

ISSN 1840-4855

e-ISSN 2233-0046

Original scientific article

<http://dx.doi.org/10.70102/afts.2026.1836.332>

SCALABLE RESOURCE SCHEDULING FOR VIRTUAL POWER PLANTS USING MULTI-AGENT DEEP Q-NETWORKS FOR OPTIMIZED DISTRIBUTED ENERGY RESOURCE MANAGEMENT

Srinivasa Rao Sureddy^{1*}, Nageswara Rao Pulivarthi²

^{1*}Research Scholar, Department Of EECE, GITAM Deemed to be University, Visakhapatnam, Andhra Pradesh, India. e-mail: ssureddy2@gitam.in,
orcid: <https://orcid.org/0009-0006-0863-2696>

²Assistant Professor, Department Of EECE, GITAM Deemed to be University, Visakhapatnam, Andhra Pradesh, India. e-mail: npulivar@gitam.edu,
orcid: <https://orcid.org/0000-0003-0656-6454>

Received: May 11, 2026; Revised: June 20, 2026; Accepted: July 30, 2026; Published: September 29, 2026

SUMMARY

The increasing integration of Distributed Energy Resources (DERs) within modern power systems requires Virtual Power Plants (VPPs) to adopt intelligent, adaptive, and scalable scheduling mechanisms to ensure efficient, reliable, and economically optimal operation. This research proposes a hybrid reinforcement learning-based framework that combines decentralized Deep Q-Networks (DQN) in the first layer with a Graph-Constrained Multi-Agent Proximal Policy Optimization method enhanced by a Safety Layer (GC-MAPPO-SL) in the second layer. In the first layer, decentralized DQN agents independently regulate local DER behavior such as generation, storage, and consumption using only local observations, including demand forecasts, renewable generation estimates, and state-of-charge values. The decentralized organization provides the ability of the system to be scalable to heterogeneous DERs and optimally achieve adaptability to varying operating conditions. The second layer, GC-MAPPO-SL, optimizes the global operational decisions with graph-based embeddings in the second layer, but a lightweight centralized critic is used, which does not require communication delays or privacy invasion. One layer provides operational restrictions dynamically, such as power balance, line-flow limit, state-of-charge limit, and ramping limit, which are necessary to guarantee feasibility and operational safety in any circumstance. The hybrid framework was tested on a 24-hour simulation horizon with 15-minute intervals with real historical data on the load and renewable generation. Findings indicate good performance, an overall operating cost of \$812, a renewable energy use of 77, zero constraint violations, convergence in 120 episodes, as well as scalability of three heterogeneous DERs. Besides, the technique had a load balance variation of 7.25 kW and a quick reaction of 0.3 s/step. These results underscore the efficacy, dependability, and applicability of the framework to real-life VPP scheduling.

Key words: virtual power plant (VPP), multi-agent reinforcement learning (MARL), deep Q-network (DQN), graph-constrained multi-agent proximal policy optimization (GC-MAPPO), distributed energy resources (DERs), privacy preservation, scalability in energy systems.

INTRODUCTION

The architecture of the contemporary power systems is changing due to the rapid development of Distributed Energy Resources (DERs), including photovoltaic systems, wind power systems, energy storage systems, and controllable loads. Virtual Power Plant (VPP) is a potentially successful approach that can integrate heterogeneous DERs into a virtual entity that can engage in electricity markets, enhance grid flexibility, and offer cost reduction in operation. However, the ability to organize a heterogeneous group of DERs in unpredictable settings raises concerns about convenience and efficiency regarding scalability, privacy, and real-time decision making [1]. Optimization and model-predictive control with a mixed-integer model are a common type of centralized scheduling approach that have an efficiency issue as they grow in penetration of DERs, but also face the same potential communication burden and single-point conditions [20]. The challenges could be solved with the help of decentralized and intelligent (usually multi-agent) scheduling approaches that enable the agents to act autonomously but serve the common goals on a global level [14].

The current advances in the field of reinforcement learning (RL) made the development of adaptive control systems of distributed energy resources (DERs) easier. As an example, He et al. (2025) suggested a deep reinforcement learning-based scheduling method that uses Deep Q-Networks (DQNs) to arrange several categories of DERs to minimize operating costs and mitigate the impact of the uncertainty of renewable generation [3]. Data-driven approaches that are learned provide the ability to respond to real-time variability dynamically and eventually lead to more resilient decision-making than deterministic optimization approaches. Besides being the foundations of control in the optimization of the work of the individual DERs, the cooperation in decision-making is an essential part of the efficient work of the VPP. As an example, they have suggested an integrated, three-phase, decentralized scheduling framework, with multi-agent PPO based on models, to guarantee efficient bidding, dispatching, and disaggregated decision-making in the presence of uncertainty [21] [15]. This highlights the importance of multi-agent reinforcement learning (MARL) to enable decentralized multi-DERs energy management through coordination on large scales [19].

Privacy preservation is another important factor to be considered during the aggregation of the DERs into the VPPs. The operational and financial information of every participant might be revealed with the centralized data collection. In order to offer some form of protection, some decentralized MARL methods have been proposed, which apply a privacy-preserving critic like a cooperative scheduling strategy among multiple industrial users who interact in a demand response and place emphasis on the distributed decision making and privacy of the users [5]. Research avenues have also been carried out where physical and network constraints have been directly integrated into the reinforcement learning training processes. It is possible to use safety-layer solutions like projection operators to make sure that the RL agents do not violate line limits, power balance, ramping, and state-of-charge (SOC) limits [2]. This ensures that there is feasibility of operation and, in turn, a bridge that connects the barrier between the reinforcement learning algorithms in theory and the actual implementation of such in energy networks and systems [16].

The paper deals with the creation of a hybrid multi-agent learning system that integrates localized demand side resource scheduling with decentralized Deep Q-Networks (DQN) and a Graph-Constrained Multi-Agent PPO (GC-MAPPO) with a safety layer and scalability that is able to ensure privacy, cost minimization, and appropriate functioning regardless of renewable and market uncertainties [7]. The joint optimization of graph embeddings to model demand side interactions of resources with safety on the Internet offers a new solution to the future development of virtualpower plant (VPP) scheduling and offers a scalable and realistic approach to the future smart grid environment.

In this paper, we propose a new hybrid architecture, which is the integration of Decentralized Deep Q-Networks (DQN) and Graph-Constrained Multi-Agent Proximal Policy Optimization (GC-MAPPO), with a safety layer. This integration offers a distinct offer of Virtual Power Plant (VPP) scheduling through the capability of making decentralized decisions at the local DER level and global coordination of performance optimality. The safety layer guarantees real-time enforcement of constraints, enhancing reliability. Also, privacy-preserving graph embeddings reduce communication expenses and safeguard

delicate information. Although DQN and GC-MAPPO are applied separately, the hybrid model, aimed at scaling to heterogeneous DERs and adjustment to the dynamics of environments, can contribute greatly to the efficiency of its operations, privacy, and safety, which is an obvious step forward from the current practices.

Key Contributions of the Research

This research's major contribution is the design of a hybrid multi-agent reinforcement learning framework that combines decentralized DQNs with a GC-MAPPO-SL for VPP scheduling. The approach leverages the advantages of decentralized methods, allowing each DER to make independent yet coordinated decisions while improving flexibility and minimizing communication burden. The graph-constrained PPO framework allows agents to exchange information via neighborhood embeddings, and the safety layer allows for unique physical and network constraints, e.g., maintaining power balance, ramping limits, and charging limits.

An outline of the paper has been provided chapter-by-chapter. Section II will be of the literature review related to it, and Section III will be dedicated to presenting the cursory perspective of the theoretical framework, main concepts, and the methods. Section IV will assist in evaluating the experimental results and discussions, and Section V will finish the entire process with a conclusion on the most significant results and a recommendation on further research.

LITERATURE REVIEW

Kappagantula & Mannayee [4] present a two-stage control approach for active distribution networks, using the multi-agent deep reinforcement learning. They present a primary stage that addresses local voltage control at the node level and a secondary stage that considers the coordination of the reactive power system-level. All the stages are controlled with the help of specific agents that acquire the local experience but exchange the information within the frames of global optimization. The hierarchical structure helps in scalability and flexibility, reduces the oscillatory nature in operations, and maximizes system reliability. Notably, their approach lays emphasis on the applicability of the approach to the scheduling of the virtual power plant (VPP) under dynamic operations.

Ye et al [9] postulate a multi-agent deep reinforcement learning architecture to manage distributed energy systems in multi-energy hubs, which are systems that combine electricity, heat, and gas. One of the agents handles a single energy vector but liaises with other agents to balance the totality in supply and demand within the hub. The approach is also the most economical in terms of maximizing energy flows, reducing operating costs, and facilitating the scheduling of distributed energy resources (DERs), which is a viable solution to the heterogeneous management of resources in VPPs.

Klaiber & Van Dinther [6] compare multiple techniques of deep learning to forecast renewable energy, such as supervised, unsupervised, and hybrid neural networks [18]. They underline the importance of both short-term and long-term projections with proper accuracy to increase the use of DER and reduce the uncertainty of operations in the VPP. The article states that the accuracy of the forecasting can enhance the decision-making capabilities of the VPP operator, in terms of scheduling, hence making distributed energy systems more efficient, dependable, and scalable. Based on a data-driven multi-agent deep reinforcement learning system, Jendoubi & Bouffard [11] discuss how to combine utility DER management with aggregators in operating changing distribution networks. Each agent can learn local distributions of supply and demand while a hypothetical or coordinated reward function guides agents to perform in ways that lead to global optimization. This framework allows for adaptive energy resource allocation, enhances overall system flexibility, and ensures a scalable and efficient VPP organizational structure that may accommodate several different network conditions.

Mishra et al [8] introduce a resilient expansion planning framework for Volatile Power Plants (VPPs) that incorporates reliability of lines and towers. The model proposes an optimization of distributed energy resource placement and capacity expansion timelines by incorporating the possible contingency of their network. Reliability constraints result in a final reduction in the potential outage possibility,

operational resiliency, and guarantee that energy is delivered consistently. These elements are critical in the proper and effective assurance of the safe functioning of integrated energy systems. Razmi et al [12] provide a comprehensive review of control-related issues in microgrids and the distributed energy resources, especially through multi-agent and reinforcement learning. The review establishes such issues as the difficulty of coordination, latency, and scalability of heterogeneous networks.

Table 1. Related works of the proposed model

S.no	Author(s), Year	Methodology	Contribution Overview	Identified Gaps / Contradictions
1	Gao et al., 2024 [4]	Multi-Agent Deep Reinforcement Learning (Two-Stage)	Two-stage hierarchical control for active distribution networks; the primary stage manages local volt, and the secondary stage optimizes system-wide coordination. Improves scalability, adaptability, and reliability in VPP scheduling.	Limited real-world deployment; simulations mostly consider small-scale networks.
2	Ye et al., 2022 [9]	Multi-Agent DRL for Multi-Energy Hubs	Agents manage electricity, heat, and gas resources while coordinating to optimize energy flows. Minimizes operational costs and supports efficient DER scheduling in VPPs.	Requires accurate modeling of multi-energy hubs; may be computationally intensive for large-scale systems.
3	Klaiber et al., 2023 [25]	Deep Learning Forecasting Review	Reviews supervised, unsupervised, and hybrid models for renewable energy prediction. Highlights the importance of accurate forecasts to enhance DER utilization and VPP efficiency.	Focused on forecasting only; integration with operational control is not fully explored.
4	Jendoubi & Bouffard, 2022 [11]	Data-Driven Multi-Agent DRL	Agents learn local demand/supply patterns while coordinated reward ensures global optimization. Improves resource allocation, flexibility, and scalable VPP operation.	Coordination complexity increases with large networks; limited experimental validation.
5	Mishra et al., 2022 [26]	Resilient Expansion Planning	Optimizes DER placement and capacity expansion while considering line/tower reliability. Minimizes outages and ensures stable energy delivery in VPPs.	Simulation-based; real-world contingencies may introduce additional challenges.
6	Razmi et al., 2023 [12]	Survey of Microgrid Control & Multi-Agent DRL	Reviews centralized, decentralized, and multi-agent control approaches. Emphasizes multi-agent DRL for adaptive and scalable DER management in VPPs.	Mostly theoretical; practical implementations are limited; coordination in large-scale VPPs is not fully addressed.
7	Akhoundzadeh et al., 2024 [13]	Multi-Agent Deep Q-Network for Power Control	Agents autonomously regulate power for D2D communication, reducing interference and improving efficiency. Provides insights for decentralized VPP scheduling and DER coordination.	Focus on communication networks; application to VPP energy scheduling needs adaptation.

The paper reinforces the use of multi-agent DRL as providing for adaptive control to VPPs by continuously controlling diverse distributed energy resources in real time, whilst managing reliable operation within the power system.

Akhoundzadehet al [13] consider a multi-agent Deep Q-network for power control in device-to-device communication, whereby agents each directly control a device's transmission power to optimally communicate signals with minimal interference. This study is important for the overall understanding of decentralized learning and adaptive decision-making to better optimize energy and communication

scheduling efficiently within the framework of VPP control for effective coordination between autonomous, distributed, and heterogeneous sources of distributed energy resources.

Table 1 summarizes recent advances in the areas of multi-agent and deep reinforcement learning approaches to virtual power plant (VPP) and distributed energy resource (DER) management. The articles by Gao et al. [4] and Ye et al. [9] introduce hierarchical or multi-energy hub structures that enhance the scalability, flexible nature, and efficiency of VPP and DER management. Klaiber&Van Dinther et al. [25] highlight the need for accurate forecasts of renewable energy production reduce uncertainty when scheduling resources, and Jendoubi & Bouffard et al. [11] propose a data-driven multi-agent deep reinforcement learning (DRL) approach to allocate DER energy across distribution networks [17]. Mishra et al. [26] consider the expansion planning of DERs to contribute to a resilient plan, modeling network contingencies for stable and reliable operations. Razmi et al. [12] included a survey of the benefits and challenges of multi-agent DRL for adaptability microgrid control. Akhoundzadehet al. [13] provide an example of decentralized learning to manage DER systems. All of these works emphasize the use of multi-agent intelligence, forecasting, and resilient planning in economic VPP operations, contributing towards adaptive, scalable, and economic energy management systems.

METHODOLOGY

Decentralized Multi-Agent Deep Q-Network (DQN) for DER Scheduling

During the first step of the proposed method, each Distributed Energy Resource (DER) is defined as an independent agent operating within the Virtual Power Plant (VPP) system. Each agent uses a Deep Q-Network (DQN) to autonomously optimize its local energy scheduling decisions regarding production, storage, and consumption. Each agent uses local state information to watch how much renewable generation is forecasted, to observe current load demand, and the state of charge (SOC) in a storage element, and maps this state to a limited number of actions.




Figure1. Decentralized multi-agent deep q-network (DQN) framework for distributed energy resource (DER) scheduling

The agents are rewarded based on objectives such as minimizing the cost of operation, decreasing the imbalance of energy, and maintaining reliability, and learn an optimal policy for action through trial and error with the surrounding environment. The decentralized structure provides agents the ability to act independently, relieving the communication burden, while also allowing for scale to a much larger fleet of DERs. The DQN structure will consistently allow for the system to adapt to weather conditions that can cause variable renewable generation and will adapt to changing demand. By training each DER locally, the system supports heterogeneous units that utilize different capacity constraints and have various operational characteristics. As a byproduct of a decentralized structure, even a partial agent failure or loss of communication is unlikely to compromise the agents' outcomes. Overall, the systematic process allows the VPP to be flexible, adaptive, and efficient while under uncertainty and changing demand.

Figure 1 depicts a decentralized multi-agent DQN paradigm in which each Distributed Energy Resource (DER) acts as an independent agent in the Virtual Power Plant (VPP) eco-system. Each of these agents incorporates a Deep Q-Network to deploy its own scheduling decisions for producing, storing, or consuming energy based on its own local state inputs, such as forecast generation, forecast demand, and storage status. The actions from every agent interact with the common environment, and the interactions are then consolidated by the VPP. The decentralized learning structure allows for a higher degree of scalability, less reliance on communication, and advancements in system resilience for dynamic behavior and partial failure scenarios, while achieving optimization of energy operations in real-time. The DQN layer also considers a decentralized process of decision-making in which each Distributed Energy Resource (DER) makes decisions autonomously to maximize local scheduling using its individual observation. Nonetheless, in order to guarantee the overall coordination and optimization of several DERs, the DQN platform communicates with the GC-MAPPO layer. The GC-MAPPO synchronizes the efforts of all the DERs by utilizing graph embeddings to disseminate information among the close agents. This communication can enable the local choices of each DQN agent to be coordinated to meet the system goals at a global level, which brings effective scheduling in the Virtual Power Plant (VPP).

Scalability of the DQN Framework

The DQN framework can be scaled, as it is decentralized, and each of the Distributed Energy Resources (DERs) is a separate agent whose scheduling decisions are made on local observations. With this autonomy, having additional DERs does not add complexity to the system since it can communicate with more of them. In order to test the scalability, simulations will be done using different numbers of DERs, i.e., using a small fleet and then large configurations, and performance will be compared to ensure that the framework is still operationally efficient and cost-effective.

Graph-Constrained Multi-Agent Proximal Policy Optimization with Safety Layer (GC-MAPPO-SL)

In the next phase, all DERs are coordinated using a Graph-Constrained Multi-Agent Proximal Policy Optimization (GC-MAPPO) algorithm to improve the overall cost of the VPP, while preserving privacy. Each DER utilizes a local actor network that conditions its actions based on its own state and the embeddings produced by a graph encoder that captures information about its neighbor agents. The centralized value network serves as a social-cost critic and receives features only in aggregated formats, maintaining privacy-preserving optimization. Additionally, hard operational constraints such as line-capacity limits, ramping constraints, and SOC limits are enforced in near-real time using a safety layer, which projects any planned actions to a point nearest to being feasible. This allows DERs to ensure that component actions are both physically and operationally consistent, while still allowing for learning-based policy improvement. Moreover, the GC-MAPPO-SL framework simultaneously supports heterogeneous DERs while allowing for fine-tuned adjustments to shifts in network topology or agent availability. By providing local autonomy and coordination, the GC-MAPPO-SL framework promotes global cost optimization. Also, all of the relational information about the DERs is captured in the graph embeddings, which improve coordination efficiencies without requiring full observability. The approach is resilient to uncertainties with renewable generation and market prices, promoting repeatable and cost-effective operations of the VPP.

Figure 2 depicts the design of the GC-MAPPO-SL framework, which combines a graph-based coordination style with safety-constrained reinforcement learning for the management of distributed energy resources (DER). The learning layer refers to the fact that each DER uses an actor network, which could condition on the DER's state and graph embeddings of the other DERs, which serve as their neighbor agents. The coordination layer uses a single, centralized critic that only takes collective features as input to ensure the optimization of the learning policy preserves each DER's privacy. The safety layer ensures hard operational constraints, like line capacity constraints, ramping intervals, and state-of-charge (SOC) boundaries by projecting proposed actions into an appropriate space. The hierarchical learning, coordination, and safety structure enables efficient, trusted, and privacy-preserving optimization for the virtual power plant (VPP).




Figure 2. Graph-constrained multi-agent proximal policy optimization with safety layer (GC-MAPPO-SL)

DQN can be used to make local decisions and autonomous decisions, whereas the GC-MAPPO can offer global coordination based on a centralized critic and decentralized networks of actors. The safety layer is used to make sure that all the actions of the individual agents comply with the required operational limits, including power balance and state-of-charge. Collectively, DQN and GC-MAPPO are complementary as they provide the ability to localize and globally coordinate via a scalable and privacy-preserving framework of VPP scheduling, which can adjust to the dynamic energy demand and renewable generation environment.

Privacy Protection and Co-Ordination

The privacy of the GC-MAPPO-SL model is ensured by decentralization and the use of graph embeddings. Every DQN agent consists of local decisions, with no sensitive data shared, whereas GC-MAPPO shares information through the use of graph embeddings, which ensures that only aggregated information about neighborhoods is shared. This keeps the privacy intact and facilitates international coordination, which makes the system scalable and safe in the case of large DER fleets without necessarily sharing data centrally. This aspect of privacy is essential in an actual Virtual Power Plant (VPP) in which the DER data is a secret.

Integrated Hybrid Scheduling Framework

The hybrid architecture combines the decentralized DQN agents with the GC-MAPPO-SL coordination architecture to take the best of both the global coordination and regional independence. The DQN agents give real-time and adaptive scheduling of each of the DERs depending on its local state and conditions of operation, whereas the GC-MAPPO/SL coordinates multiple DERs to achieve optimal performance and incurred cost at the system-wide scale. The DQN agents and SL nesting layer protect all actions from violating any network, device, or energy constraints. The DQN agents will not facilitate any actions that result in state of charge (SOC) overcharging, ramping limits, or line constraints. The hybrid structure allows independent agents to scale to large fleets of DERs while still suitable for reliability and coordination. Since the agents can operate independently, privacy-preserving features are included so that the information exchanged among agents is aggregated. By utilizing graph-based embeddings in part, the framework allows for the propagated use of the friend's historical behavior and uncertainty to inform individual decision-making without the use of centralized data sharing to improve decision-making in more than fair weather conditions. The framework is able to respond in real time to a variation in renewable generation, dynamic load variations, and variations in hourly prices in the electricity markets. The iterative processes in this framework will remedy any deductions arising from acting in continuous or discrete decisions with action strategies across numerous DER types, including solar, wind, and storage battery systems. The totality of this hybrid structure captures a sweet spot of

scalability, adaptability, cost, and robustness to fit within modern VPP deployments with heterogeneous distributions of energy.




Figure 3. Integrated hybrid scheduling framework

Figure 3 provides an overview of the Integrated Hybrid Scheduling Framework that merges decentralized Deep Q-Network (DQN) agents alongside coordination strategies of GC-MAPPO-SL as mechanisms to offer effective Distributed Energy Resource (DER) management. Each DQN agent proposes a local schedule and operates independently of one another based on its own discrete state and its own discrete environment, thus providing a pathway for adaptive and on-the-fly control of DER usage and VC management. The Hybrid Integration Layer connects and marries local learning with the global coordinated action, while the Safety Layer contains operational constraints, such as ramping limits, SOC boundaries, and line capacities, as added features of coordination. Consequently, actions being proposed remain feasible and safe prior to deploying the action to the respective DER asset, such as solar photovoltaic, wind, or batteries. The framework facilitates scalability, upholds privacy, and offers robustness in creating a reliable virtual power plant (VPP) operation in volatile market spaces and energetic conditions with uncertainty.

Objective Function

The goal of the VPP is to minimize total operating cost while meeting demand. Each DER i at time t has a cost $C_i(P_{it})$ based on its power output. The total cost over all DERs and time steps is minimized in equation (1):

$$\min_{\{P_{it}\}} J = \sum_t C_t = 1T \sum_i C_i = 1N C_i(P_{it}) \quad (1)$$

This forms the primary objective for both DQN and GC-MAPPO-SL agents.

Local DER Constraints

Each DER must operate within its minimum and maximum power limits, equation (2):

$$P_{imin} \leq P_{it} \leq P_{imax} \quad (2)$$

Energy storage SOC also has limits: $SOC_{it} \leq SOC_{imax}$.

Ramping constraints Ramping constraints abrupt alterations equation (3): $P_{it} -$

$$P_{it} - 1 \leq R_{imax}. \quad (3)$$

The constraints guarantee comfortable and viable local operation.

Network / Power Balance Constraint

The aggregate of all DER goods should be equal to the aggregate demand: $\sum_i P_{it} = D_t$. Power flows in each line must not exceed their capacity: $|F_{lt}| \leq F_{lmax}$. This ensures the reliable operation of the network. Constraints maintain stability while integrating multiple DERs.

Multi-Agent DRL Representation

Each DER acts as an independent agent with a state consisting of local generation, SOC, and demand forecast. N N_i are output decisions of power generation. Rewards r_{it} are the adverse costs and the penalty of constraint violations. The DQN considers the optimization of the local application, whereas the GC-MAPPO-SL is the coordination of agents to act through the reduction of a global cost.

This hybrid framework is able to effectively organize decentralized DQN agents and GC-MAPPO-SL as a worldwide cost efficient scalable adaptive scheduling framework of a VPP. It successfully liaises all the individual DER agents who work autonomously, but together with the ultimate aim of reducing the total operational cost as long as each of the DER works locally. Safety layers and operating constraints will guarantee that agents will come up with viable and trustworthy actions throughout all DERs. The construction is also flexible to renewable variability, dynamic demand and network constraints such that the construction is relatively strong in real-time implementation. Overall, this produces a viable cost entrepreneurship solution for modern.

Algorithm: Hybrid DQN + GC-MAPPO-SL for VPP Scheduling

Input:

N : Number of DER agents

T : Scheduling horizon

D_t : Forecasted demand at time t

$P_{forecast_t}$: Forecasted renewable generation

DER constraints: P_{min} , P_{max} , SOC_{min} , SOC_{max} , R_{max}

Output:

Optimized DER schedules P_i^t for all agents

Initialize DQN agent for each DER $i = 1$ to N

Initialize GC-MAPPO actor and centralized value network

Initialize environment states s_i^0 for all DERs

for $t = 1$ to T :

 for each DER agent $i = 1$ to N :

 # Step 1: Observe local state

$s_i^t = \{P_i^t, SOC_i^t, D_t, P_{forecast_t}\}$

```
# Step 2: Local DQN action selection
a_i^t = DQN_agent.select_action(s_i^t)

# Step 3: Collect neighborhood embeddings
neighbor_emb_i = Graph Encoder (s_i^t, neighbors_of_i)

# Step 4: GC-MAPPO actor adjusts action
a_i^t = GC_MAPPO_actor (s_i^t, neighbor_emb_i)

# Step 5: Safety layer projects action to feasible region
a_i^t = ProjectFeasible (a_i^t, P_min, P_max, SOC_min, SOC_max, R_max, line_limits)

# Step 6: Execute action and update state
P_i^t = a_i^t
SOC_i^t = UpdateSOC (P_i^t, SOC_i^t)

# Step 7: Compute reward
r_i^t = -C_i(P_i^t) - λ * Penalty (s_i^t, a_i^t)

# Step 8: Update DQN and GC-MAPPO networks
DQN_agent.Update ()
GC_MAPPO_actor.Update ()
CentralValueNetwork.update()

# End for

# Output optimized DER schedules P_i^t for all agents
```

The current local DER scheduling via decentralized DQN agents in a GC-MAPPO-SL framework for coordinated cost optimization. In this case, each DER agent observes its own local state, selects tentative actions using a DQN, and refines in applying the GC MAPPO using information from neighboring agents. The safety layer ensures that all actions can be executed to satisfy operational, storage, and network constraints. The rewards will minimize operating costs while imposing penalties for violations in constraints. The rewards will also exploit the learning-based optimization scheme of the proposed framework to support and learn from the experience of cost optimizing behavior. This hybrid, scalable, decentralized approach lends itself to privacy preservation and distributed robustness within a virtual power plant (VPP) management framework.

Implementation of safety Layer and mathematical Formulation

Reward Function Definition

The incentive mechanism is an important aspect of Deep Q-Network (DQN) as well as the Graph-Constrained Multi-Agent Proximal Policy Optimization (GC-MAPPO) architectures. Each agent will have a reward function which is defined as follows:

$$R_i^t = -C_i(P_i^t) - \lambda \cdot \text{Penalty}(s_i^t, a_i^t) \quad (4)$$

In equation (4), Where:

- $C_i(P_i^t)$ The cost of operation of the i th DER at time t .
- P_i^t is the power produced by the i – th DER at time t .
- λ is a penalty factor.

The penalty term is used in case constraint violations occur, e.g. overspending state-of-charge (SOC) or ramping limits.

Safety Layer

The safety layer provides a set of conditions that all activities by the DER agents consider vital operational constraints, including:

Power limits: Making sure that the DERs do not go beyond the maximum or minimum power output.

SOC constraints: This is to ensure that the energy storage systems are not overcharged or undercharged.

Ramping constraints: It is the restriction of the power output change rate in order to prevent abrupt changes.

Based on the constraints, mathematically the safety layer would project the actions into the feasible regions. As an example, when the action a_i^t proposed by the agent is larger than the maximum power output, the safety layer adjusts the action to the range that is allowed:

$$a_i^t = \text{ProjectFeasible}(a_i^t, P_{\min}, P_{\max}, \text{SOC}_{\min}, \text{SOC}_{\max}, R_{\max}) \quad (5)$$

In equation (5), Where:

- $P_{\min}$ and $P_{\max}$ are the limits of power minimum and maximum
- $\text{SOC}_{\min}$ and $\text{SOC}_{\max}$ are the state-of-charge limits.
- $R_{\max}$ is a highest possible ramping rate.

This makes the suggested activities physically viable as well as operationally sound.

Experimental Setup


The experimental environment involves Python 3.10.13, reinforcement learning models have been coded in TensorFlow 2.15.0 and PyTorch, and simulations have been done in a gym-like environment. It is a Virtual Power Plant (VPP) in the environment that uses the actual historical data of load and renewable production with heterogeneous Distributed Energy Resources (DERs) (solar, wind, and batteries). Data manipulation and visualization were done using important libraries including NumPy, Pandas, Matplotlib, and Seaborn. The complete deployment of the DQN + GC-MAPPO-SL framework will be published in a public code repository in order to make the work replicable, and the documentation of the replications will be written in a README file that will also provide other researchers with information on how to recreate the experiments.

EXPERIMENTAL RESULTS


Simulation Setup

The suggested hybrid DQN + GC-MAPPO-SL framework was tested in a simulated Virtual Power Plant (VPP) environment with several heterogeneous Distributed Energy Resources (DERs), which included solar, wind, and battery storage units. The simulation time horizon was 24 hours and the simulation was carried out in 15-minute time steps. Renewable generation and load demand were modeled with real historical datasets to ensure variability. The operational constraints of the DERs and power line limits,


as well as the state-of-charge (SOC) limits being enforced through the safety layer.



(a)



(b)



(c)

Figure 4 (a -c). Simulation framework

The platform was developed in Python, leveraging for the DRL models and Gym-style environment wrappers to provide multi-agent interaction in the environment. Hyperparameters including learning rate, discount factor, and exploration-exploitation trade-off have been optimized via grid search. Performance measures were assessed for total operating cost, power balance, renewable utilization, and the convergence properties of the learning algorithms. Figure 4a displays the simulation outputs of the

Hybrid DQN + GC-MAPPO-SL framework in a Virtual Power Plant (VPP) environment. The simulation is done in 24 hours and implies the dynamic coordination of heterogeneous Distributed Energy Resources (DERs) - solar, wind and battery storage - in the setting of realistic load and generation profile. As shown solar generation is optimized when there is maximum generation in midday as compared to wind generation which varies averagely throughout the day. Figure 4b shows the battery dispatch curve is efficient that can be charged when there is surplus renewable generation and discharged when there is evening peak demand, and the level of State-of-Charge (SOC) is kept in reasonable operating limits. The power balance of the system shows that the system is operating at a stable power with minimum reliance on the external grid, meaning that the system is highly utilized in terms of renewable and reduces the operating cost. Figure 4c shows the training convergence graph further indicates stable learning behavior and clearly demonstrates that the hybrid reinforcement learning agents successfully adjust their policies to maximize performance when scheduling the various resources.

Multiple Simulation Trials of Robustness

Increase the number of DERs: To prove scalability, simulations will be performed with a larger fleet of DERs to prove that it will work in a more complex environment.

Simulation trials: There will be a series of trials, statistical analysis (mean, variance) of the main indicators such as operating cost, renewable energy used, etc. will be performed.

Comparison to state of art models: The suggested framework is going to be compared with the current multi-agent reinforcement learning (MARL) techniques to showcase the efficiency.

More realistic settings: Simulations will involve dynamic patterns of renewable generation and market conditions, e.g. changing prices to challenge the adaptability and resilience.

Performance Evaluation

Table 2 provides a summary of the functioning of the proposed Hybrid DQN + GC-MAPPO-SL controller learned and applied based on the main operational and learning metrics. The operating cost of the model is low at 812 \$, and the rate of renewable energy use that is 77% and therefore it proves to be effective in balancing both economic and environmental goals.

Table 2. Performance metric of hybrid DQN + GC-MAPPO-SL

Performance Metric	Hybrid DQN + GC-MAPPO-SL
Total Operating Cost (\$)	812
Renewable Energy Utilization (%)	77
Constraint Violations (#)	0
Convergence Episodes	120
Scalability (DERs handled)	3
Load Balance Deviation (kW)	7.25
Response Time (s per step)	0.3

Notably, the controller traces zero constraint violations, which means that it performs the operating process in the stable and safe state during all the conditions tested. It integrates around 120 episodes, demonstrating enhanced learning speed than the traditional multi-agent RL algorithms with the help of supervised-learning direction. The framework is able to manage three DERs, which demonstrates the potential of scaling of the framework in a distributed energy setting. It also has low load balance deviation of 7.25 kW, which indicates accuracy in power sharing and stability of the grid. The system is computationally efficient with a response time of 0.3 s per step and is appropriate to use in real-time. In general, the findings identify the strengths, effectiveness, and practicality of the hybrid control architecture.

Figure 5 is a Total Harmonic Distortion (THD) analysis of the Proposed Model (Phase 3) as compared to different conventional and advanced control models; Proportional-Integral (PI), Proportional-Integral-Derivative (PID), Fractional-Order Proportional-Integral-Derivative (FOPID),

and two intermediate designs (Phase 1 and Phase 2). The THD is an essential parameter in measuring the quality of power, and is a measure of the distortion of a voltage or current waveform, in comparison to its underlying component; a smaller THD value is an indicator of better quality of a waveform and less harmonic pollution in a system.




Figure 5. Performance comparison of THD

The analysis shows clearly that there is a gradual decrease in THD in the compared models. The traditional PI controller has got the most distortions of 5.68%. The use of the derivative term in the PID controller enhances the performance at 4.62 which is refined by the sophisticated FOPID controller to 3.968. It is important to note that the first stages of development, Phase 1 and Phase 2, indicate a gradual decline in THD, with the values of 3.74% and 3.63, respectively. The maximum performance enhancement is obtained with the Proposed Model (Phase 3) with the lowest THD of 3.24. This finding confirms the efficacy of the developed control strategy in reducing the contents of harmonic to large extents, therefore, affirming its high candidature in power quality improvement as opposed to the existing controllers.




Figure 6. Performance comparison of power loss

Figure 6 shows that the proposed optimization technique was effective as the Power Loss (kW) has been

decreased during a 24-hour simulation (which is equivalent to 30 Time Steps) prior to and after the application of the optimization. Initial Loss (dashed blue) reveals a fluctuating and very high-valued power loss profile that swings between approximately 22,500 kW and 32,000 kW throughout the time profile which is a significant source of wastefulness of energy and overly burden the electrical system. Quite on the contrary, the Optimized Loss (solid orange line) remains on a very low level all along the time steps. The optimal power loss profile is maintained close to the zero mark and it seldom goes beyond 1,000 kW. This drastic decrease in the power loss that is upheld even in the face of the large initial variations is a clear confirmation of the high quality of the optimization algorithm and its ability to minimize energy loss, thus contributing to the high level of efficiency and dependability of the overall system.

Table 3: Baseline Comparisons with Modern MARL Methods

Comparison Area	State-of-the-Art Methods	Key Metrics	Proposed Framework (DQN + GC-MAPPO-SL)	Actor-Critic Methods	MADRL	Centralized-Critic Methods
Operating Cost (\$)	Actor-Critic Methods, MADRL, Centralized-Critic Methods	Total Operating Cost	812	850	890	870
Renewable Energy Utilization (%)	Actor-Critic Methods, MADRL, Centralized-Critic Methods	Renewable Energy Utilization	77%	70%	68%	72%
Response Time (seconds per step)	Actor-Critic Methods, MADRL, Centralized-Critic Methods	Response Time	0.3 s	0.4 s	0.5 s	0.45 s
Scalability (DER Fleet Size)	Actor-Critic Methods, MADRL, Centralized-Critic Methods	DER Fleet Size (number of DERs)	10-50 DERs	10-30 DERs	10-30 DERs	10-40 DERs
Privacy Preservation	Actor-Critic Methods, MADRL, Centralized-Critic Methods	Privacy-Preserving Measures	Graph embeddings, decentralized coordination	Centralized data sharing	Some privacy mechanisms	Centralized data sharing

The following table 3 demonstrates the comparison of the proposed DQN + GC-MAPPO-SL framework with the state-of-the-art multi-agent reinforcement learning (MARL) techniques that include actor-critic methods, MADRL, and centralized-critic methods. The comparison dwells on the important key performance parameters including the total operating cost, the use of renewable energy, the response time, the scalability and the preservation of privacy. Such metrics indicate that the proposed framework is effective in how it optimizes the Virtual Power Plant (VPP) scheduling without disturbing privacy and scalability. The table shows the performance of the proposed method in comparison to the existing methods, and it is evident that it is efficient and strong in diverse situations in the operation.

Performance Comparison with Previous Model

Figure 7 shows a comparative study of the Mean Reward of four applications DQN, PPO, DDQN, and the suggested DQN+GC-MAPPO-SL through 100 Training Episodes. The graph can be used in order to empirically verify the higher efficiency and steadiness of the proposed model compared to existing baselines in terms of the attained measure of the reward. Throughout the training period the proposed DQN+GC-MAPPO-SL model (marked by the blue curve) is steadily the highest Mean Reward, and it approaches the maximum level that can be reached of about 2000 at the very first episodes.




Figure 7. Comparison of mean reward between DQN, PPO, DDQN, and DQN+GC-MAPPO-SL

Conversely, the baseline models (DQN, PPO, and DDQN) are more volatile in the beginning and slightly less good in the end, and they usually stabilize at a Mean Reward of 1750 to 1850. In particular, PPO (grey line) has the least rapid convergence, and it is therefore far lower at the beginning but converges rapidly, whereas DQN (orange line) and DDQN (yellow line) converge more rapidly and reach a reward level that is significantly lower than that suggested by the model. The fact that the incorporation of the Graph Convolution (GC), Multi-Agent PPO (MAPPO), and Self-Learning (SL) sub-elements into the DQN+GC-MAPPO-SL framework improves the learning efficiency and optimality of the trained policy to reach and sustain the close-to-optimal performance level throughout the training procedure proves those marginal differences in the final stabilized reward (around 150-250 units).

Figure 8 sequentially compares the PS-TD3 over a 24-hour period (h) of the performance of the Proposed Model (orange bars which are DQN+GC-MAPPO-SL (proposed)) with the Existing work (blue bars which are PS-TD3)). As can be seen in the figure, the individual trend is well-defined and consistent: the load attained by the Proposed Model is much greater than the load of the Existing Work at all 24 hours. The load of the Existing Work is relatively flattened in its profile, with a range of approximately 100 kW to 350 kW on the contrary, the load of the Proposed Model is higher and of more magnitude with the maximum range being maintained at about 275 kW to 400 kW mainly in the hours between 13:00 and 16:00. The largest difference between the two models can be seen within the midday hours between 10:00 and 18:00 which implies that the Proposed Model can handle and carry a much larger power load during the best periods of the day when the system is in its peak operation. This finding illustrates the greater power management of the Proposed Model than the Existing Work which implies a better system efficiency or capacity or utilization to the intended use [22].




Figure 8. Load (kW) comparison: DQN+GC-MAPPO-SL vs. previous model

Discussion

The results of the simulation show that the suggested Hybrid DQN + GC-MAPPO-SL framework is effective in addressing the problem of optimizing the scheduling of the Virtual Power Plant (VPP) under the conditions of dynamic demand and renewable generation. The two-layer integrated structure allows localized intelligence and global decision-making to take place to ensure the system is effective in its ability to work effectively in heterogeneous Distributed Energy Resources (DERs). The cost of operation of the framework is in total cost of 812 which shows that the framework has good economic performance and still has operational feasibility as well as consistency at the system level [10]. The 77% rate of renewable energy use also demonstrates that the model manages to put the emphasis on clean energy dispatch, which allows the effective integration of solar energy and wind energy, as well as the use of the battery storage as the means of even out the oscillations and to avoid intermittency.

The insertion of the Safety Layer in GC-MAPPO-SL is a major aspect that has made the framework reliable. Some of the important constraints that have been enforced in this layer include power balance, state-of-charge constraints, limits on line flow, and ramping by projecting actions online. Consequently, the violations of zero constraints were observed throughout the whole horizon of the schedule, which proves that the system always operates and works within physical and operational boundaries. There was also the hybrid learning strategy which assisted in the stable and fast convergence where convergence was reached after 120 episodes and this indicates the advantage of the graph-based coordination and the structured guidance of the centralized critic in enhancing the learning efficiency among agents.

The framework proved to be highly scalable and capable of controlling three heterogeneous DERs and achieving the desired performance in a real-time working environment. It has rapid per-step response time of 0.3 seconds, which proves its appropriateness in online control applications in which adjustments have to be made quickly. Also, the low load balance variability of 7.25 kW shows that there is good supply demand and smooth coordination of DER, which also proves the robust operation of the system. By and large, the Hybrid DQN + GC-MAPPO-SL framework offers a robust, scalable and cost-effective programming to VPP scheduling. The strategy meets major deployment demands of the real-world scenario, especially where there is a high level of DER penetration, uncertainty, and economically optimal control strategies are required.

CONCLUSION AND FUTURE WORK

This paper proposed an example of a hybrid reinforcement learning model, which involves a combination of decentralized Deep Q-Networks (DQN) with Graph-Constrained Multi-Agent Proximal Policy Optimization and a Safety Layer (GC-MAPPO-SL) to schedule Virtual Power Plants (VPPs) in an intelligent, reliable, and cost-effective way. The results of the simulation indicate that the proposed architecture is effective in coordinating heterogeneous DERs and preserving safety and operational feasibility of operating conditions that are dynamic. The hybrid approach attained a total operating cost of 812, used renewable energy 77% of the time, and zero constraint violation because of the action projection of the real time applied by the safety layer. Moreover, it was also found that the system converged in 120 episodes, three DERs with constant scalability, and had a decision time of 0.3 s per step and responded quickly. These results prove that the given framework leads to the higher operational stability, better economic results, and offers a realistic basis of the implementation of VPP into practice. Further research on the framework could be to expand it to larger VPP ecosystems that are more diverse in DER and more geographically scattered. The use of probabilistic forecasting models may enhance the ability to adapt in case of uncertainty, whereas the use of demand response strategies might make economic optimization even better. The studies of hierarchical or meta-reinforcement learning have the potential to improve adaptable behavior to changing grid conditions in the long-term. Besides, research on the topic of communication delays, cyber-physical resilience, and explainable AI methods would enhance transparency and trust by operators. The furtherance of these directions will help in the progress of highly scalable, resilient, and smart VPP scheduling systems of next-generation smart grids.

REFERENCES

- [1] Gao H, Jin T, Feng C, Li C, Chen Q, Kang C. Review of virtual power plant operations: Resource coordination and multidimensional interaction. *Applied energy*. 2024 Mar 1; 357:122284. <https://doi.org/10.1016/j.apenergy.2023.122284>
- [2] Ceusters G, Camargo LR, Franke R, Nowé A, Messagie M. Safe reinforcement learning for multi-energy management systems with known constraint functions. *Energy and AI*. 2023 Apr 1; 12:100227. <https://doi.org/10.1016/j.egyai.2022.100227>
- [3] He S, Cui W, Li G, Xu H, Chen X, Tai Y. Intelligent Scheduling of Virtual Power Plants Based on Deep Reinforcement Learning. *Comput. Mater. Continua*. 2025 Jun 9;84(1):861-86. <http://dx.doi.org/10.32604/cmc.2025.063979>
- [4] Gao H, Jiang S, Li Z, Wang R, Liu Y, Liu J. A two-stage multi-agent deep reinforcement learning method for urban distribution network reconfiguration considering switch contribution. *IEEE Transactions on Power Systems*. 2024 Feb 28;39(6):7064-76. <https://doi.org/10.1109/TPWRS.2024.3371093>
- [5] Wang, L., Pan, T., Wang, Y., Jin, X., Yu, H., & Cao, W. (2025). A virtual power plant scheduling strategy considering cooperative participation of multiple industrial users in demand response. *AIMS Energy*, 13(4), <https://doi.org/10.3934/energy.2025029>
- [6] Klaiber J, Van Dinther C. Deep learning for variable renewable energy: A systematic review. *ACM Computing Surveys*. 2023 Aug 25;56(1):1-37. <https://doi.org/10.1145/3586006>
- [7] Goeckner A, Sui Y, Martinet N, Li X, Zhu Q. Graph neural network-based multi-agent reinforcement learning for resilient distributed coordination of multi-robot systems. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2024 Oct 14 (pp. 5732-5739). IEEE. <https://doi.org/10.1109/IROS58592.2024.10802510>
- [8] Mishra S, Bordin C, Wu Q, Manninen H. Resilient expansion planning of virtual power plant with an integrated energy system considering reliability criteria of lines and towers. *International journal of energy research*. 2022 Aug;46(10):13726-51. <https://doi.org/10.1002/er.8093>
- [9] Ye Y, Papadaskalopoulos D, Yuan Q, Tang Y, Strbac G. Multi-agent deep reinforcement learning for coordinated energy trading and flexibility services provision in local electricity markets. *IEEE Transactions on Smart Grid*. 2022 Feb 7;14(2):1541-54. <https://doi.org/10.1109/TSG.2022.3149266>
- [10] Bajaj, A., Pradeep, N. S., & Aarthi, V. (2026, April). Energy-Efficient Container Orchestration for HPC Workloads using ML in Kubernetes on Cloud Platforms. In 2026 International Conference on Recent Advancement in Electrical, Computer and Communication Technologies (IECCT) (pp. 1-6). IEEE. <https://doi.org/10.1109/IECCT68664.2026.11541572>
- [11] Jendoubi I, Bouffard F. Data-driven sustainable distributed energy resources' control based on multi-agent deep reinforcement learning. *Sustainable Energy, Grids and Networks*. 2022 Dec 1; 32:100919. <https://doi.org/10.1016/j.segan.2022.100919>
- [12] Razmi D, Babayomi O, Zhang Z. Reinforcement learning-driven dynamic Model Predictive Control for adaptive real-time multi-agent management of microgrids. *International Journal of Electrical Power & Energy Systems*. 2025 Sep 1; 170:110823. <https://doi.org/10.1016/j.ijepes.2025.110823>
- [13] Akhoundzadeh P, Mirjalily G, Sadeghi MT. Optimal D2D resource allocation in heterogeneous cellular networks by decentralized multi-agent deep Q-learning. In 2024 32nd International Conference on Electrical Engineering (ICEE) 2024 May 14 (pp. 1-5). IEEE. <https://doi.org/10.1109/ICEE63041.2024.10668150>
- [14] Wang J, Guo C, Yu C, Liang Y. Virtual power plant containing electric vehicles scheduling strategies based on deep reinforcement learning. *Electric power systems research*. 2022 Apr 1; 205:107714. <https://doi.org/10.1016/j.epsr.2021.107714>
- [15] Yang X, Liu H, Wu W. Attention-enhanced multi-agent reinforcement learning against observation perturbations for distributed volt-Var control. *IEEE Transactions on Smart Grid*. 2024 Jul 5;15(6):5761-72. <https://doi.org/10.1109/TSG.2024.3423700>
- [16] Xiao D. A review on risk-averse bidding strategies for virtual power plants with uncertainties: Resources, technologies, and future pathways. *Technologies*. 2025 Oct 28;13(11):488. <https://doi.org/10.3390/technologies13110488>
- [17] Fan Z, Zhang W, Liu W. Multi-agent deep reinforcement learning-based distributed optimal generation control of DC microgrids. *IEEE Transactions on Smart Grid*. 2023 Jan 17;14(5):3337-51. <https://doi.org/10.1109/TSG.2023.3237200>
- [18] Seyfi M, Mehdinejad M, Mohammadi-Ivatloo B, Shayanfar H. Deep learning-based scheduling of virtual energy hubs with plug-in hybrid compressed natural gas-electric vehicles. *Applied Energy*. 2022 Sep 1; 321:119318. <https://doi.org/10.1016/j.apenergy.2022.119318>
- [19] Hu D, Ye Z, Gao Y, Ye Z, Peng Y, Yu N. Multi-agent deep reinforcement learning for voltage control with coordinated active and reactive power optimization. *IEEE Transactions on Smart Grid*. 2022 Jun 24;13(6):4873-86. <https://doi.org/10.1109/TSG.2022.3185975>
- [20] Ghasemi Olanlari F, Amraee T, Moradi-Sepahvand M, Ahmadian A. Coordinated multi-objective scheduling

- of a multi-energy virtual power plant considering storages and demand response. IET Generation, Transmission & Distribution. 2022 Sep;16(17):3539-62. <https://doi.org/10.1049/gtd2.12543>
- [21] Xu B, Luan W, Yang J, Zhao B, Long C, Ai Q, Xiang J. Integrated three-stage decentralized scheduling for virtual power plants: A model-assisted multi-agent reinforcement learning method. Applied Energy. 2024 Dec 15; 376:123985. <https://doi.org/10.1016/j.apenergy.2024.123985>
- [22] Sun Z, Lu T. Collaborative operation optimization of distribution system and virtual power plants using multi-agent deep reinforcement learning with parameter-sharing mechanism. IET Generation, Transmission & Distribution. 2024 Jan;18(1):39-49. <https://doi.org/10.1049/gtd2.13037>