Multi-agent reinforcement learning for Grid Balancing using Bitcoin Mining

Ioannis T. Thomaidis
School of Social Sciences
Hellenic Open University
Patras, Greece
0009-0000-2977-8367

Panagiotis G. Giannopoulos
School of Social Sciences
Hellenic Open University
Patras, Greece
0000-0001-7530-2458

Panos Chountalas
Department of Business Administration
University of Piraeus
Piraeus, Greece
0000-0003-3538-5791

Thomas K. Dasaklis
School of Social Sciences
Hellenic Open University
Patras, Greece
0000-0003-1240-4822

***Abstract*—Maintaining balance in the electrical grid is crucial for assuring stability, minimizing energy waste and integrating renewable energy sources. This paper introduces a multi-agent reinforcement learning (MARL) framework whereby Renewable Energy Producers (REPs), BTC Miners, and Energy Manager Agents (EMAs) interact dynamically to stabilize the grid. The system utilizes energy market data and reinforcement learning algorithms to motivate miners to modify their power usage habits according to the prevailing grid circumstances. Miners, as energy consumers within the system, enable the absorption of excess renewable energy during low-demand periods and reduce use during peak times to relieve grid stress. We use the proposed MARL method in a simulated setting to assess grid stability, profitability and energy efficiency. The findings indicate that BTC mining can markedly reduce the volatility of the grid, improve the profitability of the miner and facilitate the incorporation of renewable energy.**

***Index Terms*—Bitcoin mining, Grid, Balance, Agents, Energy**

I. Introduction

The pressing need for transforming the emissions-intensive nature of fossil-fueled electrical networks into more sustainable and eco-friendly forms, has recently led to the incorporation of various renewable energy sources (RES) into the energy production mixture. In contrast to fossil-fueled energy streams that provide a consistent and reliable electricity supply, RES are intrinsically erratic due to their dependence on meteorological factors. Particularly, this dependence leads to energy surpluses, resulting in curtailment, and energy deficits, which may destabilize the system. When renewable energy (RE) surpasses real-time demand and there is inadequate storage or transmission capacity, surplus power is wasted. Conversely, during times of maximum demand, wind power could not generate enough energy to fulfill the demand, in which case grid management must depend on backup from fossil-fueled power systems, therefore resulting in more expenses and greenhouse gas emissions.

Incorporating RES into the electricity production mixture includes up-taking the risk of curtailment, a phenomenon that often results in cost deficits and hinders the widespread adoption of renewable energy. Although curtailment rates in major renewable energy markets typically range from 1.5% to 4%, these still represent significant losses of clean energy that could otherwise be effectively utilized. [1]. These costs primarily stem from both the loss of energy that could have been generated but is instead wasted and the additional expenses required to manage grid stability. Recent shortfalls attributed to curtailed energy are found in different countries worldwide. For instance, in Australia, solar curtailment rose from 4% to 7%, reaching peaks of 20% on certain days. In Germany, grid integration issues in 2024 reduced solar photovoltaic production by 2% [2]. In the United States, Texas curtailed 9% of its large-scale solar power output in 2022. China faced rising curtailment, with some provinces exceeding 10% and national rates surpassing 5% in early 2024 [3]. Wind energy curtailment has also escalated. The United Kingdom saw rates fluctuate between 3.1% and 6.7%, particularly due to transmission bottlenecks between Scotland and England.

The above curtailment figures highlight the rising need of extending the boundaries of legacy energy production and management systems towards incorporating elements that could potentially act as uninterruptible energy loads. Towards this direction, Bitcoin (BTC) mining has been recently discussed as a prominent venue capable of monetizing stranded or excess energy sources. In particular, BTC mining and the BTC blockchain infrastructure could be used as potent resource monetization instruments [4]. To this end, BTC mining has surfaced as a prospective demand-side flexibility instrument owing to its energy-intensive but remarkably adjustable characteristics. Very recent studies have investigated the viability of integrating BTC mining with RES, particularly in using surplus power that would otherwise be wasted [5]–[10].

Implementation-wise, the main efforts have focused on developing deterministic, scenario-based models that optimize the integration of BTC mining with RES based on predefined operating assumptions. For instance, studies on wind curtailment monetization [5], electricity cost management [11] and mixed-integer linear programming (MILP)-based energy scheduling [8] operate under predetermined constraints, including historical energy curtailment levels, fixed mining uptime and static tariff structures. Although significant, these studies lack the ability to adapt operations dynamically to real-time grid conditions, market variations or renewable energy intermittency. From a methodological point of view, although reinforcement learning (RL) has proved prominent towards developing self-adaptive mechanisms with enhanced adaptability, the incorporation of RL-based approaches in RES-driven energy schemes are extremely scarce.

In this paper we propose a Multi-Agent RL (MARL) system that models different types of agents, including Renewable Energy Producers (REPs), Energy Manager Agents (EMAs) and BTC Miners. The proposed MARL system operates under the dual scope of decreasing curtailment rate and enhancing the system's profitability. Particularly, our model highlights the BTC mining's dual role as an energy consumer but also as a proactive and intelligent energy-balancing system that is capable of making autonomous and data-driven decisions.

II. Literature review

Although research into renewable energy curtailment and its mitigation strategies is extensive, studies specifically exploring the integration of BTC mining into electricity grids as a flexible load remain relatively limited. It should be noted that BTC mining operations can be easily adjusted according to the availability and price of electricity and, therefore, present a possible answer to the dilemma of how to use surplus energy at times of low energy demand or how to reduce energy consumption when the grid is overloaded.

The authors in [5] conducted a quantitative scenario-based analysis to evaluate BTC mining integration with curtailed wind energy in ERCOT (2011–2020). The study projected energy availability, mining performance and financial returns under fixed curtailment percentages (e.g., 3.69% in 2020, totaling 4.21 TWh curtailed) and predefined operational constraints (e.g., 10,000 miners at 66.38% uptime, 1,000 miners at 93.33% uptime). Furthermore, [7] performed a macroeconomic equilibrium analysis to assess BTC mining's impact on renewable energy investments and electricity markets. Using a long-run equilibrium framework, the study explored how BTC mining demand influences wind and solar capacity expansion under various market conditions. The authors in [11] conducted a deterministic cost minimization study, developing a BTC mining model based on historical power market prices, current tariff structures and conventional cost reduction strategies. They projected a possible reduction in power costs of 3%-11% via wholesale energy procurement, 1%-7% via demand response and up to 20% from self-generation. Nonetheless, the methodology relied on past pricing and hence did not take into account the real-time adjustments in BTC mining energy usage in response to price fluctuations or variations in regional demand. The authors in [8] used Mixed Integer Linear Programming (MILP) scenario optimization to study the integration of BTC mining with restricted wind and solar energy in Texas. While Monte Carlo simulations (1000 and 5000 iterations) were incorporated to assess financial and technical uncertainties, the study's static modeling approach required manual re-optimization whenever market conditions changed, making it computationally intensive for large-scale grid applications. Furthermore, the study found that the C-MIN scenario resulted in a $838 million financial loss in 2020, whereas in 2021, with the same operational conditions but higher BTC prices, profits reached $166.9 million, demonstrating the model's inability to dynamically adjust to BTC market fluctuations. In [9] the authors conducted a deterministic energy scheduling analysis, introducing the Curtailment Energy Storage and Settlement (CESS) model to evaluate Bitcoin mining as a Virtual Energy Storage System (VESS) inside microgrid environments. The study revealed that BTC mining might reduce annual energy curtailment from 190.8 MWh to 0.14 MWh by using surplus renewable energy for mining activities. [10] performed a financial feasibility analysis, simulating the integration of BTC mining with residential photovoltaic systems. In this research, the use of BTC mining to utilize the excess energy produced by the rooftop solar panels led to the reduction of energy costs by 66.4% and the return on investment (ROI) of 50.8%. However, the BTC mining operations were performed with a rigid schedule, which prevented the real-time adaptive load management. [12] made a stochastic MILP-based location optimization study and concluded the best location choice for BTC mining by taking into account the power consumption, the accessibility to the power networks and the trade-off between conventional and RES. This research used stochastic probability distributions to model energy price volatility and supply uncertainty; however, it did not incorporate a dynamic learning mechanism that is able to update the BTC mining decisions in real-time based on the fluctuating grid conditions. Lastly, [13] made a financial feasibility assessment to determine the feasibility of integrating BTC mining with a solar photovoltaic (PV) system in the UAE. The study replicated stable operational conditions, assuming that surplus solar energy was sent into the grid during peak periods, while BTC mining drew power from the system during non-solar hours.

It is clear from the retrieved literature that most studies rely upon static, pre-optimized energy allocation models. Contrary to current approaches, this study reclassifies BTC mining as a variable load energy consumer capable of real-time responses to market volatility, grid stability challenges and evolving renewable energy policy. A novel, scalable and self-optimizing system is presented, enhancing energy monetization, grid flexibility and economic sustainability in unpredictable power markets using MARL decision-making.

III. High-level architecture

To address the overarching need of optimizing the performance of energy management systems, we propose an RL-centric solution, which by enabling energy trading schemes between different types of agents, facilitates the effective management of grid towards both balancing the demand and providing revenues through the incorporation of a BTC mining layer. Our method elaborates upon the capacity of miners to transform squandered energy into a valuable resource, and provides a system-wide solution to incorporate them into grid management systems.

For implementation purposes we draft an architectural approach which seamlessly integrates infrastructural and software elements, as depicted in Figure 1. Specifically, our approach models the distribution system dynamics by incorporating the following types of agents: (i) REPs, (ii) EMA, which act as the central module of intelligence, (iii) BTC miners; and (iv) grid actors, which represent the main consumers of the energy grid. Finding the optimally-performed schema, in such a systemic approach, is quite challenging for mathematical programming methods, thus our approach introduces a multi-agent approach, to model the dynamics of the main operating scenario, as described in Table I.

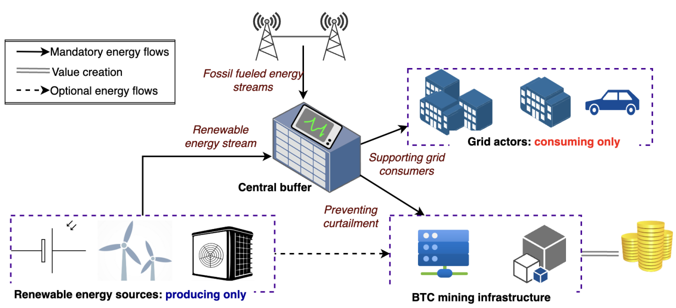

Fig. 1. A high level architecture of RL-centric energy management systems

A. Infrastructural elements

To accurately model real-world energy management systems, our approach encompasses two energy production streams: (i) RES, including solar photovoltaic panels, wind turbines, and geothermal power producers; and (ii) fossil-fueled sources. As mentioned, these two streams exhibit different characteristics, with the renewable energy stream to be subject to variability due to the erraticness of environmental conditions, while the fossil-fueled stream operates at a stable, yet pollution-intensive production rate. The energy generated by these streams is supplied to a central buffer, which represents the EMA in our approach. The EMA has a limited storage capacity, meaning that if the available energy exceeds its capacity, the excess energy should either curtailed or fed into the BTC miners. To mitigate the curtailment issue, our architecture allows both the central buffer (EMA) and REPs to be directly connected to BTC miners, enabling the efficient utilization of excess energy. Both producers integrated into our approach are considered strictly as producing agents, meaning their decision-making primarily revolves around negotiating the selling price of energy.

To simulate the stochastic nature of REPs (R_t) we use the Equation 1, in which η_t captures random fluctuations (e.g., changes in weather), and σ_R^2 represents the variance of these fluctuations.

$$R_t = R_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_R^2) \tag{1}$$

B. Grid actors

The main consumers served by the system are represented as grid actors. In our approach, grid actors encompass various entities connected to the electricity grid, mainly including residential and commercial consumers. For simulation purposes, we model grid demand (D_t) by using Equation 2, in which D_{avg} represents the average demand, A is the amplitude of the fluctuation, ω is the angular frequency (for example, $\omega = \frac{2\pi}{24}$ to represent daily cycles); and ϕ is the phase shift. This assumption allows us to capture realistic variations in energy demand, rather than relying on an over-simplified model with fixed demand rates at specific time intervals.

$$D_t = D_{avg} + A\sin(\omega t + \phi) \tag{2}$$

C. Energy traders: EMA & BTC miners

Beyond the above two types of agents, our approach also models the trading interfaces between BTC miners and EMA, to facilitate both energy stabilization and enhanced profitability for the overall energy management system. These two types of agents are trained towards finding an optimally-performed trading scheme, that safeguards a stable and robust energy distribution to consumers. Within our multi-agent approach, these agents are represented by different reward mechanisms that prioritize enhance profitability in the case of BTC miners; and grid stabilization between supply and demand, in the case of EMA. Notwithstanding their differentiated strategic orientation, when incorporated in centralized training modules, these reward mechanisms enable agents to learn optimal patterns that facilitate the energy monetization towards creating revenue by the RE surpluses, that otherwise would be curtailed (especially through BTC mining).

IV. Implementation details

This section outlines the implementation details of the proposed architecture designed to model the operating scenario presented in Table I. The primary intelligent energy trading entity is the EMA. Its main function is to monitor the energy production from REPs and match it with grid energy demand to ensure optimal allocation. As part of the centralized learning module, EMAs are integrated into a shared environment alongside BTC miners. They use system-wide data on renewable energy availability, grid demand, and surplus energy trends to optimize energy distribution via a coherent training process.

The latter class of agents consists of BTC miners, which utilize excess energy supplied by EMAs to power mining operations. The centralized module adds a financial incentive to energy management systems based on bitcoin miners' potentially endless energy resource management and revenue conversion capabilities. Particularly, it enables EMAs and BTC miners to learn from previous negotiations and experiences, leading them to mutually optimize their energy allocation strategies based on the overall system state. Once the training is completed, the system transitions to decentralized execution, where each EMA operates autonomously, making real-time decisions based solely on its local observations. This decentralized decision-making method enables the system to adapt its responses to renewable energy output and grid demand without the need of a central controller.

TABLE I
SUMMARY OF GRID BALANCING AND BTC MINING SCENARIOS

Objective	AI-driven trade-offs	System constraints
Utilize excess renewable energy for BTC mining while maintaining grid stability and preventing energy curtailment.	- Surplus Energy Determination: Difference between renewable production (R_t) and grid demand (D_t). - Dynamic Grid Demand: Modeled as $D_t = D_{avg} + A\sin(\omega t + \phi)$ to capture demand variations. - Optimal Allocation Fraction: $a_t \in [0,1]$, optimized via Q-learning based on state variables and reward feedback.	- Grid Supply Priority: During peak demand ($R_t \leq D_t$ or low X_t), limit BTC mining ($a_t \approx 0$) to ensure grid stability. - Avoid Over-Allocation: Prevent excessive BTC mining during low-demand periods if it affects grid balance. - Economic Viability: Ensure profitability without compromising grid reliability.

A. Learning module

The learning module employed in our approach aligns with the established Centralized Learning–Decentralized Execution (CL-DE) paradigm. In the first phase of centralized learning, the central EMA is trained jointly in a unified environment that incorporates information about the total system metrics such as renewable production (R_t), grid demand (D_t), and the computed surplus energy (X_t). This allows the agents to establish an integrated and synergistic strategy for energy allocation in the training phase. They update their decision policies using the well-established Q-learning technique, which iteratively refines each agent's Q-values based on the rewards they receive for executing specific actions, by following the Equation 3.

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \lambda \Big[r_t^{\text{EMA}} + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t) \Big] \tag{3}$$

In Equation 3, λ represents the learning rate and γ is the discount factor. By synthesizing the experiences (rewards) received from each type of agents in a centralized matrix, the learning module safeguards that the agents are trained towards finding the optimally-performed policy that balances the energy flows for the system.

Once centralized training is completed, the agents transition to the decentralized execution phase. During this stage, each EMA operates independently, making real-time decisions based solely on its local observation of the system state, defined as $s_t = [R_t, D_t]$. The decentralized execution framework enables faster response to fluctuations in renewable energy production and grid demand, enhancing the scalability and adaptability of the system.

B. Reward mechanisms

To safeguard an optimal performance for our training module, we delve into the operation specifics of the energy management system to design the corresponding reward mechanisms. For the centralized learning component, the reward mechanism is designed to incentivize energy allocation to BTC miners when the storage capacity in the EMA is limited while the demand of consumers has already been met. Particularly, we model the energy surplus in specific time-interval (t), by using the Equation 4a.

$$X_t = \max\{0, R_t - D_t\} \tag{4a}$$
$$E_t^{\text{mine}} = a_t X_t \tag{4b}$$

Based on this surplus, the EMA is trained to choose an allocation fraction a_t (where $a_t \in [0,1]$) that efficiently determines how much of the excess energy is directed to BTC miners, by following the Equation 4b. If renewable production is lower than grid demand, ($R_t \leq D_t$), then X_t becomes zero, and the optimal decision is to set $a_t = 0$, meaning no energy is allocated to mining and all available energy is used to meet grid demand. Conversely, when R_t exceeds D_t and X_t is positive, the EMA is incentivized to decide on an appropriate a_t that balances two conflicting goals: minimizing wasted energy and ensuring that grid demand is met.

Towards meeting this operating specifics for the EMA in the centralized learning phase, we follow the reward function presented in Equation 5, which is designed to penalize both energy waste and any deficit in meeting grid demand. Since the EMA is primarily incorporated in the centralized learning of our framework, this reward function presents also the key updating mechanism, as evident in Equation 3.

$$r_t^{\text{EMA}} = -\alpha(X_t - a_t X_t)^2 - \beta(\max\{0, D_t - (R_t - a_t X_t)\})^2 \tag{5}$$

In Equation 5, the term $a_t X_t$ represents the energy allocated to mining, $X_t - a_t X_t$ represents the portion of surplus energy that remains unused, $\max\{0, D_t - (R_t - a_t X_t)\}$ quantifies the unmet grid demand, α and β are weight parameters that balance the EMA's objectives. Specifically, α penalizes the waste of surplus energy by incentivizing its allocation to BTC mining, thereby promoting resource efficiency. Conversely, β penalizes unmet grid demand, ensuring that grid reliability is prioritized. Adjusting these parameters allows the EMA to control the trade-off between maximizing surplus energy utilization for mining and fulfilling its grid support responsibilities. Beyond the EMA, we also design a comprehensive reward mechanism for the second type of agents incorporated in the energy trading schema, specifically the BTC miners. The rewarding of these agents is primarily designed to incentivize their revenue-generating aspect. Specifically, these agents act as uninterruptible loads that negotiate the price of the excess energy with the EMA. The maximum energy load that they could potentially receive is calculated in accordance with Equation 4b.

In the decentralized execution phase, these agents negotiate energy prices to maximize their revenue margins. To incentivize this behavior, we implement the reward mechanism defined in Equation 6, in which κ represents the revenue per unit of energy consumed in mining (influenced by factors such as Bitcoin's market price and mining efficiency); and c represents the cost per unit of energy used for mining, including operational and electricity costs. According to this mechanism, BTC miners aim to optimize their trading strategy to maximize revenue across each simulated episode. Notably, these agents are prohibited from consuming energy during peak demand periods to prevent grid disturbances. However, they have the capability to lock in energy quantities at predefined prices and receive the allocated energy at a later time.

$$r_t^{\text{miner}} = \kappa(a_t X_t) - c(a_t X_t) \quad (6)$$

V. Experimental set-up and evaluation

To provide some evidence regarding the applicability of the proposed MARL framework, we orchestrate a simplistic simulation based on the operating constraints and decision factors previously described. Regarding the constants associated with the simulation of the energy supply and demand setting, we have used the Equations 1 and 2, particularly by making the following assumptions:

- The energy production scheme is hybrid, by integrating fossil-fueled and RE sourced electric power. The maximum storage capacity of EMA is $M = 120.0$ MWh.
- It is assumed that the fossil-fueled stream is not subject to volatility and its stable production capacity is $T_{\text{stable}} = 80.0$ MWh.
- The RE production is modeled by using Equation 1 and assuming: $R_{t-1} = 80.0$ MWh and standard deviation of $\sigma_R = 5.0$ MWh.
- The grid demand is modeled by using Equation 2 and assuming: $D_{\text{avg}} = 130.0$ MWh, $A = 20.0$ MWh, $\omega = \frac{2\pi}{24}$ (24-hour periodic demand); and $\phi = \frac{\pi}{4}$.
- The pricing fluctuates within the execution of an episode, starting from a base energy price of $P_{\text{base}} = 0.15$ monetary units.

Beyond the simulation parameters, to implement the Q-learning framework, based on Equation 3, we have designed a discrete action space consisting of 11 possible energy allocation fractions. The learning process follows a learning rate of $\lambda = 0.1$, a discount factor of $\gamma = 0.95$, and an exploration rate of $\epsilon = 0.1$, following an epsilon-greedy policy.

The parameters related to reward mechanisms were set in a manner that primarily incentivize the avoidance of unmet demand, by setting the following values in constants of Equation 5: α=0.05 for excess RE waste and $\beta = 0.1$ for unmet demand. Regarding the BTC mining model, the cost per MWh for mining was set at $c_{\text{cost}} = 0.05$. BTC miners adopt a pre-scheduling strategy 30% of the time, allowing them to lock in energy at a 10% discount ($0.9P_t$). When pre-scheduling, miners allocate 80% of their planned energy usage in advance.

After running our MARL framework for one hundred (100) episodes, under the above settings, we obtained the results presented in Figure 2. Running the model for such a constrained number of episodes, while limiting its overall capacity to converge to global optima, enables us a first discussion at a proof-of-concept level. The first part of this figure depicts the average demand and the total energy produced by both the REP and fossil-fueled producers. Evidently, in almost all episodes it is observed that based on the operating settings, a significant amount of RE is exceeded.

This means that if the multi-agent framework converges to a curtailment-avoidance point, this excess energy should be traded to BTC miners rather than wasted. To assess whether this functionality holds true, we refer to the third part of Figure 2. As illustrated in this part, the profitability of BTC miners follows a progressively increasing trend, which inherently translates to the fact that the framework effectively learns a policy that enables energy allocation, ensuring that surplus energy is utilized in a financially viable manner rather than being wasted.

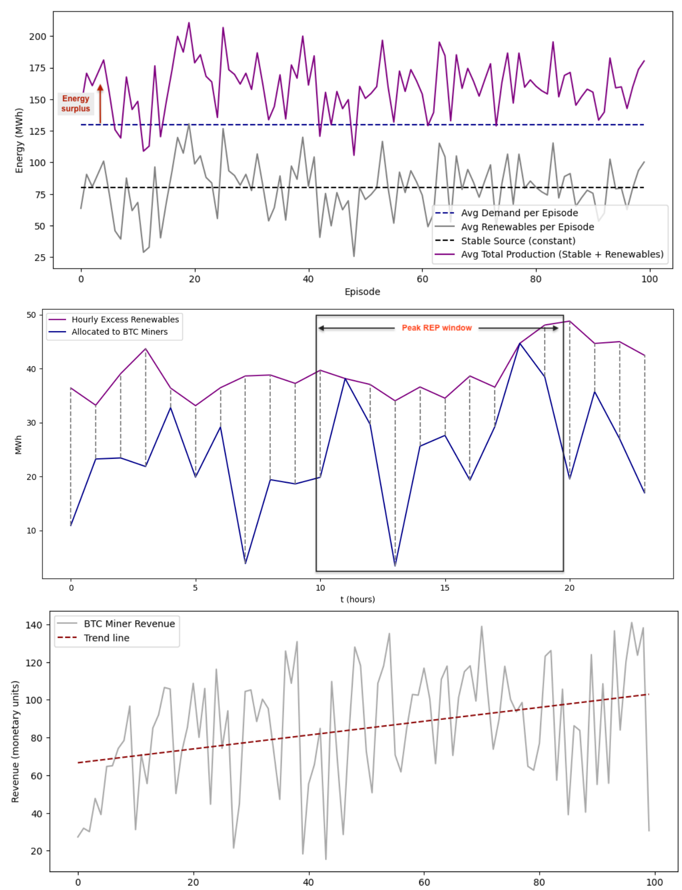

Fig. 2. An overall performance assessment of the proposed RL framework.

To better synthesize the agents' functionality, we also refer to the middle part of Figure 2, which depicts the amount of energy traded to BTC miners within a 24-hour time frame. These trades refer to the last episode of execution, capturing the maximum adaptability learned by the agents. Interestingly, these indications suggest that the framework successfully optimizes energy allocation by dynamically adjusting the traded surplus. The observed pattern reveals that the highest allocations occur within the peak REP window, signifying that the agents have learned to maximize energy utilization when renewable generation is abundant. Beyond this indication, it is also observed that the proportion of wasted surplus energy, represented by the dashed black lines, is significantly limited by at least 60%, a finding that highlights the curtailment mitigation and the enhance profitability achieved by the proposed framework.

VI. Discussion and conclusions

The integration of RES into the electrical grid poses substantial management challenges, especially because of the intermittent nature of RES, which often leads to energy curtailment, destabilization issues and energy loss. In this paper we have addressed this critical problem by proposing a MARL framework designed to dynamically balance grid load through the strategic integration of Bitcoin (BTC) mining operations. In particular, BTC miners serve as adjustable energy consumers, adapting their power usage based on real-time energy supply conditions. By employing BTC mining operations as flexible loads, the proposed system successfully managed excess renewable energy production, mitigating curtailment losses and stabilizing the grid.

Our findings have clearly underscored the potential of BTC miners to significantly reduce the volatility in energy grid while at the same time enhance their profitability. Unlike traditional scenario-based or deterministic models available in the literature, the MARL approach provides real-time adaptive capability, which is essential for responding to dynamic conditions in renewable energy generation and market prices. The results demonstrate that surplus or excess renewable energy, which would otherwise be wasted, can be effectively monetized through BTC mining activities. The experimental outcomes have underlined that agents could learn optimal allocation patterns, aligning BTC mining operations with periods of surplus production, therefore improving the economic viability and sustainability of energy systems incorporating RES.

Despite these promising results, several limitations should be kept. For instance, we conducted the simulations based on idealized assumptions and simplified parameters (e.g., fixed average grid demand and predefined pricing structures) and, therefore, the immediate applicability to real-world, complex and, arguably, volatile energy markets might be limited. Future studies should include real power grid data with dynamic pricing models to increase validation accuracy and market volatility response. The suggested MARL system is unclear about its scalability. The present technique excludes testing for large or geographically diversified grids with varying regional features and transmission restrictions. The examination of multi-regional scenarios would illustrate the framework's practical viability and resilience under actual grid settings. In addition, the processing demands of MARL hinder real-time responsiveness in practical applications. Future research should explore distributed RL techniques, including Federated Learning, to enhance training efficiency and system scalability in large-scale applications. Employing advanced economic modeling would improve understanding and practical use of MARL in energy management systems for BTC mining. Finally, exploring integration with existing infrastructures and standards like IEEE 1547 would also be a promising future research direction.

Acknowledgment

This work has been partly supported by the University of Piraeus Research Centre.

References

[1] International Energy Agency, "Renewables 2024," International Energy Agency, Paris, Tech. Rep., 2024. [Online]. Available: https://www.iea.org/reports/renewables-2024

[2] Global PV Brain Trust, "Solar pv curtailment: Drivers and solutions," Global Solar PV Brain Trust, Tech. Rep., May 2024. [Online]. Available: https://pv-thinktank.de/wp-content/uploads/2024/05/2.-Solar-PV-Brain-Trust_Second-Publication_Curtailment_May2024.pdf

[3] Global Energy Monitor, "Gem china wind & solar brief: July 2024," Global Energy Monitor, Tech. Rep., July 2024. [Online]. Available: https://globalenergymonitor.org/wp-content/uploads/2024/07/GEM-China-wind-solar-brief-July2024.pdf

[4] T. K. Dasaklis, I. T. Thomaidis, P. G. Giannopoulos, and G. T. Tsoulfas, "Rethinking bitcoin's energy use through sustainable digital business models and resources monetization: A multiple case study analysis," *Digital Business*, vol. 5, no. 1, p. 100114, 2025.

[5] J. Dorrell, A. Kancharla, and M. Ambrosia, "Green crypto mining: A quantitative analysis of the profitability of bitcoin mining using excess wind energy," *The Journal of Energy and Development*, vol. 48, no. 1, pp. 1–22, December 2023.

[6] A. Rosales, H. Millar, and A. Richardson, "Leveraging intra-provincial regulatory differences in a post-paris context: Cryptocurrency mining "reverse battery" strategy in atlantic canada," *The Extractive Industries and Society*, vol. 17, p. 101396, 2024.

[7] A. Bruno, P. Weber, and A. J. Yates, "Can bitcoin mining increase renewable electricity capacity?" *Resource and Energy Economics*, vol. 74, p. 101376, 2023.

[8] H. Niaz, J. J. Liu, and F. You, "Can texas mitigate wind and solar curtailments by leveraging bitcoin mining?" *Journal of Cleaner Production*, vol. 364, p. 132700, 2022.

[9] M. Hajiaghapour-Moghimi, E. Hajipour, K. Azimi Hosseini, M. Vakilian, and M. Lehtonen, "Cryptocurrency mining as a novel virtual energy storage system in islanded and grid-connected microgrids," *International Journal of Electrical Power & Energy Systems*, vol. 158, p. 109915, 2024.

[10] M. Hajiaghapour-Moghimi, E. Hajipour, K. Azimi Hosseini, M. Tavakkoli, S. Fattaheian-Dehkordi, M. Vakilian, and M. Lehtonen, "Hedging investments of grid-connected pv-bess in buildings using cryptocurrency mining: A case study in finland," *IEEE Access*, vol. PP, pp. 1–1, January 2023.

[11] A. Dzyuba, I. Solovyeva, and D. Konopelko, "Managing electricity costs in industrial mining and cryptocurrency data centers," *International Journal of Energy Economics and Policy*, vol. 13, no. 4, pp. 76–90, July 2023.

[12] R. Lotfi, S. Ghaboulian Zare, A. Gharehbaghi, S. Nazari, and G.-W. Weber, "Robust optimization for energy-aware cryptocurrency farm location with renewable energy," *Computers & Industrial Engineering*, vol. 177, p. 109009, 2023.

[13] A. Hakimi, M.-M. Pazuki, M. Salimi, and M. Amidpour, "Renewable energy and cryptocurrency: A dual approach to economic viability and environmental sustainability," *Heliyon*, vol. 10, p. e39765, November 2024.