A Multi-Timescale Demand Response Pricing Method for VPP Considering Privacy Protection and Renewable Energy Integration

Linhui Zhang; Xiangyu Kong*; Jin Su; Gaohua Liu
State Key Laboratory of Intelligent Power Distribution Equipment and System
Tianjin University
Tianjin, China
zhang_lin_hui@163.com; eekongxy@tju.edu.cn

Yi Qi; Jianfeng Gong
State Grid Ningxia Electric Power Company
Ningxia, China
jinpanlong@163.com

***Abstract*—To address privacy protection and renewable energy accommodation challenges faced by Virtual Power Plant (VPP) in multi-timescale Demand Response (DR), this paper proposes a VPP DR pricing method that integrates federated learning and deep reinforcement learning. To mitigate difficulties in aggregating multi-agent load data and the high risks of privacy leakage during day-ahead price decision-making, an improved horizontal federated learning framework is introduced, in which encrypted model parameters replace raw data sharing, enabling localized user-side training and secure collaborative updates. To tackle the challenges posed by renewable generation fluctuations to the continuity and agility of real-time pricing strategies, Neural Turing Machine (NTM) is employed to dynamically evaluate users' DR potential; in addition, gated recurrent units combined with an attention mechanism are used to enhance the accuracy of user behavior prediction. To improve the convergence speed and exploration quality of reinforcement learning in real-time environments, a mixed-experience multi-agent deep deterministic policy gradient (ME-MADDPG) method is proposed, enabling efficient price adjustments under rapidly changing renewable energy scenarios. Case studies demonstrate that the proposed method not only enhances user information security and response fairness, but also effectively reduces power fluctuations, improves real-time control flexibility, and significantly increases the level of renewable energy utilization.**

Keywords—demand response, price decision-making, potential assessment, federated learning, deep learning

I. Introduction

With the rapid integration of distributed energy resources and various types of flexible loads, the operation paradigm of power systems is gradually shifting from traditional centralized control to a more flexible and efficient Virtual Power Plant (VPP) model. By aggregating distributed energy resources and heterogeneous loads, VPP can enhance renewable energy accommodation while enabling precise demand-side regulation. As a key operational mechanism of VPP, Demand Response (DR) plays a critical role in stimulating user participation and improving system flexibility, in which pricing strategy serves as a central driver.

DR pricing approaches can be broadly categorized into optimization-based pricing, game-theoretic pricing, and reinforcement-learning-based dynamic pricing methods [1]. In terms of optimization theory, reference [2] improves the convergence speed and accuracy of real-time pricing algorithms through cross-iteration of Lagrange multipliers and the alternating direction method of multipliers. Reference [3] converts the optimization model into a system of equations with Lagrange multipliers based on Karush–Kuhn–Tucker (KKT) conditions, and then applies a smooth Newton method to enable fast real-time price optimization. Reference [4], considering different characteristics of user utility functions, investigates the use of piecewise linear approximation to achieve smooth real-time price computation.

Game-theoretic pricing mainly determine incentive prices based on the interactions between energy suppliers and load users. Common game-theoretic approaches can be classified into cooperative and non-cooperative games [5]. As a service aggregator, a VPP can be viewed as a global cooperative game involving multiple user entities, with the objective of achieving optimal performance for the coalition as a whole [6].

In practice, the dynamic decision-making process of electricity pricing can be formulated as a Markov decision problem. Reinforcement learning has strong advantages in modeling the interactions between users and the VPP environment, and has thus been widely adopted in recent years [7]. Reference [8] combines long short-term memory (LSTM) networks with reinforcement learning for potential assessment and decision-making. Reference [9] leverages a cloud–edge architecture to reduce the computational burden of reinforcement learning on centralized decision-making units.

However, as user-side data become increasingly high-dimensional and collected at higher frequencies, DR pricing relies more heavily on detailed user data modeling, leading to growing concerns regarding privacy leakage. This issue has become a major barrier to large-scale deployment of DR mechanisms. Meanwhile, the strong volatility and uncertainty of renewable energy output impose stricter requirements on real-time DR pricing, making multi-timescale, renewable-oriented pricing strategies particularly necessary.

To address these challenges, this paper proposes a multi-timescale DR pricing method for VPP considering both privacy protection and renewable energy accommodation. The main contributions are as follows:

(1) A day-ahead pricing decision method based on improved horizontal federated learning is proposed. A day-ahead optimization model for VPP is constructed. Encrypted parameter transmission is adopted during model training to replace the transfer of raw user data, thereby reducing privacy leakage risks arising from off-site load data sharing.

(2) A real-time pricing decision method based on ME-MADDPG is proposed. First, a multi-objective optimization model considering smooth renewable energy integration is established, and neural Turing machine (NTM) is introduced to evaluate users' response potential. A gated recurrent unit combined with an attention mechanism is jointly trained to improve the accuracy of multi-user potential assessment. Then, the mixed-experience multi-agent deep deterministic policy gradient (ME-MADDPG) algorithm is employed to achieve efficient and continuous optimization of real-time pricing decisions.

This work was supported by Technology Project of State Grid Corporation (5229NX25000F-055-ZN).

979-8-3315-8138-1

II. Multi-Timescale DR Pricing Process for VPP

This paper proposes a multi-timescale DR pricing method for VPP that accounts for both privacy protection and renewable energy accommodation, which illustrated in Fig. 1.

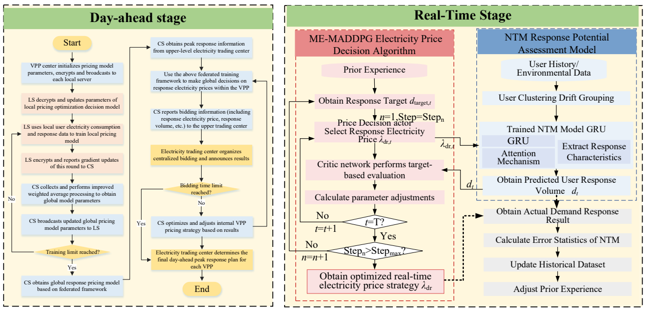

Fig. 1. Multi-timescale DR pricing process for VPP.

III. Day-Ahead DR Pricing Method for VPP Considering Privacy Protection

A. Day-Ahead DR Pricing Model

The day-ahead controllable load management of VPP focuses on incentivizing users to participate in DR, assisting the system in peak shaving and valley filling. As a profit-oriented entity, the VPP participates in the upper-level operator's load response bidding with price $\lambda_{\mathrm{vpp},t}$, while internally formulating a DR price $\lambda_{\mathrm{dr},t}$ for its users. The VPP management center earns revenue by leveraging the price difference between the two levels. The objective function is defined as:

$$\max F\left(\lambda_{\mathrm{dr}}\right)=\sum_{t=1}^{T} f\left(\lambda_{\mathrm{dr},t}\right)=\sum_{t=1}^{T}\left(f_{\mathrm{b}}\left(\lambda_{\mathrm{dr},t}\right)-f_{\mathrm{p}}\left(\lambda_{\mathrm{dr},t}\right)\right) \quad (1)$$

$$f_{\mathrm{b}}\left(\lambda_{\mathrm{dr},t}\right)=\left(\lambda_{\mathrm{vpp},t}-\lambda_{\mathrm{dr},t}\right)\sum_{i=1}^{N} d_{i,t}\gamma D_{\mathrm{M}i} \quad (2)$$

$$f_{\mathrm{p}}\left(\lambda_{\mathrm{dr},t}\right)=\lambda_{\mathrm{p}}\sum_{i=1}^{N}\left|d_{i,t}-d_{\mathrm{Aci},t}\right| \quad (3)$$

where $f\left(\lambda_{\mathrm{dr},t}\right)$ is the net revenue obtained by the VPP at time t through organizing internal DR, $f_{\mathrm{b}}\left(\lambda_{\mathrm{dr},t}\right)$ is the electricity sales revenue, and $f_{\mathrm{p}}\left(\lambda_{\mathrm{dr},t}\right)$ is the penalty function incurred when the VPP fails to achieve the scheduled load reduction. The term $\lambda_{\mathrm{dr},t}$ is the response price issued by the VPP at time t, $\lambda_{\mathrm{vpp},t}$ is the bid price submitted to the upper-level operator, and λ_{p} is the penalty price. The parameter γ is used to control the influence of the user contribution factor.

Considering the limits on user load adjustment and acceptable price fluctuations, the constraints are given as:

$$\begin{cases}\lambda_{\min,t}\le\lambda_{\mathrm{dr},t}\le\lambda_{\max,t}\\ 0\le d_{i,t}\le\min\left(\delta d_{\max,t},d_{\mathrm{target},t}\right)\end{cases} \quad (4)$$

where $\lambda_{\min,t}$ and $\lambda_{\max,t}$ is the minimum and maximum allowable DR prices; δ is the maximum load adjustment coefficient; $d_{\max i,t}$ is the maximum load of user i at time t; and $d_{\mathrm{target},t}$ is the load adjustment target of the VPP at time t.

B. Encrypted Horizontal Federated Learning Training Framework Based on Improved Weighted Averaging

The horizontal federated learning training framework employs partial homomorphic encryption (PHE) to protect user data privacy and incorporates a weighted aggregation strategy based on local model quality (W-FedAvg) into the traditional Federated Averaging (FedAvg) method. The framework consists of two layers: a global layer and a local layer. The global layer, located at the VPP management center, performs encrypted aggregation and updates of the pricing optimization model parameters using W-FedAvg. The local layer is composed of local servers at each organizational unit, which have direct access to user response data for local model training, after which encrypted gradients are transmitted to the global layer. Under global coordination, local servers with similar data characteristics collaboratively train a unified VPP pricing optimization model. Throughout this process, both local data and key model information are encrypted using PHE to ensure privacy protection.

Federated averaging is a distributed optimization method suitable for horizontal federated learning, where local servers perform gradient descent in parallel, and a central server conducts weighted aggregation. Suppose the VPP manages N units, with k_n users participating in DR in unit n, and the total number of users is K. Then, during the $(s+1)$-th iteration, the model gradient weight of the n-th local server in the FedAvg framework is:

$$\omega_{s+1}\leftarrow\sum_{n=1}^{N}\frac{k_n}{K}\omega_{(n)s+1} \quad (5)$$

where k_n is number of users participating in DR in unit n; K is total number of users; N is total number of units; s is iteration index; ω_s is the updated global model parameters after s-th iteration; and $\omega_{(n)s}$ is updated local model parameters of unit n obtained through gradient descent on local data.

In practical training, the number of users and the quality of their data vary across units. Directly using the traditional weighted averaging strategy may amplify the impact of low-quality or small-scale data, contaminating the global model parameters and degrading the overall performance of the pricing model. To address this issue, this paper introduces the W-FedAvg that aggregates gradients with enhanced differentiation based on data quality. This method augments traditional FedAvg by incorporating weight evaluation based on data accuracy and computational quality. The core idea is to divide local data into training samples and testing samples: the former are used for standard training on the initial global model, while the latter evaluate the post-training model accuracy to reflect data quality and training performance. The training accuracy q_n of the n-th local server is defined as:

$$q_n=C_n/X_n \quad (6)$$

where C_n is the number of correctly predicted samples on the testing set, and X_n is the total number of testing samples.

After obtaining the training accuracy of each local server, the traditional user-count-based weight is replaced with a quality-aware weighting factor, leading to a redefined improved weighted federated averaging algorithm:

$$\bar{\omega}_{s+1}\leftarrow\sum_{n=1}^{N}\alpha_{(n)s+1}\omega_{(n)s+1} \quad (7)$$

979-8-3315-8138-1

$$\alpha_{(n)s+1} = k_n \left(q_{(n)s+1}\right)^{-L} / \sum_{n=1}^{N} k_n \left(q_{(n)s+1}\right)^{-L} \tag{8}$$

where $\overline{\omega}_{s+1}$ is the updated global model parameters after the s-th iteration and also serves as the initial parameters for the *(s+1)*-th iteration; $\omega_{(n)s}$ is the parameter set of the n-th local model at iteration s; $\alpha_{(n)s+1}$ is the training quality factor of unit n; $q_{(n)s+1}$ is the training accuracy of the n-th local model; and L is a tunable hyperparameter used to adjust the influence of the training quality factor.

Although the above algorithm does not directly transmit users' raw response data, it still exposes intermediate computational parameters and provides no additional protection for data structure. Therefore, this paper enhances the security of the framework by incorporating PHE, which supports additive homomorphism and scalar multiplication homomorphism. For two values $\boldsymbol{u}$ and $\boldsymbol{v}$ with homomorphic encodings [[u]] and [[v]], the following properties hold:

Additive homomorphism:

$$\mathrm{Dec}_{\mathrm{sk}}\left(\left[\left[u\right]\right] \oplus \left[\left[v\right]\right]\right) = \mathrm{Dec}_{\mathrm{sk}}\left(\left[\left[u+v\right]\right]\right) \tag{9}$$

Scalar multiplication homomorphism:

$$\mathrm{Dec}_{\mathrm{sk}}\left(\left[\left[u\right]\right] \odot m\right) = \mathrm{Dec}_{\mathrm{sk}}\left(\left[\left[u \bullet m\right]\right]\right) \tag{10}$$

where Dec is the decryption function, sk is the secret key, $\oplus$ is multiplication performed on ciphertexts, and $\odot$ is the ciphertext raised to the power m.

Based on the improved weighted federated averaging method and PHE-based partial homomorphic encryption, the encrypted horizontal federated learning training framework proposed in this paper is illustrated in Fig. 2. The overall process consists of four stages: parameter assignment, model training, gradient descent, and parameter aggregation.

IV. Real-Time DR Pricing Method for VPP Considering Renewable Energy Accommodation

A. Real-Time DR Pricing Model

To mitigate the power fluctuations caused by renewable energy integration, VPP can mobilize users to participate in DR and conduct a multi-objective optimization process that accounts for smooth renewable energy accommodation. In this paper, the total energy output fluctuation rate at time t is denoted by δ_t:

$$\delta_t = (P_{\mathrm{G},t} - P_{\mathrm{G},t-1}) / \Delta t \tag{11}$$

where $P_{\mathrm{G},t}$ is the actual aggregated output.

The multi-objective optimization problem includes renewable energy accommodation revenue B_{R}, power fluctuation cost B_{F}, and user DR revenue B_{U}. Renewable energy accommodation revenue aims to reduce curtailment and increase the utilization of renewable generation. Power fluctuation cost reflects the feeder power fluctuation induced by the output of all energy resources, particularly after incorporating micro gas turbines. User DR revenue consists of users' reduced energy costs and penalties associated with response deviation.

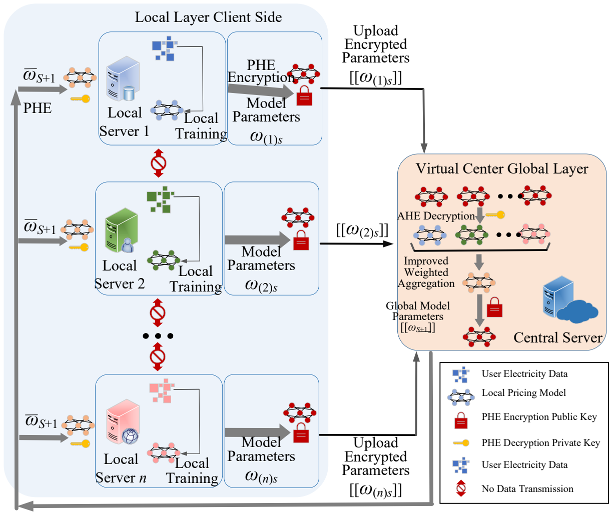

Fig. 2. Encrypted horizontal federated learning training framework based on improved weighted averaging.

The objective function and constraints are given by:

$$\max B(\lambda) = B_{\mathrm{R}} + B_{\mathrm{F}} + \eta B_{\mathrm{U}} \tag{12}$$

$$B_{\mathrm{R}} + B_{\mathrm{F}} = \sum_{t=1}^{T_{\mathrm{new}}} \left[\lambda_{\mathrm{o},t} \left(P_{\mathrm{R},t} - P_{\mathrm{Ro},t} \right) - \lambda_{\mathrm{flu}} \left| \delta_t \right| \right] \tag{13}$$

$$B_{\mathrm{U}} = \sum_{t=1}^{T_{\mathrm{dr}}} \left(\lambda_{\mathrm{o},t} - \lambda_{\mathrm{dr},t} \right) d_t - \lambda_{\mathrm{p},t} \left| d_t - d_{\mathrm{target},t} \right| \tag{14}$$

where η is the weighting factor between the two objectives; when η=0, the optimization focuses solely on smooth renewable energy accommodation. In practice, since the frequency of DR is lower than that of renewable energy fluctuation, T_{new} and T_{dr} are used to represent the number of time periods for calculating power fluctuation and DR.

At time t, $\lambda_{\mathrm{o},t}$ is the pre-response electricity purchasing price; λ_{flu} is the power fluctuation cost coefficient; $\lambda_{\mathrm{dr},t}$ is the real-time DR price; and $\lambda_{\mathrm{p},t}$ is the penalty coefficient for response deviation.

The target response $d_{\mathrm{target},t}$ of a user is determined jointly by renewable generation output and the fluctuation rate δ_t:

$$d_{\mathrm{target},t} = \begin{cases} P_{\mathrm{G},t-1} - \delta_{\mathrm{drop}} \Delta t, & \delta_t < \delta_{\mathrm{drop}} \\ P_{\mathrm{G},t-1} - \delta_{\mathrm{rise}} \Delta t, & \delta_t > \delta_{\mathrm{rise}} \\ 0, & \delta_{\mathrm{drop}} \le \delta_t \le \delta_{\mathrm{rise}} \end{cases} \tag{15}$$

where δ_{drop} and δ_{rise} denote the lower and upper bounds of the permissible power adjustment rate.

B. DR Potential Assessment Based on NTM

In VPP, users participating in DR are numerous and diverse, which limits the efficiency and accuracy of potential assessment. Therefore, this paper employs NTM to evaluate users' DR potential.

1) Mean-Shift Clustering Preprocessing

To enhance the training efficiency of the subsequent NTM model, a mean-shift clustering algorithm is first applied to group users with similar characteristics. A Gaussian kernel function $\mathrm{G}(p)$ is adopted to quantify the actual contribution of each user point to the cluster shift.

2) DR Feature Extraction Based on GRU and Attention Mechanism

User DR behavior is influenced by both external factors and internal factors. This paper adopts a Gated Recurrent Unit (GRU) combined with an attention mechanism to extract temporal features from the input time-series data.

At time *t*, a GRU unit receives the input feature x_t and the previous memory state h_{t-1}. The computation involves the update gate, reset gate, and candidate state:

$$z_{i,t} = \sigma\left(\boldsymbol{W}_z x_{i,t} + \boldsymbol{U}_z m_{i,t-1}\right) \tag{16}$$

$$r_{i,t} = \sigma\left(\boldsymbol{W}_r x_{i,t} + \boldsymbol{U}_r m_{i,t-1}\right) \tag{17}$$

$$h^*_{i,t} = \tanh\left(\boldsymbol{W} x_{i,t} + r_{i,t} \odot \boldsymbol{U} m_{i,t-1}\right) \tag{18}$$

$$\boldsymbol{h}_{i,t} = z_{i,t} \odot \boldsymbol{m}_{i,t-1} + \left(1 - z_{i,t}\right) \odot h^*_{i,t} \tag{19}$$

where $\sigma(\bullet)$ and $\tanh(\bullet)$ are activation functions; $\boldsymbol{W}_z, \boldsymbol{W}_r, \boldsymbol{W}$, $\boldsymbol{U}_z, \boldsymbol{U}_r$ and $\boldsymbol{U}$ denote the weight matrices corresponding to the update gate, reset gate, and candidate state computation. The output h_t serves as the feature vector passed to the attention mechanism and memory update module.

To enhance model's ability to capture key features under varying environmental conditions, an attention mechanism is introduced to update the memory module of the NTM. Fig. 4 illustrates the feature extraction process combining GRU and the attention mechanism.

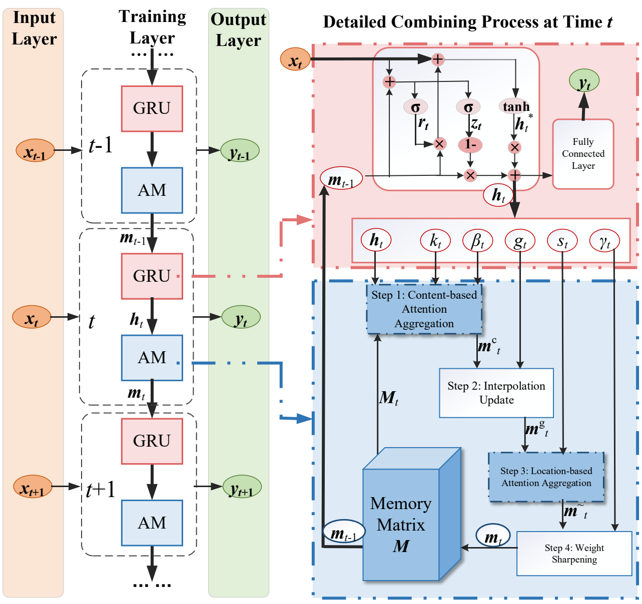

Fig. 3. DR feature extraction based on GRU and attention mechanism.

C. Real-time electricity price optimization strategy based on ME-MADDPG

Due to the limitations of the traditional MADDPG algorithm in terms of optimality and convergence , this paper adopts an improved ME-MADDPG algorithm to make electricity price decisions for VPP. The ME-MADDPG algorithm enhances MADDPG algorithm by incorporating a sample generator based on artificial potential field (APF), ensuring the generation of high-quality samples. Additionally, a dynamic hybrid sampling strategy is employed to adjust the proportion of data from different sources, increasing sample utilization and improving the efficiency of the algorithm. The framework of the algorithm is shown in Fig. 4.

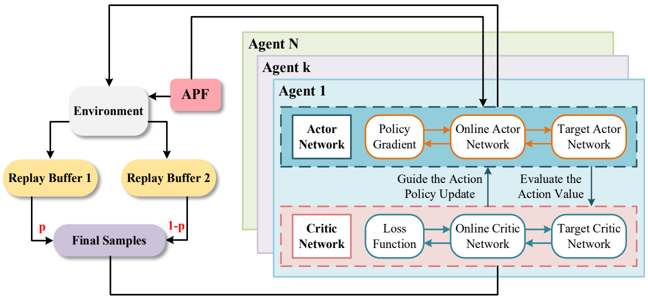

Fig. 4. ME-MADDPG algorithm framework.

V. Case Study

A. Day-Ahead DR Pricing

To verify the effectiveness of the proposed privacy-preserving day-ahead DR pricing method based on federated learning, a dataset of electricity DR from June 18, 2024 to June 17, 2025 was used. The dataset was divided into training, validation, and test sets in a 7:1:2 ratio. The test was conducted under three cases:

Case 1: Fixed electricity price at 0.75 ¥/kWh.

Case 2: Time-of-use pricing with peak, flat, and valley prices set at 0.95, 0.70, and 0.65 ¥/kWh.

Case 2: VPP-based day-ahead DR pricing, with a 24-hour pricing cycle and a granularity of 24 points per day. The day-ahead pricing range is [0.5, 1.25] ¥/kWh.

A typical day was selected for response pricing analysis. Users' load profiles under the three cases are shown in Fig. 5. VPP management center received load reduction requests for seven peak periods. Based on this, the day-ahead pricing optimization was performed with a range of [0.5, 1.25] ¥/kWh and a step size of 0.05 ¥/kWh.

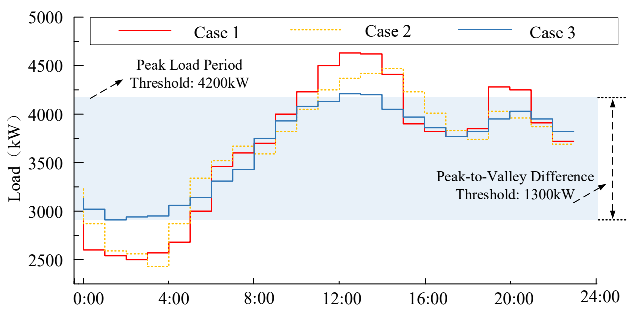

Fig. 5. Load profiles under different cases.

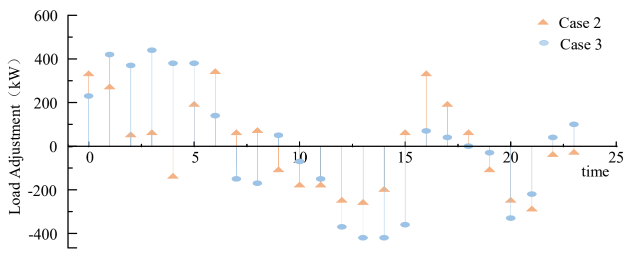

Fig. 6. User DR load adjustment.

979-8-3315-8138-1

For Case 2 and 3, the electricity prices and corresponding user load adjustments for each period are shown in Fig. 7. Case 2 cannot adjust according to actual daily load. The peak-to-valley load difference under Case 2 is 2040 kW, leaving considerable room for adjustment. In contrast, Case 3 can better stimulate users' response potential, increasing flexible load adjustment and achieving a peak-to-valley difference of 1290 kW, demonstrating effective peak shaving.

B. *Intraday DR Pricing*

To validate the proposed real-time pricing method, three cases were analyzed. Case 4 is real-time pricing; Case 5 is time-of-use pricing and Case 6 is fixed electricity price.

Users were classified into three types A, B, and C. Real-time price incentives were applied across 12 periods in a day, with target and actual response volumes shown in Fig. 8. Different user types exhibit distinct response characteristics. Group A has higher night-time response potential, reducing renewable curtailment, but limited daytime flexibility. Group B has good daytime response capability, enabling peak load reduction at lower price incentives. Group C has moderate overall response potential, evenly distributed.

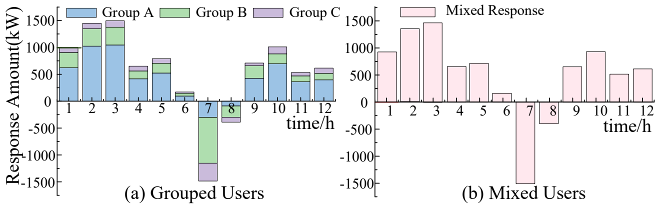

Fig. 7. User demand response load adjustment.

Fig. 9 presents response deviation rates and penalty costs. After user classification, the attention mechanism in the NTM enhances the extraction of response features for different user types, improving assessment accuracy and reducing penalty costs, thereby increasing overall optimization benefits.

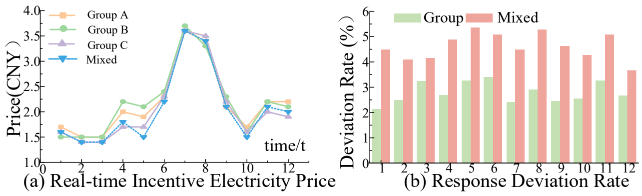

Fig. 8. Incentive real-time electricity price and response deviation rate.

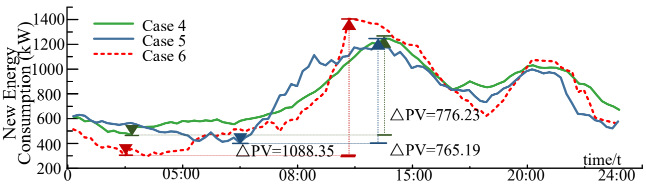

Fig. 9. Renewable energy utilization over 96 periods.

DR can be coordinated with renewable energy integration to enhance system benefits. For a given load adjustment target, real-time price optimization was conducted over 96 periods, resulting in renewable energy utilization and fluctuation rates shown in Fig. 10 and Fig. 11. During high-demand periods, electricity prices increased, prompting users to reduce load and shift consumption to periods with abundant power supply. The proposed strategy effectively suppressed renewable power fluctuations, reduced peak-to-valley differences, and improved system operational safety.

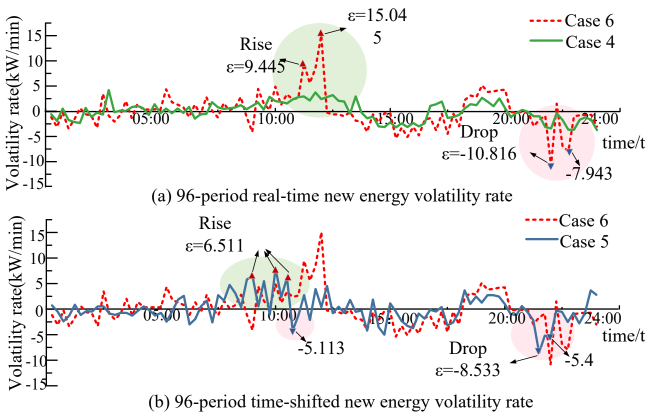

Fig. 10. Renewable energy fluctuation rates over 96 periods.

VI. Conclusion

This paper proposes a multi-timescale DR pricing method for VPP that considers both privacy protection and renewable energy accommodation. Case studies demonstrate that the day-ahead pricing optimization strategy based on encrypted horizontal federated learning can protect user privacy while improving control fairness and day-ahead response efficiency. Moreover, the real-time pricing strategy based on ME-MADDPG enhances the efficiency and accuracy of VPP real-time pricing, reduces power fluctuations, and increases the utilization of renewable energy.

Acknowledgment

This work was supported by Technology Project of State Grid Corporation (5229NX25000F-055-ZN).

References

[1] X. Sun, H. Xie, Y. Xiao, and Z. Bie, "Incentive compatible pricing for enhancing the controllability of price-based demand response", IEEE Trans. Smart Grid, vol. 15, no. 1, pp. 418–430, Jan. 2024.

[2] W. Wang, Y. Han, M. Wang, and Y. He, "Research on fair residential critical peak price: Based on a price penalty mechanism for high-electricity consumers", Applied Energy, vol. 351, p. 121892, Dec. 2023.

[3] J. Einolander, A. Kiviaho, and R. Lahdelma, "Detecting changes in price-sensitivity of household electricity consumption: The impact of the global energy crisis on implicit demand response behavior of finnish detached households", Energy and Buildings, vol. 306, p. 113941, Mar. 2024.

[4] O. B. Adewuyi and H. Aki, "Optimal planning for high renewable energy integration considering demand response, uncertainties, and operational performance flexibility", Energy, vol. 313, p. 134021, Dec. 2024.

[5] R. Cheng, G. Liu, W. Liu, Q. Shi, and Q. Chen, "A personalized network-aware price design for electric vehicle charging using a sequential stackelberg game", Applied Energy, vol. 381, p. 125145, Mar. 2025.

[6] B. Li, Q. Yang, and I. Kamwa, 'A novel stackelberg-game-based energy storage sharing scheme under demand charge', IEEE/CAA J. Autom. Sinica, vol. 10, no. 2, pp. 462–473, Feb. 2023.

[7] C. Lu, S. Li, S. Reddy Penaka, and T. Olofsson, "Automated machine learning-based framework of heating and cooling load prediction for quick residential building design", Energy, vol. 274, p. 127334, July 2023.

[8] X. Kong, D. Kong, J. Yao, L. Bai, and J. Xiao, "Online pricing of demand response based on long short-term memory and reinforcement learning", Applied Energy, vol. 271, p. 114945, Aug. 2020.

[9] Z. Ye, D. Zhang, Z.-G. Wu, and H. Yan, "A3C-based intelligent event-triggering control of networked nonlinear unmanned marine vehicles subject to hybrid attacks", IEEE Trans. Intell. Transport. Syst., vol. 23, no. 8, pp. 12921–12934, Aug. 2022.