOI: 10.1007/978-3-030-21290-2_8

28. Baplawat A, Janardhan R, M. G. R, Rafique J, Mahajan V, and Kumari M. Smart Contracts in ESG Reporting: A Fintech-Based Framework for Enhancing Corporate Governance Transparency. Lex Localis - Journal of Local Self-Government 2025; 23:5339–48. DOI: 10.52152/bkv20y90

29. Singh AK, Dugyala NR, Rahimian F, Elghaish F, and Mohandes SR. Blockchain-based Approach to Improve Environmental, Social, and Governance (ESG) Reporting in Construction Organizations. Journal of Information Technology in Construction (ITcon) 2025; 30:1497–527. DOI: 10.36680/j.itcon.2025.061

30. Mulligan C, Morsfield S, and Cheikosman E. Blockchain for Sustainability: A Systematic Literature Review for Policy Impact. Telecommunications Policy 2024; 48:102676. DOI: 10.1016/j.telpol.2023.102676

31. Jiang L, Gu Y, Yu W, and Dai J. Blockchain-based Life Cycle Assessment System for ESG Reporting. SSRN Electronic Journal 2022. DOI: 10.2139/ssrn.4121907

32. Thanasi-Boce M and Hoxha J. Blockchain for Sustainable Development: A Systematic Review. Sustainability 2025; 17:4848. DOI: 10.3390/su17114848

33. Nash JF. Non-Cooperative Games. Annals of Mathematics 1951; 54:286–95. DOI: 10.2307/1969529

34. Shapiro A. Monte Carlo Sampling Methods in Stochastic Programming. Mathematical Programming 2003; 95:559–81. DOI: 10.1007/s10107-002-0349-4

35. Shapiro A, Dentcheva D, and Ruszczyński A. Lectures on Stochastic Programming: Modeling and Theory. Philadelphia: SIAM and MPS, 2009. DOI: 10.1137/1.9780898718751

36. Sutton RS and Barto AG. Reinforcement Learning: An Introduction. 2nd ed. Cambridge, MA: MIT Press, 2018

37. Busoniu L, Babuska R, De Schutter B, and Ernst D. A Comprehensive Survey of Multiagent Reinforcement Learning. IEEE Transactions on Systems, Man, and Cybernetics, Part C 2008; 38:156–72. DOI: 10.1109/TSMCC.2007.913919

38. Wallace C, Galloway S, Burt G, et al. A Review of Multi Agent Reinforcement Learning for Smart Grids. Renewable and Sustainable Energy Reviews 2020; 115:109369. DOI: 10.1016/j.rser.2019.109369

39. Aramonte S, Huang W, and Schrimpf A. Trading in the DeFi Era: Automated Market-Makers. BIS Quarterly Review 2021 Dec. Box article in BIS Quarterly Review, December 2021. Available from: https://www.bis.org/publ/qtrpdf/r_qt2112b.htm

40. Xu J, Paruch K, Cousaert S, and Feng Y. SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) Protocols. ACM Computing Surveys 2023; 55:1–50. DOI: 10.1145/3570639

41. Bichuch M and Feinstein Z. Axioms for Automated Market Makers: A Mathematical Framework in FinTech and Decentralized Finance. arXiv preprint arXiv:2210.01227. 2022. DOI: 10.2139/ssrn.4405316

42. Adams H, Zinsmeister N, and Robinson D. Uniswap v2 Core. Whitepaper. Available from the Uniswap documentation site. 2020

43. Zhang Y, Chen X, and Park D. Formal Specification of Constant Product ($x \times y = k$) Market Maker Model and Implementation. Technical report. Preprint on constant-product automated market makers. Runtime Verification Inc., 2018