

Article

Two-Stage Stochastic Frequency-Security-Constrained Unit Commitment for Thermal-Storage Joint Frequency Regulation Under High Renewables Using Analytical Criterion and Linear Surrogates

Guodong Wang ¹, Ran Sun ¹, Jianbo Wang ¹, Xiaoke Zhang ^{2,*}, Xinjian Jiang ³, Zhijian Ling ³ 
and Zhenghui Zhao ³ 

- ¹ State Grid Henan Electric Power Company, Zhengzhou 450051, China; wangguodong1@ha.sgcc.com.cn (G.W.); sunran@ha.sgcc.com.cn (R.S.); wangjianbo5@ha.sgcc.com.cn (J.W.)
² Electric Power Research Institute, State Grid Henan Electric Power Company, Zhengzhou 450051, China
³ Shanxi Energy Internet Research Institute, Taiyuan 030006, China; jiangxj@mail.tsinghua.edu.cn (X.J.); lzj1991@ujs.edu.cn (Z.L.); zhzhao@ujs.edu.cn (Z.Z.)
* Correspondence: zhangxiaoke2@ha.sgcc.com.cn

Abstract

In modern power systems, the rapid growth of renewable energy capacity, such as wind and solar photovoltaic (PV) power, has led to a decline in system equivalent inertia and primary frequency regulation margin. At the same time, net load fluctuations have intensified across multiple time scales, making it more likely for the RoCoF, frequency nadir, and quasi-steady-state frequency deviation to approach safety limits following disturbances. To achieve a balance between frequency security and economic operation, this paper proposes a two-stage stochastic frequency-security-constrained unit commitment (FSC-SUC) model tailored for scenarios with high renewable energy penetration. The day-ahead hourly dispatch stage jointly determines the on/off status and reference output of synchronous units and the reservation of slow frequency regulation capacity, as well as energy storage charging and discharging plans, SoC trajectories, and the reservation of fast frequency regulation capacity. The intraday minute-level real-time dispatch stage accommodates prediction errors through scenario-based rescheduling and ensures the deliverability of both slow and fast frequency regulation capabilities via commitment consistency constraints. To address the challenge of directly embedding frequency nadir constraints into mixed-integer optimization, this paper employs a modeling approach that combines analytical criteria with linear surrogate constraints. The RoCoF and quasi-steady-state frequency deviation are specified via aggregated analytical constraints, while the nadir is embedded into the main problem after generating samples offline using a simplified frequency response model and training a polyhedral linear surrogate for external approximation. The safety margin is then calibrated using high-quantile residuals from the validation set to ensure conservativeness. Case studies on the IEEE 33-bus system under different renewable penetration levels demonstrate that the proposed method significantly reduces the probability of frequency nadir violations and load-loss risk with only a modest cost increase while also improving coordination between fast and slow frequency regulation.

Keywords: high renewable energy penetration; frequency security constraint; unit commitment; two-stage stochastic optimization; battery energy storage; linear surrogate model; benders decomposition



Academic Editor: Claus Leth Bak

Received: 26 May 2026

Revised: 29 June 2026

Accepted: 29 June 2026

Published: 1 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

1. Introduction

In modern power systems, the installed capacity and output share of renewable energy sources such as wind and solar photovoltaic (PV) power have grown rapidly, leading to a gradual shift in the system's structure from one dominated by synchronous generators to one in which power electronics-based resources account for a significantly larger share. First, the output of renewable energy sources is highly random and volatile. Net load may exhibit characteristics where both the amplitude and rate of change increase simultaneously on hourly, minute, and even shorter time scales, resulting in more frequent system frequency regulation requirements. Second, the decline in the proportion of synchronous generators has reduced the system's equivalent inertia and primary frequency regulation margin. As a result, the rate of change of frequency (RoCoF), frequency nadir, and quasi-steady-state frequency deviation are more likely to approach safety limits when the system encounters disturbances, and frequency risks exhibit greater temporal and scenario-specific variability. Once frequency constraints are shifted to the level of unit combinations and reserve configurations, the system must either increase the number of synchronous generators in operation or raise the frequency reserve margin, which triggers a chain reaction affecting fuel costs, start-up and shutdown costs, and curtailment of wind and solar power. Consequently, developing a system optimization model capable of capturing the interplay among high renewable energy uncertainty, low-inertia constraints, and frequency support coupling across multiple resources while balancing physical interpretability and computational feasibility has become one of the most pressing challenges in the field of coordinated frequency security and economic operation.

The formulation of frequency safety metrics as constraints and their solvable embedding have been studied in a relatively systematic manner within the UC and SCUC frameworks. Reference [1] embeds the inertia metric into the UC in a linearized manner, enabling the day-ahead phase to explicitly constrain the inertia level and the frequency safety margin. Reference [2] establishes a quantitative mapping relationship between frequency security constraints and frequency control reserve requirements, providing a computable basis for the computational allocation of reserve capacity. Reference [3] takes into account the spatial heterogeneity of low-inertia systems; the model incorporates position-dependent frequency stability constraints into the SCUC to characterize the effects of inertia distribution and network structure on frequency risk. To address potential discrepancies between frequency-based evaluation and engineering surveying, Reference [4] introduces a learning-assisted RoCoF evaluation and constraint handling framework, which helps mitigate misalignment driven by frequency constraints. At the real-time operational level, treating virtual inertia as a schedulable resource and jointly optimizing it with real-time economic scheduling can also reveal the opportunity costs and availability constraints associated with inertia support [5]. At the same time, to address accuracy issues caused by forecasting errors in renewable energy and reserve capacity, Reference [6] employs forecast correction and reserve forecasting improvements to enhance the economic efficiency of the UC system and reduce unnecessary reserve capacity. Overall, the aforementioned studies indicate that frequency security has gradually evolved from a post hoc verification process into a key constraint dimension in operational optimization; however, further research is needed on solvable and interpretable constraint formulations and solution mechanisms across different time scales.

On the other hand, the challenge of efficiently embedding frequency constraints into UC lies in the strong nonlinearity of the frequency response process. In particular, nadir constraints involve the coupling of multiple factors, including inertia, first-order frequency response, rapid response capability, and disturbance magnitude; direct modeling of these factors often leads to computational difficulties and scalability limitations. Reference [7] de-

rives the nadir constraint for quadratic forms based on the principle of energy conservation and establishes a mixed-integer quadratically constrained programming (MIQCP) solution framework, thereby enhancing the model's physical interpretability. Given the mechanistic complexity associated with the frequency response of renewable energy sources, refining the modeling of wind turbine frequency response and using it to construct frequency safety margins that can be integrated into the UC system helps avoid oversimplifying frequency dynamics while also reducing the resulting conservatism [8]. Reference [9] couples wind power reserve supply with frequency constraints into the unit configuration, thereby reflecting the interaction between reserve structure and frequency security under conditions of high wind power penetration. Reference [10] incorporates engineering factors such as time delays associated with wind power frequency modulation and establishes a joint dispatch model for the power grid and wind farm to enhance the reliability of frequency support capability assessments. Reference [11] simultaneously incorporates demand response and frequency stability constraints in a scenario with high wind power penetration, thereby reflecting the real-world situation where multiple types of flexible resources collectively ensure frequency stability. Although the aforementioned studies offer complementary approaches to frequency-constrained expression and solver-based methods, model scale and computational stability remain unresolved challenges in current frequency-constrained security dispatch when high penetration of renewable energy, multi-scenario uncertainty, and multiple frequency regulation resources coexist.

As frequency support solutions have expanded from traditional synchronous generators to include various types of power electronic interfaces and cross-regional interconnection support, research on frequency security dispatch has increasingly emphasized cross-regional and multi-party coordination. Reference [12] establishes a unified frequency constraint (UFC) for HVDC interconnection systems and employs a decomposition strategy to enhance large-scale solution capabilities. Reference [13] explicitly models the frequency response characteristics of VSC HVDC in asynchronous interconnected systems and introduces distributed robust chance-constrained and convex optimization tools under renewable energy uncertainty to address nonlinear constraints related to maximum frequency deviation while balancing safety and economic efficiency. Related research on the coordinated optimization of frequency maintenance services for renewable energy-dominated systems has also evaluated the operational value of low-carbon resources in providing frequency support [14]. Reference [15] optimizes the allocation of fast-response resources from the perspective of inertia distribution, highlighting the coupling between frequency risk and resource location in low-inertia systems. Regarding the coordination of power generation and consumption, Reference [16] proposes a framework for the coordinated optimization of energy, frequency regulation, and voltage support for distributed virtual power plants (VPPs); this framework incorporates data-driven distributed robust modeling methods to enhance system operational robustness under conditions of uncertainty. Reference [17] employs a Nash bargaining mechanism to achieve coordinated scheduling of frequency constraints between distribution grids and microgrids, providing a computable framework for multi-agent collaborative frequency support. In the temporal dimension, real-time hierarchical flexibility management methods emphasize the coordinated dispatch capability of flexibility resources across different time scales [18], while in the spatial dimension, the modeling of virtual energy storage plants that integrate frequency regulation with congestion management reflects the necessity of jointly considering frequency services and network constraints [19]. The above studies have expanded the modeling scope of frequency security dispatch from the dual perspectives of resource types and system structure; however, they have also brought the trade-off between uncertainty, time scale, and solvability into sharper relief.

In research on combined thermal-storage frequency regulation, energy storage not only provides rapid response capabilities but may also alter the operating modes and reserve structures of thermal power plants, thereby affecting the overall economic efficiency and frequency security of the system. Reference [20] proposes a system optimization approach that replaces governor reserves with energy storage and quantifies the impact of energy storage on the operating costs and frequency support capabilities of conventional units. To enhance the realism of energy storage capacity modeling, megawatt-scale battery energy storage power curves are used to characterize the boundaries of available power as a function of state of charge (SoC) and to support the joint evaluation of frequency regulation and trading [21]. Regarding participation in the frequency regulation market, a dynamic bidding model for hybrid energy storage systems participating in intraday frequency regulation markets reveals the coupling relationship between SoC constraints and profit maximization [22]. Modeling the participation of flexibility resources, such as battery swap stations, in the frequency regulation market under conditions of uncertainty also provides a systematic evaluation pathway for new forms of energy storage [23]. In terms of multi-service coupling, energy storage primarily addresses the coordination of services across different time scales in research on the integrated application of peak shaving and frequency regulation [24], while the feasibility and benefit trade-offs of frequency regulation under arbitrage planning constraints reveal the constraining role of energy budgets on frequency regulation compliance [25]. A significant body of research has already incorporated energy storage systems into the frequency regulation framework. Reference [26] incorporates the uncertainty in the primary availability of energy storage into the frequency constraint (UFC) to characterize the impact of energy storage performance risks on the satisfaction of frequency security constraints. Reference [27] optimizes the configuration of frequency regulation energy storage systems based on an equivalent full-cycle model, demonstrating how cycle costs constrain configuration and operational strategies; furthermore, site selection optimization from the perspective of frequency stability can also improve frequency regulation efficiency [28]. The recent economic dispatch model for shared energy storage stations provides a computational framework for user-side shared energy storage participation in system dispatch [29], while the incorporation of frequency constraints into autonomous system planning indicates that frequency security has been integrated across both operational and planning levels [30]. At the market–system interface, the two-stage robust clearing model for energy and reserve in wind–fire systems proposed in Reference [31] highlights the impact of joint energy and reserve clearing on system safety and economic efficiency under conditions of uncertainty. Reference [32] proposes an incentive-compatible market mechanism for primary frequency regulation ancillary services and incorporates nadir and quasi-steady-state frequency deviation constraints into the clearing process, thereby providing a mechanism foundation for the fair settlement of various types of frequency resources. The chain-based framework for distributed energy storage aggregators participating in the frequency regulation market covers key aspects such as daily quotes and intraday capacity allocation [33], while the characterization of VPP feasibility regions and joint market bidding methods provides more comprehensive modeling tools for aggregating resources to deliver frequency regulation capabilities at the system level within this framework [34]. While the aforementioned studies have proposed feasible solutions in areas such as frequency constraint modeling, resource allocation, and market organization, most of them still exhibit a certain degree of disjointedness in terms of the time scale. Specifically, there is insufficient feedback coupling between day-ahead commitments and real-time adjustments, making it difficult to ensure that the frequency support commitments of thermal power and energy storage remain both economically

viable and feasible under conditions where renewable energy forecast errors and credible disturbances coexist.

Based on a review of the relevant studies mentioned above, we can identify common issues that require further resolution in the optimization of combined thermal-storage frequency regulation systems under conditions of high renewable energy penetration. First, the formulation of frequency security constraints at the day-ahead UC level must establish a consistent balance between physical realism and computational solvability; otherwise, it may lead to safety-driven conservative power generation and an increase in curtailed renewable energy. Second, the frequency regulation capacity of energy storage is constrained by SoC coupling and energy budget limitations. Its ability to consistently meet commitments under various scenarios requires coordinated optimization with the start-stop operations and reserve structure of thermal power plants to avoid situations where commitments are feasible in the day-ahead market but unfeasible in real time. Third, existing research often links day-ahead planning with real-time frequency regulation deployment through rule-based or weakly coupled approaches in many scenarios, lacking systematic modeling that integrates day-ahead commitments with real-time corrections within a unified optimization framework. To address these issues, this paper proposes a two-stage stochastic frequency-security-constrained unit commitment model, FSC-SUC. The proposed model and method have three main innovations. First, a two-stage stochastic framework covering day-ahead commitments and real-time corrections is established. This framework incorporates renewable energy forecast errors and credible disturbances in the form of scenarios, and simultaneously constrains frequency safety indicators such as RoCoF, frequency nadir, and quasi-steady-state frequency deviation across all scenarios. Second, we integrate energy storage SoC coupling, frequency regulation capacity commitments, and thermal power plant start-up and ramping constraints into a unified collaborative optimization process, allowing the division of labor between thermal power and energy storage for slow and fast frequency regulation tasks to emerge endogenously from the optimization process. Third, we propose a solvable modeling approach that combines analytical frequency constraints with data-driven linear approximation, along with a decomposition-based solution strategy to improve computational efficiency across multiple scenarios, thereby achieving a more reasonable trade-off between frequency security and renewable energy integration. The remainder of this paper is organized as follows. Section 2 presents the modeling of the system and frequency security indicators and constructs the two-stage stochastic FSC-SUC model. Section 3 proposes a decomposition-based solution method. Section 4 verifies the effectiveness and advantages of the proposed method through case studies analyzing different renewable energy penetration rates and disturbance scenarios. Section 5 concludes and provides an outlook.

2. Two-Stage Stochastic Frequency and Security Constrained Unit Commitment for Thermal-Storage Joint Frequency Regulation

This section develops a two-stage stochastic frequency-security-constrained unit commitment model for high-renewable-penetration systems, in which thermal units and energy storage jointly provide frequency support. The overall approach of the model centers on a hierarchical structure comprising day-ahead planning and real-time correction. During the day-ahead phase, the model simultaneously determines the start-up and shut-down status and reference output of thermal power plants, as well as the reservation of slow frequency support capacity. It also establishes the reference charging and discharging plans for energy storage, the SoC trajectory, and the reservation of fast frequency support capacity. This transforms frequency support capacity from a resource that is called upon *ex post facto* into an assessable resource that can be committed during the planning pe-

riod and delivered during the execution period. During the real-time correction phase, scenario-based rescheduling is performed at 5 min intervals to address forecasting errors in renewable energy and load, thereby achieving error absorption and sustainable compliance with the SoC under commitment consistency constraints while also satisfying frequency safety criteria such as RoCoF, frequency nadir, and quasi-steady-state frequency deviation. To balance physical interpretability and computational feasibility, this chapter employs a strategy that combines analytical criteria with linear surrogate constraints for frequency security modeling. Specifically, RoCoF and quasi-steady-state frequency deviation are expressed as aggregated analytical constraints, while the frequency nadir is approximated using a polyhedral linear surrogate constructed from samples generated offline via a simplified frequency response model. This approach transforms the complex dynamic feasible region into a set of linear inequalities that can be embedded into a mixed-integer optimization model. Based on the modeling framework described above, this chapter will present, in sequence, the system architecture and time-scale settings; the model for unit commitment and frequency support capacity reservation during the day-ahead phase; the 5 min real-time correction and rescheduling model; and a unified method for embedding frequency security criteria and linear surrogate constraints. This provides a comprehensive mathematical foundation for subsequent decomposition-based solutions and case studies.

2.1. System Structure and Time-Scale Design

The study considers a power system consisting of a set of thermal units $\mathcal{G}$, a set of battery energy storage systems B , renewable generation, and system load. To better reflect the operation and dispatch practices of electricity markets, the day-ahead stage adopts an hourly time set T , while the real-time correction stage adopts a 5 min time set K . A mapping $t(k)$ is defined to indicate the hourly interval to which each 5 min interval k belongs. The two-stage time-scale setting is designed to meet two modeling objectives simultaneously. First, the discrete operational characteristics of thermal units, such as unit commitment and minimum up/down time constraints, imply that commitment decisions should not be modeled at a fine-grained time scale; otherwise, the number of binary variables would increase significantly, leading to a substantial rise in computational complexity. Second, the SoC energy budget constraint of energy storage exhibits strong temporal coupling. Under high penetration of renewable energy, forecast errors may accumulate over multiple 5 min intervals. If corrections are performed only at the hourly scale, situations may arise where energy balance is satisfied while the SoC trajectory becomes unsustainable. Therefore, the hierarchical structure with an hourly layer and a 5 min correction layer can control the scale of discrete decisions while realistically embedding the SoC dynamic constraints into the error realization process, thereby explicitly verifying the sustainable performance of joint frequency regulation by thermal units and energy storage.

System uncertainty is represented by a scenario set S . Each scenario $s \in S$ provides the 5 min resolution renewable generation forecast error trajectory $\varepsilon_{k,s}^{RES}$ and load forecast error trajectory $\varepsilon_{k,s}^D$, with an associated scenario probability π_s , satisfying $\sum_{s \in S} \pi_s = 1$. The credible disturbance set C is constructed according to the conventional N-1 criterion, which can be represented by the outage of the largest online generating unit or an equivalent maximum active power deficit event, denoted as ΔP_c . In terms of model semantics, error scenarios are mainly used to characterize the real-time energy balance pressure and the sustainability of the SoC trajectory, while credible disturbances are used to establish the verification boundary of frequency security constraints. Together, they determine the effectiveness of frequency security constraints and the model's coverage of risks under high penetration of renewable energy. The coupling mechanism of the two-stage decision can be summarized as the consistency between commitment and execution. In the first stage, at the hourly scale,

the model determines the unit commitment status $u_{g,t}$, baseline output $p_{g,t}$, and reserved primary frequency response capacity $r_{g,t}^{\text{PFR}}$, as well as the baseline charging and discharging schedules of energy storage $p_{b,t}^{\text{ch}}, p_{b,t}^{\text{dis}}$, the SoC trajectory $e_{b,t}$, and the reserved fast frequency response capacity $r_{g,t}^{\text{FFR}}$. In the second stage, at the 5 min scale, redispatch is performed for each scenario. The adjustment ranges of $p_{g,k,s}$ and $p_{b,k,s}^{\text{ch}}, p_{b,k,s}^{\text{dis}}$ must be constrained by the reserved capacities determined in the first stage, thereby ensuring that the primary and fast frequency response capabilities are not arbitrarily deployed ex post but are explicitly defined in the planning stage and strictly delivered during execution. This consistency loop is the key to incorporating thermal unit primary frequency regulation and energy storage fast frequency regulation into a unified system optimization framework.

The structural diagram of the two-stage stochastic unit commitment framework with joint frequency regulation by thermal units and energy storage remains unchanged in its overall design. The input layer consists of data and parameters, the optimization layer includes the day-ahead master problem and scenario-based subproblems, and the output layer provides economic performance, security metrics, and resource stress statistics. The connection from the master problem to the subproblems represents commitment transfer, while the connection from the subproblems to the master problem corresponds to cut generation and feedback. The detailed block diagram is shown in Figure 1.




Figure 1. Framework of two-stage stochastic frequency-secure UC.

2.2. Day-Ahead Stage Model

In the day-ahead stage, the objective is to minimize the day-ahead operating cost and the expected real-time correction cost. To facilitate two-stage decomposition and maintain the structural stability of the master problem, an outer-layer objective formulation is adopted:

$$\min C^{\text{DA}} + \eta \quad (1)$$

Here, η is a proxy variable for the expected real-time correction cost, which is progressively approximated by cut planes within the decomposition framework, reflecting the marginal impact of different commitment decisions on real-time adjustment costs. C^{DA} consists of start-up and shut-down costs, fuel costs, reserved capacity costs for primary and fast frequency response, energy storage operating costs, and renewable curtailment penalties. It is formulated in a linear or piecewise linear form to ensure that the master problem remains a mixed-integer linear program. This design avoids the direct inclusion of high-dimensional scenario variables in the master problem. Instead, the

impact of uncertainty is gradually incorporated through η and cut planes during the iterative process, which is consistent with the standard modeling paradigm of two-stage stochastic programming.

$$\sum_{g \in \mathcal{G}} p_{g,t} + \sum_{b \in \mathcal{B}} (p_{b,t}^{\text{dis}} - p_{b,t}^{\text{ch}}) + P_t^{\text{RES,DA}} - c_t = D_t^{\text{DA}}, \forall t \quad (2)$$

Equation (2) describes the energy balance at the hourly scale. To avoid introducing significant uncertainty in the day-ahead stage due to forecast deviations, this study incorporates error terms and performs redispatch in the real-time stage, while the day-ahead stage forms the baseline schedule solely based on forecast values. The introduction of the curtailment variable c_t ensures that the day-ahead model remains feasible when renewable generation is significantly in surplus and thermal units are constrained by minimum output limits. Meanwhile, it provides a statistical basis for subsequent analysis of renewable energy accommodation capability.

The operational constraints of thermal units are modeled using a standard UC formulation. The output limits and start-up/shut-down logical constraints are expressed as:

$$P_g^{\min} u_{g,t} \leq p_{g,t} \leq P_g^{\max} u_{g,t}, \forall g, t \quad (3)$$

$$u_{g,t} - u_{g,t-1} = y_{g,t} - z_{g,t}, \forall g, t \quad (4)$$

Here, $u_{g,t} \in \{0, 1\}$ indicates whether unit g is online, while $y_{g,t}$ and $z_{g,t}$ denote start-up and shut-down actions, respectively. The minimum up/down time constraints are given as:

$$\sum_{\tau=t-U_g^{\min}+1}^t y_{g,\tau} \leq u_{g,t}, \forall g, t \quad (5)$$

$$\sum_{\tau=t-D_g^{\min}+1}^t z_{g,\tau} \leq 1 - u_{g,t}, \forall g \quad (6)$$

The ramping constraints at the hourly scale are given as:

$$p_{g,t} - p_{g,t-1} \leq R_g^{\uparrow} u_{g,t-1} + S_g^{\uparrow} y_{g,t}, \forall g, t \quad (7)$$

$$p_{g,t-1} - p_{g,t} \leq R_g^{\downarrow} u_{g,t} + S_g^{\downarrow} z_{g,t}, \forall g, t \quad (8)$$

Equations (3)–(8) characterize the operational characteristics of thermal units under constraints of minimum output, ramping limits, and start-up/shut-down logic. Since the provision of primary frequency response by thermal units depends on their adjustable margins and ramping capabilities, the above operational constraints determine the actual feasible boundary of primary frequency support and also define the opportunity cost of maintaining synchronous units online under high renewable penetration scenarios.

The primary frequency regulation capability of thermal units is represented by the reserved primary frequency response margin:

$$0 \leq r_{g,t}^{\text{PFR}} \leq P_g^{\max} u_{g,t} - p_{g,t}, \forall g, t \quad (9)$$

Equation (9) directly couples primary frequency support with the baseline output, implying that a higher baseline output results in a smaller upward regulation reserve available for frequency response.

This constraint, on the one hand, ensures the physical feasibility of frequency regulation capability commitments; on the other hand, it provides a structural basis for cost representation of primary frequency support, since increasing $r_{g,t}^{\text{PFR}}$ typically requires either

reducing the baseline output or keeping more units online, thereby affecting fuel costs and start-up and shut-down costs.

The energy storage constraints are used to simultaneously characterize power limits and energy budget constraints. To avoid simultaneous charging and discharging within the same time interval, a binary variable $V_{b,t}$ is introduced.

$$0 \leq p_{b,t}^{\text{ch}} \leq P_b^{\text{ch,max}} v_{b,t}, 0 \leq p_{b,t}^{\text{dis}} \leq P_b^{\text{dis,max}} (1 - v_{b,t}) \quad (10)$$

The SoC dynamics and boundary constraints of the energy storage system are given as:

$$e_{b,t+1} = e_{b,t} + \eta_b^{\text{ch}} p_{b,t}^{\text{ch}} \Delta T - \frac{1}{\eta_b^{\text{dis}}} p_{b,t}^{\text{dis}} \Delta T, \forall b, t \quad (11)$$

$$E_b^{\text{min}} \leq e_{b,t} \leq E_b^{\text{max}}, \forall b, t \quad (12)$$

Equations (10)–(12) enable the energy storage system to form a feasible baseline charging/discharging schedule and SoC trajectory in the day-ahead stage. Since the fast frequency response capability in the real-time stage heavily depends on the available SoC margin, the feasibility of the day-ahead SoC trajectory directly affects the deliverability of fast frequency response commitments. Describing fast frequency support solely through power limits neglects the constraining role of SoC dynamics; therefore, this study incorporates SoC constraints as a core component of the fast frequency capability budget. The reserved fast frequency response capability of energy storage is subject to both power and energy budget constraints. To ensure deliverability over the fast frequency response duration τ_t^{FFR} , the following formulation is constructed.

$$0 \leq r_{b,t}^{\text{FFR}} \leq P_b^{\text{dis,max}} (1 - v_{b,t}) - p_{b,t}^{\text{dis}} + p_{b,t}^{\text{ch}}, \forall b, t \quad (13)$$

$$r_{b,t}^{\text{FFR}} \leq \frac{e_{b,t} - E_b^{\text{min}}}{\tau_t^{\text{FFR}}}, \forall b, t \quad (14)$$

Equation (13) characterizes the instantaneous power availability of fast frequency response, while Equation (14) represents the energy sustainability of fast frequency response over a given duration. Together, they form the capability envelope of fast frequency support, allowing fast frequency commitments to be dynamically adjusted with the SoC state. The introduction of this capability envelope transforms fast frequency support from a conventional fixed-capacity variable into a deliverable capacity subject to energy budget constraints, thereby providing a consistent commitment boundary for subsequent 5 min scenario-based redispatch.

2.3. Five-Minute Real-Time Correction Model

In the real-time correction stage, redispatch is performed at a 5 min resolution for each scenario to accommodate forecast errors and maintain consistency with the day-ahead commitments. For scenario s and time interval k , the power balance constraint is given as:

$$\sum_g p_{g,k,s} + \sum_b (p_{b,k,s}^{\text{dis}} - p_{b,k,s}^{\text{ch}}) + P_k^{\text{RES,DA}} + \varepsilon_{k,s}^{\text{RES}} - c_{k,s} = D_k^{\text{DA}} + \varepsilon_{k,s}^D - \ell_{k,s}, \forall k \quad (15)$$

Equation (15) explicitly incorporates forecast errors into the real-time power balance, making redispatch variables responsible for absorbing these deviations. The variables $c_{k,s}$ and $\ell_{k,s}$ have two meanings. First, as a feasibility recovery mechanism for extreme scenarios, they ensure that the subproblem remains solvable and can return comparable objective values. Second, the high penalty coefficient of $\ell_{k,s}$ explicitly quantifies infeasibility in the objective function, thereby driving the day-ahead stage through cut-plane feedback

to increase reserved capacity or adjust the SoC trajectory, forming an endogenous response to extreme risks. Thermal units satisfy output limits and ramping constraints at the 5 min time scale:

$$p_g^{\min} u_{g,t(k)} \leq p_{g,k,s} \leq p_g^{\max} u_{g,t(k)}, \forall g, k \quad (16)$$

$$p_{g,k,s} - p_{g,k-1,s} \leq R_g^{\uparrow,5m}, p_{g,k-1,s} - p_{g,k,s} \leq R_g^{\downarrow,5m} \quad (17)$$

The redispatch range is constrained by the reserved capacity determined in the day-ahead stage:

$$-r_{g,t(k)}^{\text{PFR}} \leq p_{g,k,s} - p_{g,t(k)} \leq r_{g,t(k)}^{\text{PFR}}, \forall g, k, s \quad (18)$$

Equation (18) concretizes the primary frequency response commitment into a real-time dispatchable regulation boundary. This boundary constraint prevents arbitrary expansion of the thermal units' adjustment range in the real-time stage, thereby ensuring the auditability of primary frequency response commitments and making the day-ahead decision on $r_{g,t}^{\text{PFR}}$ directly impact real-time feasibility.

The power mutual exclusivity, SoC update, and boundary constraints of energy storage at the 5 min time scale are given as:

$$0 \leq p_{b,k,s}^{\text{ch}} \leq p_b^{\text{ch,max}}, 0 \leq p_{b,k,s}^{\text{dis}} \leq p_b^{\text{dis,max}} \quad (19)$$

$$e_{b,k+1,s} = e_{b,k,s} + \eta_b^{\text{ch}} p_{b,k,s}^{\text{ch}} \Delta\tau - \frac{1}{\eta_b^{\text{dis}}} p_{b,k,s}^{\text{dis}} \Delta\tau \quad (20)$$

$$E_b^{\min} \leq e_{b,k,s} \leq E_b^{\max}, \forall b, k, s \quad (21)$$

where $\Delta\tau = 5/60$ h. To ensure the executability of fast frequency response commitments, the increment of energy storage discharging is kept consistent with the reserved fast frequency capacity determined in the day-ahead stage.

$$0 \leq p_{b,k,s}^{\text{dis}} - p_{b,t(k)}^{\text{dis}} \leq r_{b,t(k)}^{\text{FFR}}, \forall b, k, s \quad (22)$$

Equation (22) ensures that the deployment of fast frequency response from energy storage does not exceed the day-ahead commitment, and together with Equation (14), it guarantees that the SoC margin can sustain the required fast frequency response duration. Under scenarios with persistent forecast bias, the SoC dynamics gradually tighten the available fast frequency boundary, allowing the model to capture the sustainable regulation capability and operational limits of energy storage under high renewable penetration.

2.4. Simplified Frequency Response Model and Linear Proxy Constraint for Frequency Nadir

The frequency security criteria in this study include RoCoF, frequency nadir, and steady-state frequency deviation. To balance tractability and physical consistency, RoCoF and steady-state frequency deviation are modeled using aggregated analytical constraints, while the nadir is embedded via an outer-approximation approach based on linear proxy constraints. The core idea of this strategy is to transform the complex nonlinear feasible region of frequency dynamics into a set of linear inequalities, enabling compatibility with the mixed-integer structure of unit commitment and the dual-based cut generation mechanism in Benders decomposition. At the same time, safety margins are calibrated to ensure auditable control of approximation errors. Offline sample generation and labeling are based on a simplified system frequency response model. The frequency deviation $\Delta f(t)$ satisfies the swing equation:

$$\frac{2H_t}{f_0} \frac{d\Delta f(t)}{dt} = \Delta P_m(t) - \Delta P_e(t) - D_t \Delta f(t) \quad (23)$$

where H_t denotes the equivalent inertia and D_t denotes the equivalent damping. $\Delta P_e(t)$ includes the credible disturbance deficit ΔP_c and the equivalent term of forecast errors.

The primary frequency response of thermal units is approximated by a first-order dynamic model:

$$T_g \frac{dP^{\text{PFR}}(t)}{dt} + P^{\text{PFR}}(t) = K_g \Delta f(t), 0 \leq P^{\text{PFR}}(t) \quad (24)$$

The fast frequency response of energy storage is modeled using a simplified delayed ramp-up representation:

$$P^{\text{FFR}}(t) = \begin{cases} 0, & t < \tau_d \\ \min(\kappa(t - \tau_d), R_t^{\text{FFR}}), & t \geq \tau_d \end{cases} \quad (25)$$

Within the window $[0, T_w]$, numerical integration of Equations (23)–(25) is performed to obtain $\Delta f(t)$, and the frequency nadir deviation Δf^{nadir} is extracted as the training label. The simplified frequency response model generates training samples that cover the operating region and reveal the boundary relationship between aggregated features and frequency nadir, rather than replacing high-fidelity simulation. We set the parameters based on engineering experience, and their robustness is further validated through sensitivity analysis in the case studies.

In the online optimization, an aggregated feature vector is defined:

$$x_t = [H_t, R_t^{\text{PFR}}, R_t^{\text{FFR}}, \Delta P_c]^T \quad (26)$$

where the RoCoF constraint is formulated using an equivalent lower bound:

$$H_t = \sum_{g \in \square} u_{g,t} H_g, R_t^{\text{PFR}} = \sum_g r_{g,t}^{\text{PFR}}, R_t^{\text{FFR}} = \sum_b r_{b,t}^{\text{FFR}} \quad (27)$$

$$H_t \geq \Gamma(\Delta P_c), \forall t, c \quad (28)$$

where the quasi-steady-state frequency deviation constraint is expressed by the lower bound of the total primary frequency regulation:

$$R_t^{\text{PFR}} \geq \alpha \Delta P_c, \forall t, c \quad (29)$$

where $\Gamma(\cdot)$ and α are determined by parameters such as frequency thresholds, nominal frequency, and equivalent damping, maintaining linear tractability and establishing an explicit mapping with the disturbance magnitude.

The frequency nadir constraint is approximated using a polyhedral linear outer approximation:

$$w_m^\top x_t + b_m \geq \rho_m, \forall t, c, m \in \mu \quad (30)$$

where w_m and b_m denote the hyperplane parameters, and ρ_m represents the safety margin. The detailed hyperplane coefficients are provided in Appendix A. The polyhedral outer approximation can improve boundary tightness by increasing the number of hyperplanes or by partitioning the feature space for piecewise fitting. Meanwhile, the safety margin is calibrated using high-quantile residuals from a validation set to maintain statistical conservatism. This set of constraints, together with Equations (28) and (29), forms the frequency security constraint system, enabling frequency security requirements to be incorporated into the day-ahead master problem in a linear form and providing an interpretable constraint-driven mechanism for unit commitment decisions, primary frequency reserve allocation, and fast frequency reserve allocation.

3. Training of Linear Surrogate Constraints and Two-Stage Decomposition Method

3.1. Implementation Essentials of Scenario Generation and Reduction

The construction of the scenario set S determines the extent to which the two-stage model covers uncertainty. The scenarios in this paper are formed by combining error trajectories with credible disturbances. The forecast error scenarios are generated from historical data of a regional power system with a 5 min resolution over 365 days. The raw data consists of 105,120 data points ($365 \times 24 \times 12$) for both wind and PV generation. The prediction errors are modeled using a Gaussian mixture model (GMM) with three components to capture the asymmetric error distribution. The temporal correlation is preserved by sampling full-day error trajectories rather than independent time points. After scenario reduction using the k-means clustering algorithm with $k = 30$, the retained scenarios capture 92% of the total variance of the original dataset. The extreme scenarios (top 5% of absolute error trajectories) are mandatorily retained to ensure tail risk coverage. This approach keeps the scenario set size within a computable range without significantly compromising risk coverage. In this study, error trajectories are generated by sampling historical prediction-error data and constructing an initial scenario library that preserves the temporal characteristics of renewable and load forecast errors. Given that the 5 min correction sequences are long and the initial number of scenarios is typically large, a direct solution would result in a significant increase in the size of subproblems and the parallel processing load; therefore, scenario reduction is necessary. The goal of scenario reduction is to retain a set of scenarios that can represent both mean characteristics and tail risk, while preserving the temporal correlation of the error sequences as much as possible. In practice, we use k-means clustering to group error trajectories, and a central trajectory can be selected from each group to represent the typical scenario, with the probability weight determined by the proportion of samples within the group. To avoid losing tail risks after reduction, we establish retention rules for extreme samples in the error sequence. For example, upper and lower deviation trajectories at high percentiles can be mandatorily retained, and the probability weights can be renormalized. This approach keeps the scenario set size within a computable range without significantly compromising risk coverage, thereby ensuring the iterative efficiency and result stability of the decomposition solution.

The credible disturbance set C is constructed using the conventional N-1 principle. In practice, the maximum unit trip or the equivalent maximum active power shortfall event is often selected as the reference disturbance, and several typical contingencies may be further incorporated to enhance robustness. The disturbance magnitude ΔP_c is incorporated into the RoCoF, quasi-steady-state frequency deviation, and nadir proxy constraints, establishing a clear mapping between frequency safety constraints and accident magnitude. This imposes explicit constraints on the frequency support capacity demand during the planning phase.

3.2. Offline Training Process and Technical Details for Linear Surrogate Constraints

The training of linear surrogate constraints involves sample generation, frequency response calculation, feature engineering, hyperplane fitting, and safety margin calibration. The sample generation stage must cover all possible combinations of inertia levels and frequency support capacity within the feasible region of the main problem. Specifically, the inertia H_t is determined by the set of online thermal power generation units and takes discrete values, while RFR_{slow} and RFR_{fast} are continuous variables subject to the constraints in Equations (9), (13) and (14). To balance coverage and training efficiency, we adopt a hierarchical sampling strategy. The training dataset consists of 50,000 samples generated using the hierarchical sampling strategy described above. Specifically, 20 representative

unit-commitment combinations are enumerated to cover the feasible range of equivalent inertia, and 2,500 Latin-hypercube samples are generated for each combination in the continuous space of PFR and FFR capacity. An independent validation set of 10,000 samples is used for accuracy assessment and safety-margin calibration.

First, at the start–stop state level, several typical unit combinations are enumerated or sampled to obtain representative H_t levels. Then, at each H_t level, Latin hypercube sampling is performed on RFR_{slow} , RFR_{fast} , and ΔP_c to ensure that the samples are uniformly distributed in the continuous space. To improve the accuracy of the surrogate constraints in the critical region, denser sampling can be performed near the boundaries where the frequency threshold is met or approached, focusing the training on identifying the boundaries of the feasible region rather than redundantly covering internal points.

The frequency response calculation stage employs the simplified frequency response model defined by Equations (23)–(25). Physically interpretable ranges must be specified for the parameters T_g , K_g , D_t , τ_d , κ , and T_w . T_g and K_g can be derived from the first-order dynamics and droop characteristics of a typical thermal power plant speed control system; D_t can be set based on empirical load damping values and verified for its impact through sensitivity analysis; τ_d and κ can be set based on the response delay and ramping capability of the energy storage converter; and T_w is set to the time window during which the frequency nadir primarily occurs. The parameters of the simplified frequency response model are set based on typical values from the literature and are summarized in Table 1. For each sample point, $\Delta f(t)$ is integrated (or the nadir is extracted directly from the simulated frequency trajectory), and Δf_{nadir} is compared with the threshold Δf_{min} to obtain a feasibility label. This process generates a supervised learning dataset, with x_t as the input and feasibility (feasible or infeasible) or the regression value of Δf_{nadir} as the output.

Table 1. Parameters of the simplified frequency response model.

Parameter	Symbol	Value	Unit	Description
Governor time constant	T_g	0.2	s	Turbine-governor response delay
Governor gain	K_g	20	pu/pu	Droop gain (inverse of droop setting)
Load damping coefficient	D_t	1.0	pu/Hz	Frequency-sensitive load fraction
BESS response delay	τ_d	0.1	s	Converter response time
BESS ramp rate	κ	1.5	pu/s	Fast power injection rate
Nadir observation window	T_w	5	s	Time window for nadir detection

T_g and K_g are derived from the first-order dynamics and droop characteristics of a typical thermal power plant speed control system; D_t is set based on empirical load damping values; τ_d and κ are set based on the response delay and ramping capability of the energy storage converter; and T_w is set to the time window during which the frequency nadir primarily occurs.

To validate the simplified frequency response model (Equations (23)–(25)), we compared its output with a high-fidelity simulation model implemented in MATLAB/Simulink R2023b, which includes detailed turbine-governor dynamics, AGC loops, and converter response models. The comparison was conducted over 1000 randomly sampled operating points covering the feasible region.

The results indicate that the simplified model slightly underestimates the frequency nadir by an average of 0.03 Hz (relative error < 3%) compared to the high-fidelity simulation. This underestimation is conservative from a security perspective (i.e., it slightly overestimates the risk) and is accounted for by the safety margin ε calibrated from the validation set. The correlation coefficient between the simplified model and high-fidelity

simulation is 0.98, confirming that the simplified model captures the dominant dynamics sufficiently for surrogate training purposes.

In the hyperplane fitting stage, the principle of outward approximation is adopted. If a classification-based fitting method is used, the objective is to construct a set of linear inequalities that places infeasible points as far to the left of the hyperplane as possible and feasible points to the right, while imposing a significant penalty for misclassifying infeasible points to ensure conservativeness. If a regression-based fitting method is adopted, $\Delta f_{nadir} = \theta_{x_t}^T + \theta_0$ is first fitted, and then a linear constraint is formed using $\Delta f_{nadir} \geq \Delta f_{\min} + \rho$. To improve approximation accuracy, we partition the feature space, with hyperplanes fitted separately at different H_t levels or within different ΔP_c intervals, thereby forming a polyhedral outward approximation and obtaining the constraint set in Equation (30). This partitioned fitting can significantly alleviate the issues of excessive conservativeness or local inaccuracies caused by a single hyperplane in global approximation. The fitted surrogate achieves an R^2 of 0.96 on the validation set, with an RMSE of 0.042 Hz and an MAE of 0.031 Hz. The validation classification accuracy is 97.3%, and the false-negative rate after safety-margin calibration is 0.8%. These indicators demonstrate that the surrogate is sufficiently accurate for screening frequency-nadir security while retaining the conservativeness required by the UC model.

$$x_t = [H_t, R_t^{PFR}, R_t^{FFR}, \Delta P_c]^T \quad (31)$$

$$\Delta f_{nadir}, t^{pred} = \theta_{x_t}^T + \theta_0 \quad (32)$$

$$\Delta f_{nadir}, t^{pred} \geq \Delta f_{\min} + \rho \quad (33)$$

$$w_m^T x_t + b_m \geq \rho_m, m = 1, 2, \dots, M \quad (34)$$

The safety margin calibration phase is performed using an independent validation set. We calculate the surrogate prediction error or classification margin for samples in the validation set, and a high percentile is selected as the value of ρ_m , thereby statistically ensuring that the linear surrogate remains an outer approximation of the true feasible region. This calibration method offers clear auditability, explains the source of conservatism in the surrogate constraints, and allows for a trade-off between conservatism and efficiency in subsequent applications by adjusting the percentile level.

3.3. Master and Subproblem Formulation for Two-Stage Decomposition

The two-stage stochastic FSC-SUC has a structure comprising a first-stage mixed-integer decision-making process and multi-scenario continuous rescheduling. To improve solution efficiency at large scales, this paper employs Benders decomposition. The master problem corresponds to the day-ahead stage model, which includes Equations (2)–(14) and the frequency security constraints (28) through (30), and introduces η to denote the expected real-time correction cost. At each iteration, the master problem outputs the day-ahead decision z , where z uniformly represents $(u, p, r_{PFR}, p_{ch}, p_{dis}, e, r_{FFR}, c)$.

For each scenario subproblem, after the master problem decisions are fixed, the 5 min rescheduling model in Equations (15)–(22) is solved. The subproblem is a linear programming problem, and its optimal value $Q_s(z)$ represents the minimum adjustment cost and penalty that satisfies the power balance, ramping, SoC dynamics, and capacity consistency constraints under scenario s . Since the subproblems are independent across different scenarios, they can be solved in parallel, meaning that computational efficiency is primarily influenced by the number of scenarios and the scale of each individual subproblem. The dual variables of the subproblems provide marginal sensitivity information regard-

ing the master problem's decisions, which can be used to generate cutting planes and progressively approximate η .

Under the multi-cut strategy, an optimal cutting plane can be constructed for each iteration i , such that the lower bound η on the current solution is progressively tightened. The cutting plane can be written as:

$$\eta \geq \sum_{s \in S} \pi_s \left[Q_s(z^i) + \lambda_s^{(i)T} (z - z^i) \right] \quad (35)$$

where $\lambda_s^{(i)}$ is the projection of the subproblem dual vector onto the coupling variables. This form captures the local linearization of the real-time correction cost with respect to the day-ahead decisions, thereby encouraging the master problem in subsequent iterations to prefer start-up and reserve schemes that are more conducive to reducing the correction cost.

If a subproblem is infeasible in the strict sense, this paper employs a feasibility mechanism to ensure that the subproblem still returns a finite objective value and provides a direction for improvement. In the implementation, penalty terms for load reduction $\ell_{k,s}$ and additional power curtailment $c_{k,s}$ ensure that Equation (15) remains globally feasible. Meanwhile, we set the penalty coefficients significantly higher than other control costs, making infeasibility explicitly reflected in the objective function. This, in turn, drives the master problem to increase reserve capacity or adjust the SoC trajectory under cutting-plane feedback to restore strict feasibility. This strategy avoids the implementation complexity associated with dual-ray solving and enhances the algorithm's stability in multi-scenario parallel environments.

Convergence criteria may be based on the relative gap between the objective value of the master problem and the lower bound obtained from the subproblems or on a stopping condition where the improvement in the lower bound η falls below a threshold over several consecutive iterations. To improve convergence speed, looser surrogate constraints or a smaller set of hyperplanes can be used in the initial stage to quickly obtain a feasible solution. The precision of the surrogate constraints and the number of hyperplanes can then be gradually increased to tighten the boundaries, thereby forming a step-by-step approximation process from coarse to fine.

The flowchart of linear surrogate constraint training and the two-stage decomposition solution is shown in Figure 2. It includes scenario construction and reduction, linear surrogate training and margin calibration, Benders master problem and parallel subproblem iterations, and convergence output.

3.4. Implementation and Reproducibility Remarks

To ensure model reproducibility, the implementation is organized into three parts: system physical parameters, algorithmic parameters, and the decoupled offline surrogate-training module. First, all parameters are categorized into two types based on their source: system physical parameters and algorithm parameters. System physical parameters include unit output limits, ramping limits, minimum up/down times, inertia constants, energy storage power, energy limits and efficiency, and SoC safety limits. Algorithm parameters include the number of scenarios and reduction strategies, the sample size and partitioning method for linear surrogate training, the safety margin percentile, and the termination threshold and maximum number of iterations for the decomposition iteration. Second, the surrogate constraint training module is decoupled from the two-stage solution module. Surrogate training is performed offline and outputs a set of linear inequality coefficients $\{w_m, b_m, \rho_m\}$ that can be directly embedded into the optimization model. During the online solution stage, only this set of coefficients is read, thereby ensuring that the model solution does not depend on additional simulation calls. Third,

in terms of computational implementation, the scenario subproblems can be solved in parallel in a multi-core or cluster environment. The scale of the subproblems is primarily determined by $|r|$, $|G|$, and $|B|$. Using a 5 min correction interval allows the subproblem scale to be kept within an acceptable range while reflecting the granularity of market scheduling. Fourth, frequency thresholds and parameter selections should align with engineering codes. The settings for the RoCoF threshold, minimum frequency threshold, and quasi-steady-state frequency deviation threshold directly affect the compactness of Equations (28)–(30); therefore, necessary sensitivity analyses should be conducted in the case studies to avoid the over-reliance of conclusions on a single parameter setting.




Figure 2. Flowchart of linear surrogate constraint training and two-stage decomposition solution.

4. Case Studies


4.1. System and Data Setup

This chapter uses the IEEE 33-bus distribution network as a case study platform to verify the feasibility of the two-stage stochastic frequency-security-constrained unit commitment framework in mitigating frequency risks and providing fast and slow frequency support under high renewable energy penetration. The IEEE 33-bus system is used as the primary mechanism-verification platform because the frequency security constraints are expressed through aggregated inertia and frequency-response capability rather than through topology-specific parameters. The case study has a 24 h scheduling horizon; in the day-ahead stage, start-up and reserve commitment configurations are performed at a 1 h resolution, while in the real-time stage, scenario-based corrections for prediction errors are applied at a 5 min resolution. The frequency security constraints cover three types of indicators: RoCoF, frequency nadir, and quasi-steady-state frequency deviation. The frequency nadir is incorporated into the day-ahead master problem via the linear surrogate outer approximation method developed in Chapter 3, while the fast frequency support capability of energy storage is constrained by both power limits and SoC energy budgets. The system resource deployment and types are shown in Table 2.

Table 2. System resource deployment and parameters.

Resource	Bus	Type	Pmax_MW	Pmin_MW	H_s
DG1	1	Synchronous DG	2	0.3	6
DG2	18	Synchronous DG	1.6	0.25	5
DG3	30	Synchronous DG	1.2	0.2	4
BESS	24	Battery storage	1.6	NaN	NaN
PV	18	PV	NaN	NaN	NaN
Wind	30	Wind	NaN	NaN	NaN

Load and renewable energy are modeled using typical daily normalized curves scaled according to system size. Load exhibits steep ramp rates during morning and evening peak periods, while solar power reaches its peak at noon and wind power output fluctuates gradually throughout the day. This temporal pattern results in a potential trough in net load around noon and a sharp rise during the evening peak, thereby creating significant differences in inertia and frequency regulation requirements within a single day. The renewable energy output and system load are shown in Figure 3.

**Figure 3.** Renewable energy output and system load.

Real-time uncertainty is characterized using 5 min forecast error scenarios. The cumulative temporal effects of these errors alter the evolution of real-time correction power and the SoC of energy storage systems, thereby affecting fast frequency deliverability and the level of exposure to frequency nadir risk in tail-end scenarios. The generated uncertainty scenarios are shown in Figure 4.

The penetration rates for renewable energy were set at 0.30, 0.60, and 0.90, respectively. As the penetration rate increases, the net load at midday decreases significantly and approaches zero during certain periods, while the net load during the evening peak remains at a relatively high level. This makes the constraints imposed by the energy storage budget on the feasibility of fast frequency regulation even more pronounced at high penetration rates. The net load curves for different penetration rates are shown in Figure 5.




Figure 4. Forecast error scenarios. The colored lines denote different stochastic forecast error scenarios generated from historical data using a Gaussian mixture model.




Figure 5. Net load curves at different permeability levels.

4.2. Comparison Schemes and Evaluation Methods

To verify the impact of two-stage stochastic modeling, commitment consistency constraints, and SoC energy budget constraints on frequency security and deliverability, this chapter presents four comparative scenarios. The baseline scenario considers only energy balance and output constraints, without making an explicit commitment to frequency security. The frequency security-constrained unit selection scenario introduces RoCoF, quasi-steady-state frequency deviation, and minimum frequency constraints in the day-ahead stage but does not establish a deliverability closed-loop corresponding to scenario-based real-time correction. The two-stage stochastic schemes introduce 5 min scenario-based rescheduling in the real-time stage and use commitment consistency constraints to limit the magnitude of resource calls. One variant does not explicitly impose SoC energy budget constraints, while the other further incorporates the SoC energy budget into the fast frequency commitment boundary, thereby establishing a sustainable fast frequency support capability. The modeling differences among the schemes are summarized in Table 3.

Table 3. Modeling differences among the schemes.

Method	Freq Constraints	Two-Stage	SoC Budget	Key Feature
Base UC	0	0	0	Energy balance only
FSC UC	1	0	0	Day-ahead frequency constraints
FSC SUC noSoC	1	1	0	Two-stage with consistency, no SoC budgeting
FSC SUC	1	1	1	Two-stage with consistency and SoC budgeting

To quantify the actual benefit of the proposed FSC-SUC, we compare it against three state-of-the-art methods from the recent literature:

- Method A: Frequency-constrained UC with quadratic nadir constraint [7];
- Method B: Two-stage robust UC with frequency constraints [31];
- Method C: Frequency-constrained UC with learning-assisted RoCoF evaluation [4].

The comparison results are shown in Table 4.

Table 4. Performance comparison with state-of-the-art methods (high-penetration case).

Method	Total Cost (\$)	Nadir Violation Prob.	Load Shed (MWh)	Computation Time (s)
Method A	3,900	4.5%	0.002	650
Method B	4,050	1.2%	0	8,500
Method C	3,880	8.3%	0.01	1,200
FSC-SUC (Proposed)	3,920	0%	0	1,847

The results in Table 4 indicate that while Method A is computationally cheaper, it suffers from a higher violation probability. Method B provides comparable robustness but at a significantly higher computational cost due to the robust optimization formulation. The proposed method achieves the best security performance (zero violation probability) with a moderate and acceptable computational overhead.

The assessment is conducted comprehensively across three dimensions: economic viability, renewable integration, and frequency security. Economic viability is measured by intraday operating costs, incorporating costs associated with reserve allocation and penalties for extreme scenarios. Renewable integration is measured by the amount of curtailed renewable energy. Load loss risk is measured by the amount of load loss. Frequency security is primarily assessed using high-percentile statistics of RoCoF, frequency nadir, and quasi-steady-state deviation, with the probability of the frequency nadir exceeding the threshold introduced to characterize the level of tail risk exposure. All of the above statistics are calculated based on the 5 min scenario SoC evolution to determine fast frequency deliverability, ensuring that the impact of coupling fast frequency commitments with the energy budget on the risk assessment is explicitly reflected.

4.2.1. Validity of Linear Surrogate Constraints for Frequency Nadir

Within the threshold neighborhood, the surrogate output maintains a stable, monotonically consistent relationship with the frequency nadir in the dynamic simulation, as shown in Figure 6. In this figure, the dashed line represents the fitted linear surrogate boundary (decision threshold separating secure and insecure operating regions) after safety-margin calibration. This indicates that the surrogate is capable of capturing the dominant trends in nadir risk, as equivalent inertia and slow and fast frequency regulation capabilities vary. After introducing a safety margin, the deviation distribution of samples near the threshold converges toward the conservative side, thereby reducing the probability of misclassifying

unsafe operating points as safe. This provides an auditable basis for the linear surrogate outward approximation embedding strategy proposed in Section 3.




Figure 6. Surrogate validation: surrogate score vs. maximum frequency deviation.

Under a fixed disturbance magnitude, the required FFR capacity to meet the frequency nadir threshold decreases as the equivalent inertia increases, and shifts upward or downward with varying PFR levels, as illustrated in Figure 7. In this figure, the blue, green, and purple solid curves represent the required FFR capacity under low, medium, and high PFR levels, respectively. Additionally, the dashed line indicates the frequency nadir security boundary separating feasible and infeasible operating regions. This pattern reflects the division of labor between fast and slow frequencies during dynamic processes. Fast frequency components primarily suppress peak deviations during the initial phase of a disturbance, while slow frequency components determine the subsequent recovery slope and the persistence of deviations. Therefore, as slow frequency capacity increases, the dependence of the nadir constraint on fast frequency components is systematically reduced; the fast frequency demand curve shifts downward overall, and diminishing marginal returns emerge. This phenomenon provides a physical explanation for the conclusion in the subsequent two-stage model that “fast and slow frequency coordination is endogenously formed through optimization.”




Figure 7. FFR requirement curves for nadir security under different PFR levels.

4.2.2. Reserve Allocation and Frequency Tail Risk Under Medium Penetration Rate

Under operating conditions with a penetration rate of 0.60, the configurations of slow frequency and fast frequency reserves exhibit significant differences within a single day. As illustrated in Figure 8, the dashed line indicates the baseline case without frequency-security constraints (or without SoC energy-budget constraints). Slow frequency reserves increase as the load level rises, corresponding to the temporal changes shown in Figure 8a; fast frequency reserves, however, increase significantly during periods when net load growth and the risk of error accumulation are more pronounced, as shown in Figure 8b. In the baseline scenario, which considers only energy balance, the levels of both types of reserves are generally low, lacking explicit protection against low inertia and credible disturbances. After introducing day-ahead frequency constraints, slow frequency reserves are systematically increased to meet quasi-steady-state and frequency nadir requirements. With the further introduction of two-stage closed-loop control, the fast frequency allocation exhibits stronger time-of-day adaptability. Without energy budget constraints, the fast frequency commitment at the day-ahead level may fail to be fully honored during SoC dip scenarios; however, after incorporating the SoC energy budget, the upper bound of the fast frequency commitment becomes more stable, thereby reducing tail risk exposure during critical time periods.




Figure 8. Comparison of PFR and FFR reserve capacities at a 0.6 penetration level. The dashed line indicates the baseline case without frequency-security constraints (or without SoC energy-budget constraints).

Figure 9a to Figure 9c present the 95th percentile statistical results for RoCoF, frequency nadir, and quasi-steady-state deviation, respectively, where the dashed line represents the corresponding security limit (frequency safety threshold). The baseline scheme exceeded the limits for RoCoF and frequency nadir during multiple time periods, indicating that frequency risks are significantly amplified under conditions of low inertia and uncertainty coupling. The current frequency constraint scheme effectively reduces RoCoF and improves the frequency nadir; however, in the absence of a closed-loop fulfillment mechanism, the system's fast frequency fulfillment capability is easily weakened by SoC evolution, resulting in insufficient safety margins during certain time periods. The two-stage closed-loop approach explicitly propagates error impacts to the controllable domain through scenario-based rescheduling, further reducing tail risk; when the SoC energy budget is incorporated, the systematic deviation between fast frequency commitments and deliverability is suppressed, and both the frequency nadir and quasi-steady-state deviation exhibit more stable safety margins.




Figure 9. Hourly 95th-percentile statistics of frequency security metrics. The dashed line represents the corresponding security limit (frequency safety threshold).

Figure 10 illustrates the hierarchical characteristics of the frequency deviation trajectory under critical disturbances, specifically manifesting as differences in peak suppression capability during the initial phase of the disturbance and variations in the recovery phase, where the dashed line denotes the frequency security threshold (minimum acceptable frequency level). In scenarios where the energy budget is not considered, the two-stage scheme is constrained by the SoC drop in certain cases, resulting in reduced fast frequency injection capability during the initial phase of the disturbance and peak deviations approaching the threshold. When the energy budget is incorporated, fast frequency capability

remains at a more stable and achievable level during critical periods, and the margin of peak deviation is further expanded.




Figure 10. Comparison of frequency deviation trajectories under disturbance scenarios. The dashed line denotes the frequency security threshold (minimum acceptable frequency level).

4.2.3. Comprehensive Comparison Across Penetration Rates

The costs, curtailed generation, load losses, and probability of exceeding the frequency nadir threshold across different penetration rates are shown in Figure 11a through Figure 11d, respectively. As penetration increases, curtailment and load loss in the baseline scheme rise in tandem, while the probability of frequency nadir violation increases from 0.18 at low penetration to 0.62 at high penetration (Tables 5 and 6), reflecting the coupled amplification effect of frequency risk and supply reliability risk under conditions of low inertia and high uncertainty. The day-ahead frequency constraint scheme can significantly reduce the probability of frequency deviation, but it remains sensitive to SoC evolution under high penetration rates. The two-stage closed-loop control improves error absorption and reduces tail risk; however, without energy budget constraints, the out-of-bounds probability remains 0.22 and the load loss is 0.008 MWh at high penetration rates. After incorporating the SoC energy budget, the probability of exceeding limits drops to 0 for all three penetration levels, and load losses are eliminated. Meanwhile, the 95th percentile of the frequency nadir decreases from 0.95 Hz to 0.50 Hz under high penetration, representing a reduction of approximately 47.4%. Compared to the day-ahead frequency-constrained scheme, the 95th percentile of the frequency nadir further decreased from 0.60 Hz to 0.50 Hz, a reduction of approximately 16.7%, with the corresponding cost increase falling within the range of 1% to 2% (e.g., increasing from 3,850 to 3,920 under high penetration). The above results indicate that energy budget constraints primarily reduce tail risk by limiting unsustainable fast frequency commitments and stabilizing their deliverability, while the cost is mainly reflected in the opportunity cost associated with reserve and energy retention.

The comprehensive results are summarized in Table 5 according to the dimensions of economic viability and supply reliability, while indicators related to frequency risk and feasibility are summarized in Table 6. These two tables address different evaluation aspects, making it easier to verify the relationships among cost implications, grid integration capacity, supply risk, and frequency risk.




Figure 11. Comprehensive comparison under different penetration levels.

Table 5. Comparison of economic and energy indicators for different methods under various penetration rates.

Penetration	Method	Cost	Curtailment_MWh	LoadShed_MWh
Low	Base UC	5,200	0.2	0.005
Low	FSC UC	5,450	0.25	0
Low	FSC SUC noSoC	5,420	0.22	0.001
Low	FSC SUC	5,520	0.28	0
Medium	Base UC	4,800	0.8	0.05
Medium	FSC UC	4,680	0.9	0
Medium	FSC SUC noSoC	4,660	0.85	0.002
Medium	FSC SUC	4,745	0.92	0
High	Base UC	4,100	1.6	0.08
High	FSC UC	3,850	1.8	0.001
High	FSC SUC noSoC	3,830	1.7	0.008
High	FSC SUC	3,920	1.85	0

To further quantify the economic impact of frequency security constraints, the total operating cost is decomposed into base energy-production cost, reserve-holding/opportunity cost, renewable-curtailment-related cost, and load-shedding penalty. Compared with Base UC, the frequency-security-constrained schemes increase the reserve-holding component because additional synchronous PFR and BESS FFR margins must be retained. After the SoC energy budget is introduced, part of the BESS capacity is deliberately preserved for deliverable fast response, which slightly increases the opportunity cost of storage operation. However, this increase is compensated by lower load-shedding penalties and lower frequency-risk exposure, especially under high renewable penetration. Therefore, the proposed FSC-SUC does not reduce cost by relaxing security requirements; instead, it reallocates a small portion of the operating cost to auditable frequency-security margins and converts this cost into measurable reductions in tail risk and supply-interruption risk.

Table 6. Comparison of frequency security and tail risk indicators for different methods under various penetration rates.

Penetration	Method	MaxDev95_Hz	ViolationProb	SoC_p5
Low	Base UC	0.65	0.18	NaN
Low	FSC UC	0.52	0.04	NaN
Low	FSC SUC noSoC	0.54	0.06	NaN
Low	FSC SUC	0.47	0	NaN
Medium	Base UC	0.8	0.45	NaN
Medium	FSC UC	0.55	0.08	NaN
Medium	FSC SUC noSoC	0.57	0.12	NaN
Medium	FSC SUC	0.49	0	NaN
High	Base UC	0.95	0.62	NaN
High	FSC UC	0.6	0.15	NaN
High	FSC SUC noSoC	0.63	0.22	0.21
High	FSC SUC	0.5	0	0.36

It is important to clarify that the zero nadir violation probability reported in Table 4 reflects the model's performance within the finite scenario set S constructed from historical data and the specified credible disturbance set C . This does not imply absolute elimination of all frequency risks in practice. The following factors contribute to this outcome:

- (1) The safety margin ε calibrated from the validation set introduces intentional conservatism, pushing the operating points away from the constraint boundaries;
- (2) The SoC energy budget constraint explicitly prevents over-commitment of FFR capacity that cannot be sustained over the required duration;
- (3) The scenario set, while representative, is finite and does not cover all possible extreme events outside the historical observations.

To further validate robustness against unseen uncertainties, we conducted an out-of-sample test using 1000 holdout scenarios not included in the training or scenario set. The out-of-sample violation probability for the proposed method is 0.3%, compared to 4.2% for the FSC-SUC without SoC budgeting and 12.5% for the Base UC. These results confirm that the proposed model is not overfitting to the specific scenario set but provides a reliable frequency security guarantee in practice.

4.2.4. Scalability Verification on the IEEE 118-Bus System

To further demonstrate the scalability and applicability of the proposed FSC-SUC framework beyond the IEEE 33-bus system, an additional case study is conducted on the IEEE 118-bus test system. The same four methods, namely Base UC, FSC UC, FSC SUC noSoC, and FSC SUC, are compared under low, medium, and high renewable penetration levels. The evaluation metrics include total operating cost, renewable curtailment, load shedding, and the probability of frequency nadir violation.

The comparison results show that the overall trends observed in the IEEE 33-bus system remain valid for the larger IEEE 118-bus system. As renewable penetration increases, the Base UC scheme exhibits a clear rise in both reliability risk and frequency-security risk. In particular, load shedding increases from 0.06 MWh at low penetration to 1.18 MWh at high penetration, while the nadir violation probability increases from 0.21 to 0.68, indicating that the lack of explicit frequency security constraints becomes increasingly problematic under low-inertia and high-uncertainty conditions.

After introducing day-ahead frequency security constraints, the FSC UC scheme significantly reduces the frequency nadir violation probability to 0.05, 0.10, and 0.18 under the three penetration levels, respectively. However, its performance remains constrained by the absence of a closed-loop realization mechanism and the lack of explicit SoC sustain-

ability control. The FSC SUC noSoC scheme improves coordination between day-ahead commitment and real-time redispatch, but under medium and high penetration levels, it still exhibits residual violation probabilities of 0.14 and 0.24, respectively, together with nonzero load shedding.

By contrast, the proposed FSC SUC scheme eliminates load shedding and reduces the nadir violation probability to zero across all three penetration levels. This result indicates that the combination of two-stage stochastic coordination, commitment consistency, and SoC-based fast-frequency budgeting remains effective in the larger system. The improvement is especially pronounced under high renewable penetration, where the nadir violation probability decreases from 0.68 in the Base UC case to 0 in FSC SUC.

From an economic perspective, the proposed FSC SUC scheme incurs a moderate increase in operating cost compared with FSC UC and FSC SUC noSoC, reflecting the opportunity cost of maintaining more deliverable frequency-support capability and a more conservative SoC trajectory. However, this additional cost is accompanied by improved frequency-security performance and the elimination of load-shedding risk, which demonstrates a favorable trade-off between security and economy in the large-scale test system as well.

The comprehensive results are summarized in Table 7 according to the dimensions of economic viability and supply reliability, while indicators related to frequency risk and feasibility are summarized in Table 8.

Table 7. Comparison of economic and energy indicators for different methods under various penetration rates in the IEEE 118-bus system.

Penetration	Method	Cost	Curtailement_MWh	LoadShed_MWh
Low	Base UC	68,400	3.4	0.06
Low	FSC UC	70,600	3.8	0.01
Low	FSC SUC noSoC	70,000	3.6	0.02
Low	FSC SUC	71,400	4.1	0
Medium	Base UC	61,100	8.8	0.52
Medium	FSC UC	60,400	9.4	0.02
Medium	FSC SUC noSoC	60,000	9.1	0.04
Medium	FSC SUC	61,500	9.8	0
High	Base UC	54,800	17.6	1.18
High	FSC UC	52,900	18.9	0.06
High	FSC SUC noSoC	52,000	18.2	0.11
High	FSC SUC	53,700	19.5	0

Table 8. Comparison of frequency security and tail risk indicators for different methods under various penetration rates in the IEEE 118-bus system.

Penetration	Method	MaxDev95_Hz	ViolationProb	SoC_p5
Low	Base UC	0.72	0.21	NaN
Low	FSC UC	0.57	0.05	NaN
Low	FSC SUC noSoC	0.60	0.07	NaN
Low	FSC SUC	0.52	0	NaN
Medium	Base UC	0.89	0.49	NaN
Medium	FSC UC	0.63	0.10	NaN
Medium	FSC SUC noSoC	0.66	0.14	NaN
Medium	FSC SUC	0.56	0	NaN
High	Base UC	1.08	0.68	NaN
High	FSC UC	0.70	0.18	NaN
High	FSC SUC noSoC	0.75	0.24	0.19
High	FSC SUC	0.60	0	0.34

Overall, the IEEE 118-bus results confirm that the proposed FSC-SUC framework is not limited to the IEEE 33-bus distribution-level case. Instead, it preserves the same qualitative advantages in a larger and more complex system, thereby supporting the practical scalability of the proposed modeling and solution approach.

4.3. Computational Performance of Benders Decomposition

The computational performance of the Benders decomposition is evaluated on a workstation with an Intel Xeon 3.6 GHz CPU and 64 GB RAM, using CPLEX 12.10 as the solver. The convergence behavior for the 30-scenario case is illustrated in Figure 12 (convergence curve where the upper and lower bounds converge within 12 iterations).




Figure 12. Comprehensive comparison under different penetration levels in the IEEE 118-bus system.

Table 9 summarizes the computational metrics. For the 30-scenario case, the algorithm converges in 12 iterations, with a total solution time of 1847 s. The master problem (MILP) accounts for approximately 65% of the total time, while the scenario subproblems (LPs) are solved in parallel across 12 cores, consuming the remaining 35%.

Table 9. Computational performance of the Benders decomposition.

Number of Scenarios	Iterations	Master Time (s)	Subproblem Time (s)	Total Time (s)	Monolithic Status
10	8	520	180	700	Solved (2,100 s)
30	12	1,200	647	1,847	Memory Overflow
50	15	2,100	1,050	3,150	Memory Overflow
100	18	4,500	2,800	7,300	Not attempted

For comparison, solving the monolithic formulation (without decomposition) for the same instance resulted in memory exhaustion due to the explosion of scenario variables. The decomposition reduces peak memory usage from approximately 32 GB to 6 GB, demonstrating its practical necessity for the proposed two-stage model.

4.4. Summary

This chapter validates the effectiveness of a two-stage stochastic frequency-security-constrained unit commitment framework on the IEEE 33-bus system. The results demonstrate that linear surrogates can provide frequency nadir risk boundaries consistent with dynamic evaluations within the threshold neighborhood and form conservative upper bounds through margin correction, thereby supporting the linear solvable embedding of frequency nadir constraints in day-ahead mixed-integer models. The two-stage closed-loop explicitly feeds back the impact of prediction errors on the adjustable domain and tail risk to the day-ahead commitment configuration, while the commitment consistency constraint ensures the enforceability of fast and slow frequency support capabilities during the real-time stage. Upon further introduction of the SoC energy budget constraint, fast frequency commitments are restricted to a sustainable fulfillment range. Under high penetration rates, the probability of frequency nadir violation can be reduced from 0.62 to 0, achieving approximately a 47.4% reduction in the 95th percentile of the frequency nadir, while keeping the cost increment within an auditable range of 1% to 2%. These results collectively validate the engineering effectiveness and physical interpretability of the three key mechanisms discussed in this paper: linear surrogate embedding, two-stage closed-loop control, and energy budget constraints.

4.5. Sensitivity Analysis

To evaluate the robustness of the proposed FSC-SUC, we conduct systematic sensitivity analyses on four key parameters: the frequency nadir threshold, the number of scenarios, the safety margin percentile, and the BESS capacity. The results are summarized in Table 10.

Table 10. Sensitivity of total cost and nadir violation probability to parameter variations.

Parameter	Variation Range	Total Cost (High Penetration)	Nadir Violation Prob.
Frequency Nadir Threshold	−0.1 Hz	3,890	0.0%
	Baseline	3,920	0.0%
	+0.1 Hz	3,850	1.2%
Scenario Count	10 scenarios	3,880	0.5%
	30 scenarios (baseline)	3,920	0.0%
	100 scenarios	3,950	0.0%
Safety Margin (Percentile)	90th	3,870	1.8%
	95th (baseline)	3,920	0.0%
	99th	3,980	0.0%
BESS Capacity	−40%	3,850	2.1%
	−20%	3,890	0.4%
	Baseline	3,920	0.0%

The results in Table 10 indicate that the proposed model maintains zero violation probability across a wide range of parameters except when thresholds are relaxed or capacities are severely undersized. The cost variation remains within 2–3% across all tested configurations, confirming the auditable conservatism and robustness of the proposed framework.

5. Conclusions

This paper addresses the issues of increased frequency security risks and insufficient availability of frequency regulation resources under conditions of high penetration of renewable energy by proposing a two-stage stochastic frequency-security-constrained unit

commitment model, FSC-SUC. The model performs joint optimization at the day-ahead hourly level to determine synchronous generator start-up/shutdown and reference output, slow frequency regulation reserves, energy storage SoC trajectories, and fast frequency regulation reserves. At the real-time 5 min scale, it employs scenario-based rescheduling to accommodate errors and uses commitment consistency constraints to ensure that slow and fast frequency regulation capabilities are treated as deliverable resources. To address the bottleneck of difficulty in directly embedding frequency nadir constraints into mixed-integer optimization, this paper proposes a unified modeling framework that combines analytical criteria with linear surrogate-based approximation. It calibrates the safety margin using high-percentile residuals from the validation set to achieve auditable conservative control. In the IEEE 33-bus case study, the proposed method effectively reduces the probability of frequency nadir violations and the risk of load loss under various renewable energy penetration rates while enhancing the sustainable fulfillment capability of the SoC trajectory. Compared to approaches that rely solely on day-ahead frequency constraints or lack energy budget constraints, the model significantly improves the stability of the frequency safety margin with limited cost increases.

Future work can be expanded in the following directions. First, higher-fidelity dynamic frequency models or data-driven verification steps can be incorporated into the iterative correction mechanism for surrogate constraints to further improve the tightness of the surrogate boundaries and reduce unnecessary conservatism. Second, new frequency support resources such as multi-regional interconnected systems, virtual synchronous generators, and flexible DC systems can be integrated into a unified framework to study cross-regional frequency security constraints and resource coordination. Third, ancillary service market mechanisms can be combined with risk measurement methods to construct a frequency security optimization model that integrates market clearing and operational dispatch, thereby enhancing the practical applicability of the proposed method.

Author Contributions: Conceptualization, G.W.; Methodology, X.J.; Validation, R.S.; Investigation, Z.L.; Visualization, X.Z.; Writing—original draft, Z.Z.; Writing—review & editing, J.W.; Supervision, G.W.; Funding acquisition, X.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Science and Technology Project of State Grid Henan Electric Power Company (Grant No. 52170226000G).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: Authors Guodong Wang, Ran Sun and Jianbo Wang were employed by the company State Grid Henan Electric Power Company; Author Xiaoke Zhang was employed by the company State Grid Henan Electric Power Company Electric Power Research Institute. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Appendix A

The full set of surrogate coefficients (α_h, β_h) for Equation (30) is provided in the Table A1. Table A1 summarizes the hyperplane structure for the case with a disturbance magnitude of 0.6 pu. The feature space is partitioned by H_{sys} into four regions, resulting in eight hyperplanes.

Table A1. Summary of linear surrogate hyperplane coefficients.

Region	H_{sys} Range	Number of Hyperplanes	α_h Range	β_h
1	2.0–3.5 s	2	0.15–0.28	0.42–0.55
2	3.5–5.0 s	2	0.12–0.22	0.38–0.48
3	5.0–6.5 s	2	0.08–0.18	0.32–0.42
4	6.5–8.0 s	2	0.05–0.14	0.28–0.36

References

- Kim, S.-E.; Chun, Y.-H. Modeling of Linearized Generator Inertia Constraints for Unit Commitment. *Energies* **2024**, *17*, 1120. [\[CrossRef\]](#)
- Liu, L.; Hu, Z.; Wen, Y.; Ma, Y. Modeling of Frequency Security Constraints and Quantification of Frequency Control Reserve Capacities for Unit Commitment. *IEEE Trans. Power Syst.* **2024**, *39*, 5781–5791. [\[CrossRef\]](#)
- Tuo, M.; Li, X. Security-Constrained Unit Commitment Considering Locational Frequency Stability in Low-Inertia Power Grids. *IEEE Trans. Power Syst.* **2023**, *38*, 4134–4145.
- Shikuma, R.; Fujimoto, Y.; Orihara, D.; Kikusato, H.; Taoka, H.; Hayashi, Y. Impact of Virtual Synchronous Generators on Frequency-Constrained Unit Commitment: A Neural Network-Assisted Framework for Assessing Rate of Change of Frequency. *IEEE Access* **2025**, *13*, 117716–117731. [\[CrossRef\]](#)
- She, B.; Li, F.; Ai, X.; Wen, J. Virtual Inertia Scheduling (VIS) for Real-Time Economic Dispatch of IBR-Penetrated Power Systems. *IEEE Trans. Sustain. Energy* **2024**, *15*, 938–951.
- Chen, X.; Liu, Y.; Wu, L. Towards Improving Unit Commitment Economics: An Add-On Tailor for Renewable Energy and Reserve Predictions. *IEEE Trans. Sustain. Energy* **2024**, *15*, 2625–2638. [\[CrossRef\]](#)
- Li, P.; Su, H.; Zhao, X.; Qi, J.; Wei, H.; Bai, X. Formulating a quadratic frequency nadir constraint in unit commitment based on energy conservation law. *Electr. Power Syst. Res.* **2024**, *228*, 110086. [\[CrossRef\]](#)
- Yu, J.; Yong, P.; Yu, J.; Yang, Z. Frequency security constraint in unit commitment with detailed frequency response behavior of wind turbines. *Energy* **2024**, *313*, 133735.
- Jiang, B.; Guo, C.; Chen, Z. Frequency constrained unit commitment considering reserve provision of wind power. *Appl. Energy* **2024**, *361*, 122898. [\[CrossRef\]](#)
- Huo, J.; Zhang, Y.; Liu, Y.; Lin, W.; Sun, Y. Joint Dispatch Model for Power Grid and Wind Farms Considering Frequency Modulation Delay. *Energies* **2025**, *18*, 6263. [\[CrossRef\]](#)
- Qian, M.; Wang, J.; Yang, D.; Yin, H.; Zhang, J. An Optimization Strategy for Unit Commitment in High Wind Power Penetration Power Systems Considering Demand Response and Frequency Stability Constraints. *Energies* **2024**, *17*, 5725. [\[CrossRef\]](#)
- Jiang, S.; Hu, Q.; Li, F.; Bai, L.; Dong, Y.; Zhao, X. Coordinative frequency-constrained unit commitment model for HVDC interconnected AC systems. *Int. J. Electr. Power Energy Syst.* **2022**, *141*, 108176. [\[CrossRef\]](#)
- Xu, D.; Wu, Z.; Zhu, L.; Guan, L. A novel frequency constrained unit commitment considering VSC-HVDC's frequency support in asynchronous interconnected system under renewable energy Source's uncertainty. *Electr. Power Syst. Res.* **2025**, *238*, 111098. [\[CrossRef\]](#)
- Baig, A.M.; Badesa, L.; Wang, Y.; Strbac, G. Co-optimising frequency-containment services from zero-carbon sources in electricity grids dominated by Renewable Energy Sources. *Int. J. Electr. Power Energy Syst.* **2025**, *164*, 110375. [\[CrossRef\]](#)
- Maureira, G.C.; Benigni, A.; Monti, A. Optimal allocation of fast frequency response in power systems considering inertia distribution. *Sustain. Energy Grids Netw.* **2025**, *44*, 101905. [\[CrossRef\]](#)
- Esfahani, M.; Alizadeh, A.; Amjadi, N.; Kamwa, I. A distributed VPP-integrated co-optimization framework for energy scheduling, frequency regulation, and voltage support using data-driven distributionally robust optimization with Wasserstein metric. *Appl. Energy* **2024**, *361*, 122883. [\[CrossRef\]](#)
- Zhou, Z.; Wang, Z.; Zhang, Y.; Wang, X. Nash Bargaining-Based Coordinated Frequency-Constrained Dispatch for Distribution Networks and Microgrids. *Energies* **2024**, *17*, 5661. [\[CrossRef\]](#)
- Bagherinezhad, A.; Hosseini, M.M.; Parvania, M. Real-Time Hierarchical Energy Flexibility Management of Integrated Hybrid Resources. *IEEE Trans. Smart Grid* **2024**, *15*, 2348–2360.
- Wang, X.; Zhao, T.; Parisio, A. Frequency regulation and congestion management by Virtual Storage Plants. *Sustain. Energy Grids Netw.* **2022**, *29*, 100586. [\[CrossRef\]](#)
- Shin, W.; Lee, C.; Chung, I.-Y.; Lim, J.; Youn, J.; Rhie, Y.; Hur, K.; Shim, J.W. Reserve replacement from governor to energy storage system on conventional generator for operating-cost reduction. *Appl. Energy* **2022**, *324*, 119697. [\[CrossRef\]](#)
- Koltermann, L.; Cortés, M.C.; Figgenger, J.; Zurmühlen, S.; Sauer, D.U. Power curves of megawatt-scale battery storage technologies for frequency regulation and energy trading. *Appl. Energy* **2023**, *347*, 121428. [\[CrossRef\]](#)

22. Tang, R.; Xu, Q.; Fang, J.; Jiang, A. A dynamic bidding strategy of hybrid energy storage system participating in day-ahead frequency regulation market. *J. Energy Storage* **2022**, *56*, 106161. [\[CrossRef\]](#)
23. Wang, Z.; Hou, S. Optimal participation of battery swapping stations in frequency regulation market considering uncertainty. *Energy* **2024**, *302*, 131815. [\[CrossRef\]](#)
24. Li, S.; Xu, Q.; Huang, J. Research on the integrated application of battery energy storage systems in grid peak and frequency regulation. *J. Energy Storage* **2023**, *59*, 106459. [\[CrossRef\]](#)
25. Xie, R.; Wang, Y.; Zhang, S.; Lin, B.; Chen, Q.; Wang, F.; Wang, X.; Chen, Y.; Xia, B. BESS Frequency Regulation Strategy on the Constraints of Planned Energy Arbitrage Using Chance-Constrained Programming. *Energy Rep.* **2022**, *8*, 73–80. [\[CrossRef\]](#)
26. Abiari, H.; Daemi, T.; Saeid, S.A. Frequency-constrained unit commitment under uncertain PFR of energy storage systems. *J. Energy Storage* **2023**, *73*, 108848. [\[CrossRef\]](#)
27. Huang, W.; Zheng, Q.; Hu, Y.; Huang, Y.; Zhou, S. Optimization of Frequency Modulation Energy Storage Configuration in Power Grid Based on Equivalent Full Cycle Model. *Energies* **2024**, *17*, 2120. [\[CrossRef\]](#)
28. Lee, J.; Han, S.; Lee, D. Optimizing the Location of Frequency Regulation Energy Storage Systems for Improved Frequency Stability. *Batteries* **2023**, *9*, 592. [\[CrossRef\]](#)
29. Tian, C.; Zhang, Q.; Mei, D.; Chen, E.; Li, Z.; Zhang, X. Day-Ahead Economic Dispatch Optimization for Industrial Consumers Utilizing Shared Energy Storage Stations. *Processes* **2025**, *13*, 3964. [\[CrossRef\]](#)
30. Ranjbar, H.; Ahmadian, A.; Elkamel, A. Frequency-constrained autonomous microgrid planning for mining industry applications. *Appl. Energy* **2025**, *396*, 126201.
31. Vieira, M.P.; Martins, A.C.P.; Soler, E.M.; Balbo, A.R.; Nepomuceno, L. Two-stage robust market clearing procedure model for day-ahead energy and reserve auctions of wind–thermal systems. *Renew. Energy* **2023**, *218*, 119276. [\[CrossRef\]](#)
32. Li, K.; Wei, L.; Fang, J.; Ai, X.; Cui, S.; Zhu, M.; Wen, J. Incentive-compatible primary frequency response ancillary service market mechanism for incorporating diverse frequency support resources. *Energy* **2024**, *306*, 132339. [\[CrossRef\]](#)
33. Xu, T.; Sun, J.; Meng, H.; Wang, R.; Ji, Y.; Zhang, Y.; Song, P.; Xiang, J. Frequency regulation market participation of distributed energy storage systems. *Int. J. Electr. Power Energy Syst.* **2025**, *172*, 111099. [\[CrossRef\]](#)
34. Zhan, J.; Huang, M.; Sun, X.; Chen, Z.; Zhang, Y.; Xie, X.; Chen, Y.; Qiao, Y.; Ai, Q. Aggregation Method and Bidding Strategy for Virtual Power Plants in Energy and Frequency Regulation Markets Using Zonotopes. *Energies* **2025**, *18*, 2458. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.