Stage-wise hybrid nested Benders' decomposition-stochastic dual dynamic programming for virtual power plants

Youngdoo Choi[a], Seokwoo Kim[a], Sanghyeon Bae[b,*], Dong Gu Choi[a,*]

[a]*Department of Industrial and Management Engineering, Pohang University of Science and Technology (POSTECH), Pohang, Gyeongbuk 37673, Republic of Korea*
[b]*Department of Business Administration, Ajou University, Suwon, Gyeonggi 16499, Republic of Korea*

Abstract

Participants in energy markets make sequential decisions across multiple time horizons under uncertainty, leading to large-scale multistage stochastic optimization problems. Stochastic dual dynamic programming is widely used for its tractability, but its application to modern energy markets is challenged by nested dependencies induced by participation across multiple interrelated markets under increasing uncertainty from distributed energy resources. To address this, we introduce a stage-wise hybrid nested Benders' decomposition-stochastic dual dynamic programming approach by nodalizing early-stage uncertainties. We apply the approach to a bidding and operation problem for a virtual power plant in Day-ahead and Intraday electricity markets, where strong nested dependencies exist between Day-ahead prices and Intraday uncertainties. We reformulate the problem by incorporating bidding decisions into the state space and by introducing a penalty term in the final stage to address recourse infeasibility. Numerical experiments on a high-renewable system demonstrate that the proposed method converges and outperforms conventional approaches. Furthermore, the performance gains over stochastic dual dynamic programming increase as the level of the nested dependency increases.

Keywords: OR in Energy, Auctions/bidding, Nested Benders' decomposition, Stochastic dual dynamic programming

1. Introduction

Electricity markets are typically designed as complex multistage systems, in which various participants, including transmission operators, system operators, and generators, make decisions

*Corresponding Authors
Email addresses: chlduden11@postech.ac.kr (Youngdoo Choi), seokwookim@postech.ac.kr (Seokwoo Kim), azureharry@ajou.ac.kr (Sanghyeon Bae), dgchoi@postech.ac.kr (Dong Gu Choi)

across multiple time horizons ranging from long-term planning to day-ahead and real-time operations. The structure allows participants to progressively adjust their decisions, starting from early commitments and refining them as dispatch time approaches and uncertainty is gradually resolved.

The decision problems in such a multistage decision-making process naturally require large-scale stochastic optimization models involving multiple sources of uncertainty. Solving these problems in a monolithic manner is computationally intractable, motivating the use of decomposition-based approaches. Stochastic Dual Dynamic Programming (SDDP) has been one of the most widely adopted methods, as it provides tractable approximations of value functions and reduces computational complexity.

However, the direct application of SDDP to modern energy market problems has become increasingly challenging. The growing penetration of distributed energy resources (DERs), including renewable energy sources, energy storage systems (ESS), and demand response (DR), has increased volatility and uncertainty in prices, generation, and curtailment. As a result, market participants must increasingly consider participation across multiple interrelated markets, including Day-ahead(DA), Intraday(ID), capacity, and ancillary service markets.

Jointly optimizing decisions across these markets leads to nested multi-settlement problems. In addition, electricity prices, demand, and renewable generation exhibit nested dependencies across multiple time scales, including daily, monthly, and yearly patterns. These features violate the stage-wise independence assumption underlying classical SDDP. While several extensions of SDDP have been proposed to capture stage-wise dependencies (see, e.g., Chapter 14 of Füllner and Rebennack (2025)), they are typically limited to Markovian settings, where dependencies are usually represented through adjacent-stage transitions. In contrast, stochastic processes across interconnected energy markets are difficult to capture under such a Markovian setting, as they may exhibit nested dependencies between random variables separated by two or more stages. Consequently, both the decision structure and the stochastic processes exhibit nested and interdependent characteristics, significantly complicating the application of standard SDDP.

To overcome the challenge, we propose a variant of SDDP, referred to as *stage-wise* hybrid nested Benders' decomposition/Stochastic dual dynamic programming (hybrid NBD/SDDP), that addresses stochastic processes with nested dependency structures. While the proposed method is generally applicable to a broad class of multistage decision problems in energy markets, we focus on a bidding and operation problem of a virtual power plant (VPP) in DA and ID electricity markets on Jeju Island, South Korea. In the problem, the dependency between DA prices and ID

stochastic parameters, hereafter referred to as the DA-ID dependency, exhibits a pronounced nested dependency structure, further intensified by the high penetration of renewable energy sources.

The main contributions of this work are as follows. First, we present a computationally tractable multistage stochastic optimization model for joint bidding and energy storage system (ESS) operations across DA and ID electricity markets. To capture the intertemporal dependence of bidding decisions and their temporal coupling with operations, we incorporate DA and ID bidding decisions into the state space through appropriate state-variable reformulations. Moreover, we employ a tailored scenario tree structure that captures DA–ID dependency by nodalizing only the DA stage and constructing recombining trees for the ID stages conditional on each DA price node.

Second, to capture the nested dependencies, we propose an extended version of the hybrid NBD/SDDP algorithm tailored to our problem setting. Under a conditional stage-wise independence assumption, the number of expected cost-to-go (ECTG) functions scales with the product of the number of DA-stage nodes and the number of ID stages. Moreover, to address computational challenges, we introduce an approximation scheme in which the number of ID-stage recombining scenario trees is reduced, thereby enabling a tractable approximation with lower computational effort.

Lastly, we conduct numerical experiments using real-world data from the Jeju Island electricity market in South Korea where the joint occurrence of negative prices and curtailment reinforces the DA–ID dependency. We first demonstrate the convergence of the hybrid NBD/SDDP algorithm and its superiority in terms of mean profit compared to alternative solution methods. Second, we investigate the trade-off between computational effort and solution quality in our approximation scheme. Finally, by varying the degree of DA–ID dependency in the scenario tree, we show that hybrid NBD/SDDP and SDDP perform similarly when the dependency is weak, while the advantage of hybrid NBD/SDDP becomes more pronounced as the dependency level increases.

The remainder of this paper is organized as follows. Section 2 reviews the related works and emphasizes the distinctions between our study and the existing studies. Section 3 describes our problem settings, and Section 4 formulates the multistage stochastic model and presents its dynamic programming reformulation. Section 5 proposes the hybrid NBD/SDDP approach and its approximation scheme. Section 6 reports the experimental setup and computational results. Section 7 concludes with a discussion of limitations, practical implications, and future research directions.

2. Related Works

Section 2.1 reviews the relevant literature on SDDP and its applications to energy systems, including extensions that relax the stage-wise independence assumption. Section 2.2 surveys previous studies on bidding and operation problems in DA and ID electricity markets. At the end of each subsection, we highlight how our work differs from the existing literature.

2.1. Related works on stochastic dual dynamic programming

Sequential decision-making problems, such as multistage bidding and operation problems, can be formulated using either a Stochastic Dynamic Programming (SDP) model or a Multistage Stochastic Programming (MSP) framework. However, both approaches suffer from the curse of dimensionality, making them difficult to apply to large-scale scenario tree problems. To address this challenge, Pereira and Pinto (1985, 1991) proposed Stochastic Dual Dynamic Programming (SDDP) and successfully applied it to hydro and thermal planning problems. The method exploits the stage-wise independence assumption to significantly reduce the number of expected cost-to-go (ECTG) functions that need to be approximated, and has since been widely applied in the energy sector, including hydro power scheduling (Abgottspon et al., 2014; Hjelmeland et al., 2018), economic dispatch and unit commitment problems (Cordera et al., 2023; Zou et al., 2018), as well as the multistage DA and ID market bidding problems (Wozabal and Rameseder, 2020; Shinde et al., 2022).

A major drawback of SDDP is that it relies on the stage-wise independence assumption, which often does not hold in real-world applications. To address stage-wise dependency, either a time series-based approach (TS-SDDP) or a Markov chain-based approach (MC-SDDP) can be used; see Löhndorf and Shapiro (2019) and Chapter 14 of Füllner and Rebennack (2025) for details. However, MC-SDDP can only be applied when the underlying stochastic processes can be represented through Markovian stage-to-stage transitions. Although TS-SDDP extends applicability to nested settings by augmenting the state space, it requires convexity when stochastic parameters are treated as decision variables and assumes that the underlying process follows an AR(n) model, which is not guaranteed in practice.

To address these limitations, Rebennack (2016) proposed a hybrid NBD/SDDP method that partitions the stochastic parameters into two independent components: one that is stage-wise independent (or has a simple dependency structure) and another that captures stage-wise dependent uncertainty. A scenario tree is constructed only for the latter component to account for the depen-

dency structure. Building on this idea, we extend the methodology from the random-variable level to the stage level to capture nested dependency structures. If the stochastic processes in later stages are conditionally independent across time given the realization of an earlier-stage variable, then it suffices to construct a scenario tree only for the earlier stage. In this case, different realizations of the earlier-stage variable induce different conditional distributions of the subsequent processes, thereby resulting in different ECTG functions.

In this study, we introduce an extended version of hybrid NBD/SDDP and apply it to a multistage stochastic bidding and operation problem in DA and ID electricity markets. In the problem, the dependence between DA and ID prices for a given delivery period induces a nested dependencies, as ID prices for earlier delivery periods are sequentially realized between the DA commitment and the realization of the corresponding ID price. To this end, we develop a *stage-wise* hybrid NBD/SDDP framework in which only the DA stage decisions are nodalized, while the remaining uncertainties, such as ID prices, generation, and curtailment, are treated as conditionally stage-wise independent given the DA prices.

2.2. Related works on multistage VPP bidding and operation Problem

Many studies have investigated the optimal bidding and operation problem for VPPs in the DA market using various methodological frameworks. Early studies employed deterministic optimization models (Mashhour and Moghaddas-Tafreshi, 2011; Pandžić et al., 2013a), two-stage frameworks (Fazlalipour et al., 2019; Pandžić et al., 2013b), or three-stage formulations (Morales et al., 2010; Heredia et al., 2018). However, these approaches fall short of capturing the full complexity of the DA and multiple ID markets, which inherently require a more detailed multistage modeling framework.

To address these limitations, recent studies have adopted more detailed multistage modeling frameworks. Kim et al. (2022) applied stochastic dynamic programming (SDP) to solve the multistage bidding and ESS operation problem in the South Korean electricity market. Pan and Guan (2022) proposed an integrated stochastic self-scheduling strategy for two-settlement electricity markets, formulating a bidding and operation problem within a multistage stochastic programming framework. However, both studies relied on an approximated scenario tree due to the curse of dimensionality, which may compromise solution quality. Kim and Choi (2024) addressed the multistage stochastic bidding problem for a VPP by using a sample robust approach. Nonetheless, this work employed a linear decision rule-based optimization model, again due to dimensionality

challenges. Finnah et al. (2022) employed approximate dynamic programming (ADP) to address multistage bidding and ESS operation problems. However, these studies relied on policy function approximation (PFA) and value function approximation (VFA), respectively, instead of deriving analytical optimization solutions.

To overcome such computational complexity, some studies have employed SDDP in multistage stochastic DA and ID bidding and operation problems. Wozabal and Rameseder (2020) addressed the optimal bidding of a VPP in the Spanish DA and auction-based ID electricity markets using a variant of SDDP, known as Approximate Dual Dynamic Programming(ADDP). Shinde et al. (2022) also used an SDDP model for VPP trading in continuous ID electricity markets, incorporating thermal, wind, and hydro units. However, none of these studies jointly consider DA and ID bidding decisions together with ESS operations within an integrated multistage framework under nested dependencies between DA and ID prices.

In this paper, we apply our proposed algorithm to a multistage DA and ID bidding problem for a VPP with ESS operations. The price data are obtained from the electricity market in Jeju Island, South Korea, where a high penetration of renewable energy leads to frequent negative prices, substantial price volatility, and recurrent curtailment events. These characteristics induce strong dependencies between DA and ID prices, and curtailment. Therefore, this setting provides a compelling case study for evaluating the performance of stage-wise hybrid NBD/SDDP under strong nested DA–ID dependency.

3. Problem Description

This section first introduces our problem settings in Section 3.1, followed by a detailed description of the modeling assumptions and setup in Section 3.2.

3.1. Bidding and ESS operations of a VPP in DA and ID electricity markets

The inherent output variability of DERs has led to an increase in transaction volumes in short-term electricity markets (Neuhoff et al., 2016), thereby reinforcing trading structures across multiple time frames, including DA and ID markets. ID markets enable participants to adjust their DA positions through a market-based mechanism. Deviations in physical delivery from the final position at the ID market's gate closure are subject to imbalance penalties. This market structure compels participants with DERs to make interdependent decisions under uncertainty across multiple time horizons.

However, the limited capacity and high output variability of DERs often hinder direct participation in DA and ID markets. A recently emerging brokerage model, known as the VPP, facilitates the market participation of small-scale DERs by aggregating their output through a cloud-based IT platform. The VPP operates collectively in the electricity market as a single power plant, representing the combined energy resources of its constituent DERs.

Building on this market structure and participation model, we consider the following problem setting. We study a VPP participating in both the DA and ID electricity spot markets on Jeju Island, South Korea. Given that solar PV and wind resources account for approximately 51% of the installed capacity on Jeju Island (Korea Power Exchange (KPX), 2025), negative electricity prices can occur when renewable generation is overly supplied, potentially leading to the curtailment.

Under the market structure, electricity spot trading operates within a multi-settlement framework consisting of a single DA market covering all 24 delivery hours and 24 ID markets corresponding to each delivery hour. Each market has distinct opening and closing times. The DA market opens at 8:00 and closes at 11:00 on the preceding day, with auction results published at 17:00. Following the announcement of auction results for the DA market, market participants can revise their bidding profiles in the corresponding ID markets for each of the 24 delivery hours. Each ID market opens after the 17:00 publication of DA market results, closes one hour before the associated delivery hour, and publishes the settled price 15 minutes before delivery. For example, the ID market for the 00:00–01:00 delivery closes at 23:00 on the day before delivery, and the ID price is announced at 23:45. Table 1 summarizes the overall operations of both DA and ID markets.

Table 1: Overview of operations for Jeju Island's electricity markets

Markets	DA	ID1	ID2	$\cdots$	ID24
Market opening	08:00(D-1)	17:00(D-1)	17:00(D-1)	...	17:00(D-1)
Market closing	11:00(D-1)	23:00(D-1)	00:00	...	22:00
Publication of results	17:00(D-1)	23:45(D-1)	00:45	...	22:45
Planning horizon	00:00–24:00	00:00-01:00	01:00-02:00	...	23:00-24:00
Number of traded hours	24	1	1	...	1

In each market, the VPP submits an offer in a price–quantity pair. The offer is accepted if the market-clearing price is greater than or equal to the offer price; otherwise, it is rejected. If accepted, the VPP is obligated to deliver the offered quantity and is compensated at the market-clearing price. During the planning horizon of each ID market, the VPP dispatches the output of DERs in accordance with the cleared quantity. Ex post, any deviation between metered output and the cleared quantity is settled through the imbalance mechanism (i.e., imbalance charges or penalties). Therefore, by strategically operating the ESS, the VPP can reduce such deviations and

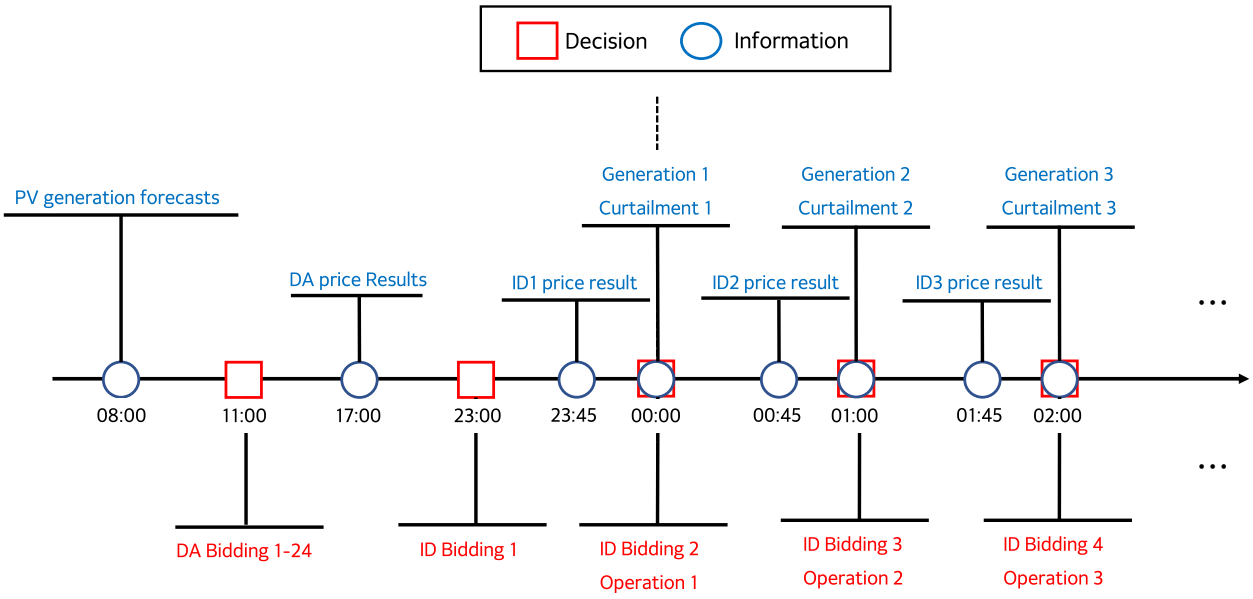

Figure 1: Flow of decisions of a VPP in DA and ID electricity markets.

minimize curtailment.

3.2. Assumptions and setups

We consider a VPP consisting of PV generators and ESSs and jointly optimizing two portfolios. The PV portfolio represents the combined output of all constituent generators, while the ESS portfolio captures the aggregate charging and discharging capabilities of the fleet. Over a one-day horizon, the initial SoC is fixed at 50%, and a daily cycle constraint ensures that it returns to this level at the end of the day. This reflects typical market conditions in which electricity prices are higher during nighttime hours, particularly in systems with high renewable penetration, making it optimal for the ESS to gradually discharge from the evening (18:00, D-1) to the early morning (06:00).

Bidding decisions are taken immediately before each market's closing time, summarized in Table 1. Operational decisions (PV generation dispatch and ESS charging/discharging) are taken at the start of each delivery hour; e.g., the 00:00–01:00 decision is made at 00:00. Note that this difference in decision timing creates a mismatch between ID bidding and operations. For instance, at 04:00, the ID bid for 05:00–06:00 is submitted while the 04:00 operation is executed. We address this mismatch by state variable reformulation in Section 4. Figure 1 illustrates the overall problem settings.

For the bidding decision variables of the VPP, we consider only the bidding quantities, since bidding prices are typically intended to reflect generators' marginal costs, which are close to zero

for renewable resources. The realized generation may be subject to curtailment when renewable generation is oversupplied. The deviation of ID bidding quantities from the corresponding decisions in DA markets is bounded within a prescribed range since the primary role of ID markets is to facilitate the VPP to adjust and rebalance its PV generation dispatch and ESS charging/discharging to help the stability of the grid.

To focus on the DA–ID dependency, we assume that the ID price at each time period depends only on the corresponding DA price. Although ID prices may still exhibit conditional dependencies across time periods given the DA price profile in practice, such dependencies are not modeled in this study. Accordingly, conditional on the DA price profile, ID prices are assumed to be stage-wise independent across time periods.

4. Model Formulation

In this section, we outline the MSLP model for our problem. The model has the following three key features. First, to capture the DA-ID dependency, we employ a tailored scenario tree structure in which the first stage consists of DA price profile nodes, and for each DA-stage node, a stage-wise independent ID recombining scenario tree follows. To the best of our knowledge, this type of scenario tree construction for modeling DA–ID dependency has not been previously explored in studies using SDDP-type algorithms. Second, to make the resulting formulation amenable to SDDP-type algorithms, we further apply state-variable reformulations. Since the DA bidding profile determined in the DA stage remains effective until the final stage, the corresponding information must be delivered across subsequent stages via state variables. In addition, the temporal mismatch between ID bidding decisions and ESS operations necessitates an additional state-variable reformulation to transmit ID bidding decisions to the next stage. Lastly, in order to avoid recourse infeasibility caused by the constraint in the last stage where the SoC value is fixed at 50% of its capacity, we penalize deviations from this target value. This penalty propagates from the last stage across all stages and variables, so that as the number of iterations increases, the SoC value in the last stage converges to the desired value.

4.1. Notations and scenario tree

We begin by introducing the notation for the DA and ID scenario trees. Let $\mathbb{T} = \{1, 2, 3, \ldots, 24\}$ denote the set of time periods and let $T = |\mathbb{T}|$ denote its cardinality. For the DA scenario, define $\mathcal{P}_{DA} = \{P_{DA}^1, \ldots, P_{DA}^{N_{DA}}\}$ as the set of DA price profiles, where $P_{DA}^\ell = (P_{DA,1}^\ell, \ldots, P_{DA,T}^\ell)$

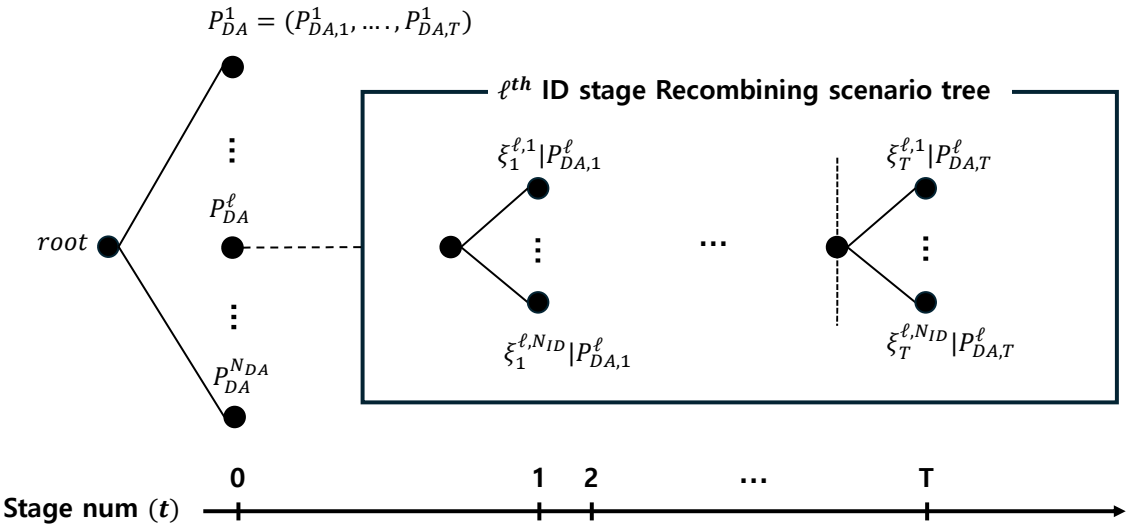

Figure 2: Illustration of the scenario tree structure consisting of DA price profiles and the corresponding ID-stage recombining scenario trees.

represents a DA price profile. Given a DA price scenario P_{DA}^{ℓ}, the associated ID recombining scenario tree is denoted by $\mathcal{T}^{\ell}$. For each tree $\mathcal{T}^{\ell}$, each stage t consists of N_{ID} random vectors, each given by $\xi_t^{\ell,j} = \{P_{ID,t}^{\ell,j}, \delta_{E,t}^{\ell,j}, \delta_{C,t}^{\ell,j}\}$, where $P_{ID,t}^{\ell,j}$ denotes the ID price, while $\delta_{E,t}^{\ell,j}$ and $\delta_{C,t}^{\ell,j}$ denote the generation/DA forecast ratio and curtailment ratio, respectively. The index set is denoted by $[N] := \{1, 2, \ldots, N\}$ (e.g., $[N_{DA}] := \{1, 2, \ldots, N_{DA}\}$).

We assume that both DA price profiles and ID-stage random vectors are equally likely, i.e., $\Pr(P_{DA} = P_{DA}^{\ell}) = \frac{1}{N_{DA}}$ and $\Pr(\xi = \xi_t^{\ell,j}) = \frac{1}{N_{ID}}$. The scenario tree is depicted in Figure 2, and other model parameters and decision variables are summarized in Table 2.

4.2. Multistage stochastic programming model

4.2.1. DA stage

In the DA stage, the producer selects an optimal DA bidding profile for the next delivery day to maximize expected profit over one DA settlement and twenty-four ID settlements.

$$(\mathcal{V}) \quad \max \quad 0 + \mathcal{Q}_{DA}(x_{DA}) := \frac{1}{N_{DA}} \sum_{\ell \in [N_{DA}]} Q_0^{\ell}(x_{DA}; P_{DA}^{\ell}) \tag{1a}$$

$$\text{s.t.} \quad q_{DA,t} \in [0,\, E_{DA,t} + B] \qquad \forall t \in \mathbb{T} \tag{1b}$$

In the above formulation, $\mathcal{Q}_{DA}(x_{DA})$ denotes the expected cost-to-go(ECTG) function, which represents the expected profit after DA bidding. The DA bidding profile is carried to the next stage as a state variable $x_{DA} = q_{DA}$, since they are used in the DA market clearing at 17:00.

Table 2: Summary of the notation.

Notation	Definition
Fixed Parameters	
E_{DA}	Day ahead PV generation forecast profile (kWh)
B	PCS capacity of the ESS (kW)
ν_c, ν_d	Charging and discharging efficiency of the ESS
$\overline{s}, \underline{s}$	Maximum and minimum ESS storage capacity (kWh)
w	Bound on deviation of ID bidding amount from DA bidding amount (kWh)
γ	Over and under penalty coefficient (KRW/kWh)
Stochastic Parameters	
P^{ℓ}_{DA}	DA price profile (KRW/kWh) for $\ell \in [N_{DA}]$
$P^{\ell}_{ID,t}$	ID price (KRW/kWh) at time t for $\ell \in [N_{DA}]$
$\delta^{\ell}_{E,t}$	Generation/DA forecast ratio at time t for $\ell \in [N_{DA}]$
$\delta^{\ell}_{c,t}$	Curtailment ratio at t for $\ell \in [N_{DA}]$
Decision Variables	
q_{DA}	DA bidding amount profile (kWh)
$q^{\ell}_{ID,t}$	ID bidding amount at time t (kWh) for $\ell \in [N_{DA}]$
s^{ℓ}_t	State of Charge(SoC) at time $t+1$ (kWh) for $\ell \in [N_{DA}]$
g^{ℓ}_t	PV production amount at time t (kWh) for $\ell \in [N_{DA}]$
c^{ℓ}_t	ESS charging amount at time t (kWh) for $\ell \in [N_{DA}]$
d^{ℓ}_t	ESS discharging amount at time t (kWh) for $\ell \in [N_{DA}]$
u^{ℓ}_t	Dispatch amount at time t (kWh) for $\ell \in [N_{DA}]$
$q^{\ell}_{D,t}$	Dispatch instruction at time t (kWh) for $\ell \in [N_{DA}]$

Constraint (1b) impose limits on the bidding quantity. Each bidding amount is constrained to lie between 0 and the sum of the generation forecast and the ESS capacity.

$$Q^{\ell}_0(x_{DA}; P^{\ell}_{DA}) := \max \quad \sum_{t\in\mathbb{T}} P^{\ell}_{DA,t}\, q_{DA,t} \;+\; \mathcal{Q}^{\ell}_0\big(x^{\ell}_0\big) \tag{2a}$$

$$\text{s.t.} \quad q^{\ell}_{ID,1} \in [0,\, B] \tag{2b}$$

The above formulation describes the time period from the realization of the DA price to the closure of the ID0 market. Given a DA price P^{ℓ}_{DA} and a DA bidding profile $x_{DA} = q_{DA}$, the DA market is cleared and the generator yields a profit of $\sum_{t\in\mathbb{T}} P^{\ell}_{DA,t}\, q_{DA,t}$. The objective function combines this DA profit with the expected value of the subsequent stage, represented by the ECTG function $\mathcal{Q}^{\ell}_0(\cdot)$, defined in (3). During this stage, the ID1 bidding decision must be determined according to the bidding rules specified in (2b). The bidding amounts for the DA markets q_{DA} must be carried through all 24 ID stages as a state variable, as they are used both for the double-settlement profit at each corresponding ID market. The ID1 bidding quantity is likewise passed forward as part of the state variable, ensuring their use in the profit calculation for ID1 in the next

stage. Hence, the state variable delivered to the next stage is $x_0^\ell = (q_{DA}, q_{ID,1}^\ell)$.

4.2.2. ID stage

The ECTG function after realizing P_{DA}^ℓ can be formulated as

$$\mathcal{Q}_0^\ell(x_0^\ell) := \frac{1}{N_{ID}} \sum_{j \in [N_{ID}]} Q_1^\ell(x_0^\ell;\, \xi_1^{\ell,j}) \tag{3}$$

where for each stage $t \in \mathbb{T}$,

$$Q_t^\ell(x_{t-1}^\ell;\, \xi_t^{\ell,j}) := \max \quad (q_{DA,t} - q_{ID,t}^\ell)\, P_{ID,t}^{\ell,j} - \gamma \left| u_t^\ell - q_{D,t}^\ell \right| + \mathcal{Q}_t^\ell(x_t^\ell) \tag{4a}$$

$$\text{s.t.} \quad q_{ID,t+1}^\ell \in \left[q_{DA,t+1} - w,\, q_{DA,t+1} + w\right] \tag{4b}$$

$$s_t^\ell = s_{t-1}^\ell + c_t^\ell - d_t^\ell \tag{4c}$$

$$s_t^\ell \in [\underline{s}, \overline{s}] \tag{4d}$$

$$c_t^\ell \le g_t^\ell \tag{4e}$$

$$g_t^\ell \le E_{DA,t} \cdot \delta_{E,t}^{\ell,j} \tag{4f}$$

$$u_t^\ell = g_t^\ell + d_t^\ell - c_t^\ell \tag{4g}$$

$$q_{D,t}^\ell = q_{ID,t}^\ell \cdot (1 + \delta_{C,t}^{\ell,j}) \tag{4h}$$

$$q_{ID,t}^\ell, c_t^\ell, d_t^\ell, u_t^\ell \ge 0 \tag{4i}$$

The above formulation covers the period from the realization of the ID price to the closure of the ID t market. The objective function consists of three parts: the double-settlement adjustment, the imbalance penalty, and the ECTG function $\mathcal{Q}_t^\ell(x_t^\ell)$.

The double-settlement term $(q_{DA,t} - q_{ID,t}^\ell)\, P_{ID,t}^{\ell,j}$ accounts for the difference between the awarded quantities in the DA and ID markets, multiplied by the ID price, in order to eliminate double counting between DA and ID revenues. The imbalance penalty term $-\gamma \left| u_t^\ell - q_{D,t}^\ell \right|$ penalizes deviations between actual dispatch amount and the dispatch instruction, for both over- and under-delivery.

Constraint (4b) describe the bidding rules for the IDt market. Generation limits, ESS operation, curtailment, and dispatch are represented in (4d)–(4h). The charging amount cannot exceed the available production, and the production amount cannot be larger than the PV generation amount. The variables that must be carried forward to the next stage are represented as the state variable $x_t^\ell = (q_{DA},\, q_{ID,t+1}^\ell,\, s_t^\ell)$. Note that in the experiments, the previous and current-stage values of the DA bidding amount profile q_{DA} can be discarded, since they are not required for subsequent stages.

4.3. Reformulation

Imbalance penalty: The imbalance penalty described in (4a) can be generalized as $-\gamma|u-q|$. We introduce two new continuous variables $\overline{\mu} = \max(u-q, 0)$ and $\underline{\mu} = \max(q-u, 0)$. There is no need to introduce a binary variable or big-M constant, since this term is to be minimized. Thus, it can be reformulated as follows:

$$\begin{cases} \overline{\mu} \geq 0, \quad \overline{\mu} \geq u - q, \\ \underline{\mu} \geq 0, \quad \underline{\mu} \geq q - u. \end{cases} \tag{5}$$

SoC value of the last stage: The SoC value in the last stage must be fixed at $0.5S$. However, directly imposing the constraint $s = 0.5S$ in (4) may lead to recourse infeasibility when solving the problem sequentially during the forward pass. To address this issue, we introduce a penalty term $\beta\,|s - 0.5S|$ in the objective function of (4). This penalizes deviations from the desired value and enables the impact of terminal SoC deviations to propagate to earlier-stage decisions during the iterative algorithmic procedure.

We now establish the following lemma, which shows that the multistage model satisfies relatively complete recourse.

Lemma 1. *Model* (1)–(4) *satisfies relatively complete recourse.*

Proof. Since problem (2) trivially has relatively complete recourse, it suffices to show that problem (4) admits a feasible solution for any stage n.

For any realization of the uncertainty $\xi_t^{\ell,j} = (P_{ID,t}^{\ell,j}, \delta_{E,t}^{\ell,j}, \delta_{C,t}^{\ell,j})$ and any feasible state $x_{t-1}^{\ell} = (q_{DA}, q_{ID,t}^{\ell}, s_{t-1}^{\ell})$, where $q_{DA} \in [0, E_{DA,t} + B]$, $s_{t-1}^{\ell} \in [\underline{s}, \overline{s}]$, $\delta_{E,t}^{\ell,j} \geq 0$, and $\delta_{C,t}^{\ell,j} \in [-1, 0]$.

Consider the following feasible solution:

$$q_{ID,t}^{\ell} = q_{DA}, \quad c_t = d_t = g_t = u_t = 0, \quad s_t^{\ell} = s_{t-1}^{\ell}.$$

It is straightforward to verify that this solution satisfies all constraints in (4) for any admissible realization of ξ_t^{ℓ}. □

5. Solution Approach

Following the notation convention in Zou et al. (2019), let y_t^{ℓ} denote the local decision variables and z_t^{ℓ} the auxiliary continuous variables associated with each stage problem Q_t^{ℓ}. The auxiliary

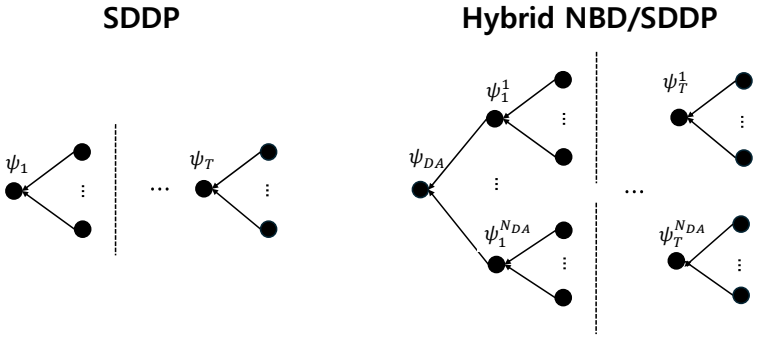

Figure 3: Illustration of the difference between SDDP and hybrid NBD/SDDP algorithms.

variables appear in the constraint $z_t^\ell = x_{t-1}^\ell$ and are relaxed when computing the cutting-plane coefficients in the backward steps. In addition, let $f_t^\ell(x_t^\ell, y_t^\ell, z_t^\ell;\, \xi_t^\ell)$ represent the local objective function and $X_t^\ell(\xi_t^\ell)$ the feasible region. Then, for each stage $t \in \mathbb{T}$, the problem Q_t^ℓ can be formulated as follows.

$$Q_t^\ell(x_{t-1}^\ell;\, \xi_t^\ell) := \max \quad f_t^\ell(x_t^\ell, y_t^\ell, z_t^\ell;\, \xi_t^\ell) \;+\; \mathcal{Q}_t^\ell(x_t^\ell) \tag{6a}$$

$$\text{s.t.} \quad (x_t^\ell, y_t^\ell, z_t^\ell) \in X_t^\ell(\xi_t^\ell) \tag{6b}$$

$$z_t^\ell = x_{t-1}^\ell \tag{6c}$$

Since Q_0^ℓ has a different structure from Q_t^ℓ for $t \in \mathbb{T}$, we denote $\bar{\mathbb{T}} := \mathbb{T} \cup \{0\}$, $\xi_0 = P_{DA}$ and $x_{-1} = x_{DA}$ for notational convenience.

5.1. Stage-wise hybrid NBD/SDDP

Typical SDDP-type algorithms assume stage-wise independence. However, due to the DA–ID dependencies in our problem, SDDP-type algorithms cannot be applied directly. Specifically, different realizations of the DA price P_{DA}^ℓ may lead to different ECTG functions in the corresponding ID stages; that is, $\mathcal{Q}_t^\ell(\cdot) \neq \mathcal{Q}_t^{\ell'}(\cdot)$ for $\ell \neq \ell'$. Thus, we approximate the ID-stage ECTG functions separately for each DA-stage node.

Once the DA price P_{DA}^ℓ is realized, we assume that the ID recombining scenario tree $\mathcal{T}^\ell$ satisfies stage-wise independence; that is, the ECTG functions depend only on the DA price profile P_{DA}^ℓ and the stage. Thus the number of ECTG functions that need to be approximated through the algorithm is $N_{DA}T + 1$. Figure 3 illustrates the difference between the SDDP and hybrid NBD/SDDP algorithms.

5.1.1. Forward step

Each iteration of hybrid NBD/SDDP begins by sampling a subset of scenario from each $\mathcal{T}^\ell$ for every $\ell \in [N_{DA}]$. In the forward step, the algorithm starts by solving the DA-stage problem using an approximate concave piecewise linear ECTG function. The DA-stage problem at iteration I is formulated as follows:

$$(\underline{\mathcal{V}}(\psi^I_{DA})) \quad \max_{x_{DA} \in X_{DA}} \quad 0 + \psi^I_{DA}(x_{DA}) \tag{7}$$

Where the approximate ECTG function $\psi^I_{DA}(\cdot)$ is defined as

$$\psi^I_{DA}(x_{DA}) := \max \left\{ \theta_{DA} \leq \frac{1}{N_{DA}} \sum_{\ell \in [N_{DA}]} \left(v_0^{i,\ell} + \left(\pi_0^{i,\ell}\right)^\top x_{DA} \right), \forall i \leq I \right\} \tag{8}$$

The function $\psi^i_{DA}(\cdot)$ represents the concave upper approximation of the true expected cost-to-go function, $\mathcal{Q}_{DA}(\cdot)$.

Next, for each DA price profile P^ℓ_{DA}, given a feasible solution $\tilde{x}_{DA}$ obtained from $\underline{\mathcal{V}}(\psi^I_{DA})$ and a sampled scenario path $\tilde{\Omega}^\ell = \left\{\tilde{\xi}^\ell_t\right\}_{t \in \mathbb{T}}$ drawn from $\mathcal{T}^\ell$, the algorithm proceeds stage by stage from $t = 0$ to $t = T$ along $\tilde{\Omega}^\ell$, solving an approximate dynamic programming recursion at each stage. The stage problem $\underline{Q}^\ell_t$ in the I-th iteration takes following form:

$$\underline{Q}^\ell_t(\tilde{x}^\ell_{t-1}, \psi^{I,\ell}_t; \tilde{\xi}^\ell_t) := \max_{x_t, y_t, z_t} \quad f^\ell_t(x_t, y_t, z_t; \tilde{\xi}^\ell_t) \; + \; \psi^{I,\ell}_t(x_t) \tag{9a}$$

$$\text{s.t.} \quad (x_t, y_t, z_t) \in X^\ell_t(\tilde{\xi}^\ell_t) \tag{9b}$$

$$z_t = \tilde{x}^\ell_{t-1} \tag{9c}$$

The concave upper-approximation ECTG function $\psi^{i,\ell}_t(\cdot)$ is defined as:

$$\psi^{I,\ell}_t(x_t) := \max \left\{ \theta_t \leq \frac{1}{N_{ID}} \sum_{j \in [N_{ID}]} \left(v^{i,\ell,j}_{t+1} + \left(\pi^{i,\ell,j}_{t+1}\right)^\top x_t \right), \forall i \leq I \right\} \tag{10}$$

where $\psi^{i,\ell}_T \equiv 0$ for all i.

After completing a forward iteration, we obtain a feasible solution $\tilde{x}_{DA}$ and $\left\{\tilde{x}^\ell_t\right\}_{t \in \bar{\mathbb{T}}}$ for every $\ell \in [N_{DA}]$.

5.1.2. Backward Step

For each $\ell \in [N_{DA}]$, the backward step begins at the final stage T. Given the solution $\tilde{x}^{\ell}_{T-1}$ and the parameters of all N_{ID} branches $\left\{\xi^{\ell,j}_{T}\right\}_{j\in[N_{\mathrm{ID}}]}$, solving $\underline{Q}^{\ell}_{T}\left(\tilde{x}^{\ell}_{T-1}, \psi^{I,\ell}_{T}; \xi^{\ell,j}_{T}\right)$ for each j yields the Benders' cut coefficients $\left(v^{I,\ell,j}_{T}, \pi^{I,\ell,j}_{T}\right)$. Here, the gradient value $\pi^{I,\ell,j}_{T}$ corresponds to the dual variable associated with the constraint $z^{\ell}_{T} = x^{\ell}_{T-1}$. Then the inequalities are then aggregated into a single inequality to add one inequality in (10). Consequently, the upper approximation of $\mathcal{Q}^{\ell}_{T-1}(\cdot)$ is updated from $\psi^{I,\ell}_{T-1}$ to $\psi^{I+1,\ell}_{T-1}$. The backward pass then proceeds stage-wise from $T-1$ to $t=1$.

Once $\psi^{I+1,\ell}_{0}$ has been obtained for every $\ell \in [N_{DA}]$, and given the solution $\tilde{x}_{DA}$ and DA price profiles $\left\{P^{\ell}_{DA}\right\}_{\ell\in[N_{DA}]}$, Solving $\underline{Q}^{\ell}_{0}\left(\tilde{x}_{DA}, \psi^{I+1}_{0}; P^{\ell}_{DA}\right)$ yields the Benders' cut coefficients $\left(v^{I,\ell}_{0}, \pi^{I,\ell}_{0}\right)$ for all $\ell \in [N_{DA}]$, and they are aggregated to add one inequality in (8). Then, the upper approximation of $\mathcal{Q}_{DA}(\cdot)$ is updated from $\psi^{i}_{DA}(\cdot)$ to $\psi^{i+1}_{DA}(\cdot)$. A complete description of the forward pass and the backward pass of hybrid NBD/SDDP algorithm is given in Algorithm 1.

Since Lemma 1 establishes relatively complete recourse, the Benders' decomposition relies solely on optimality cuts, with no need for feasibility cuts. Such optimality cut coefficients are calculated as :

$$\pi_t = \arg\min_{\pi} R_t, \quad v_t = \underline{Q}_t\left(\tilde{x}_{t-1}, \psi^{i}_{t}, \xi_t\right) - \pi_t^{\top}\tilde{x}_{t-1}$$

where R_t denotes the dual problem of $\underline{Q}_t\left(\tilde{x}_{t-1}, \psi^{i}_{t}, \xi_t\right)$ obtained by relaxing the constraint $z_t = \tilde{x}_{t-1}$.

5.1.3. Stopping criteria and finite convergence of hybrid NBD/SDDP

A typical stopping criterion for SDDP-type algorithms is to terminate when the gap between the upper bound (UB) and the lower bound (LB) becomes sufficiently small. LB is estimated using the sample mean profit, while UB is obtained by solving the first-stage problem (7) with the approximated ECTG function of the DA stage (ψ^{i+1}_{DA}). Unlike the standard SDDP algorithm, which requires generating multiple samples from a single recombining scenario tree, the proposed approach does not require additional sampling, as N_{DA} samples are already available from the day-ahead price profiles. Nevertheless, since the lower bound is a statistical estimate, care must be taken when using the UB–LB gap as a stopping criterion.

Regarding finite convergence, we first show that the proposed hybrid NBD/SDDP can be interpreted as a reduced form of the NBD algorithm.

Lemma 2. *The stage-wise hybrid NBD/SDDP algorithm is a reduced form of the NBD algorithm.*

Algorithm 1 The stage-wise hybrid NBD/SDDP Method

Initialize: LB $\leftarrow -\infty$, UB $\leftarrow \infty$, $i \leftarrow 1$, and the initial upper approximations $\{\psi_{DA}^1(\cdot)\}$ and $\{\psi_t^{1,\ell}(\cdot)\}_{t\in\bar{\mathbb{T}}} \quad \forall \ell \in [N_{DA}]$
while stopping criterion is not satisfied **do**
  Sample scenario $\tilde{\Omega}^\ell = \{\tilde{\xi}_1^\ell, \ldots, \tilde{\xi}_T^\ell\} \quad \forall \ell \in [N_{DA}]$
  ▷ Forward step
  for $\ell = 1, \ldots, N_{DA}$ **do**
    Solve forward problem $\underline{\mathcal{V}}(\psi_{DA}^i)$
    Collect solution $\tilde{x}_{DA}$
    for $t = 0, \ldots, T$ **do**
      Solve forward problem $\underline{Q}_t^\ell(\tilde{x}_{t-1}^\ell, \psi_t^{i,\ell}, \tilde{\xi}_t^\ell)$
      Collect solution $\tilde{x}_t^\ell, \tilde{y}_t^\ell, \tilde{z}_t^\ell$
    end for
    $LB \leftarrow \sum_{t\in\bar{\mathbb{T}}}, f(\tilde{x}_t^\ell, \tilde{y}_t^\ell, \tilde{z}_t^\ell; \tilde{\xi}_t^\ell)$
  end for
  ▷ Backward step
  for $\ell = 1, \ldots, N_{DA}$ **do**
    for $t = T, \ldots, 1$ **do**
      for $j = 1, \ldots, N_{ID}$ **do**
        Solve $\underline{Q}_t^\ell\left(\tilde{x}_{t-1}^\ell, \psi_t^{i+1,\ell}, \xi_t^{\ell,j}\right)$ and its dual problem.
        Collect cut coefficients $(v_t^{i,\ell,j}, \pi_t^{i,\ell,j})$
      end for
      Add cut in the form of (10) to $\psi_{t-1}^{i,\ell}$ to get $\psi_{t-1}^{i+1,\ell}$
    end for
  end for
  for $\ell = 1, \ldots, N_{DA}$ **do**
    Solve $\underline{Q}_0^\ell\left(\tilde{x}_{DA}^\ell, \psi_0^{i+1,\ell}, P_{DA}^\ell\right)$ and its dual problem.
    Collect cut coefficients $(v_0^{i,\ell}, \pi_0^{i,\ell})$
  end for
  Add cut in the form of (8) to ψ_{DA}^i to get ψ_{DA}^{i+1}
  Solve $\underline{\mathcal{V}}(\psi_{DA}^i)$ and update UB with the optimal value
  $i \leftarrow i + 1$
end while

Proof. Consider a composite scenario tree constructed by combining the DA nodes $\mathcal{P}_{DA}$ with the corresponding scenario trees $\mathcal{T}^\ell$ for each P_{DA}^ℓ. This yields a large-scale (possibly non-recombining) scenario tree representation of the problem.

Even if the conditional independence of ID parameters given the DA node does not hold, the resulting problem can still be formulated and solved using NBD on this expanded scenario tree. Therefore, the proposed hybrid NBD/SDDP algorithm can be viewed as a computationally reduced implementation of this full NBD formulation. □

Based on this result, we establish finite convergence of the proposed method.

Proposition 1. *The stage-wise hybrid NBD/SDDP algorithm converges to an optimal solution of* (1)–(4) *in a finite number of iterations.*

Proof. The result follows from Theorem 2 in Zou et al. (2019), which establishes finite convergence of the NBD algorithm under three conditions: (i) the sampling procedure in the forward step is performed with replacement, (ii) the cuts generated in the backward step are valid, tight, and finite, and (iii) for a given previous-stage solution and ECTG approximation, the nodal problem is solved consistently to the same optimal solution.

In our setting, condition (i) is satisfied by the sampling procedure used in the forward pass, condition (ii) follows from the fact that Benders' cuts provide tight and valid cuts in the MSLP problem, and condition (iii) holds since the nodal problems are solved using a deterministic solver (e.g., Gurobi), which ensures consistent selection among optimal solutions. Therefore, the required conditions are satisfied, and the result follows. □

5.2. Approximation

Since each DA price profile $P_{DA}^{\ell} = \left(P_{DA,1}^{\ell}, \ldots, P_{DA,T}^{\ell}\right)$ consists of $T = 24$ elements, N_{DA} must be sufficiently large; however, this renders the algorithm computationally intractable. To address this, Our aim is to run the approximated hybrid NBD/SDDP algorithm on the original DA profile nodes $\mathcal{P}_{DA}$ and the reduced scenario trees $\left\{\hat{\mathcal{T}}^k\right\}_{k\in[K]}$ to obtain the DA-stage ECTG function $\hat{\psi}_{DA}$.

First, we reduce $\mathcal{P}_{DA}$ to $\hat{\mathcal{P}}_{DA} := \left\{\hat{P}_{DA}^1, \ldots, \hat{P}_{DA}^K\right\}$ with K representative elements. We denote the index set of $\hat{\mathcal{P}}_{DA}$ by $[K] := \{1, \ldots, K\}$, and define the ID recombining scenario tree $\hat{\mathcal{T}}^k$ corresponding to each $\hat{P}_{DA}^k \in \hat{\mathcal{P}}_{DA}$. Due to the mismatch between the number of given DA nodes (N_{DA}) and the reduced ID scenario trees (K), we allocate the nearest scenario tree to each DA price profile by finding the closest DA price profile within the reduced DA nodes $\hat{\mathcal{P}}_{DA}$. Then, we run Algorithm 1 as if each corresponding ID scenario tree for each DA price profile were its nearest scenario tree. Subsequently, we obtain an ECTG function $\hat{\psi}_{DA}$, and by using it, we can also obtain the DA-stage solution $\tilde{x}_{DA}$.

Unfortunately, the ID-stage ECTG approximations $\{\hat{\psi}_t^k\}_{t\in\bar{\mathbb{T}}}$ cannot be directly used in the evaluation phase, because a realized DA price P_{DA} sampled from $\mathcal{P}_{DA}$ does not necessarily belong to the reduced DA price set $\hat{\mathcal{P}}_{DA}$. Therefore, we employ the algorithm solely to obtain DA-stage solutions. After the realization of the DA price P_{DA}, the subsequent ID stage problems are solved using a standard SDDP to get the ID-stage ECTG approximations $\{\psi_t^k\}_{t\in\bar{\mathbb{T}}}$. The complete approximation procedure is given in Algorithm 2.

When selecting the approximation parameter K, an inherent trade-off must be considered. The value of K controls the balance between approximation accuracy and computational tractability.

Algorithm 2 Approximation

Reduce $\mathcal{P}_{DA} = \left\{P_{DA}^1, \ldots, P_{DA}^{N_{DA}}\right\}$ to $\hat{\mathcal{P}}_{DA} = \left\{\hat{P}_{DA}^1, \ldots, \hat{P}_{DA}^K\right\}$
Generate ID scenario trees $\left\{\hat{\mathcal{T}}^k\right\}_{k\in[K]}$ from $\hat{\mathcal{P}}_{DA}$
Allocate the nearest scenario tree $\hat{\mathcal{T}}^k$ to each DA price node P_{DA}^ℓ where $k = \arg\min_k \left\{\left\|P_{DA}^\ell - \hat{P}_{DA}^k\right\|\right\}$.
Run Algorithm 1 on $\mathcal{P}_{DA}$ and $\left\{\hat{\mathcal{T}}^k\right\}_{k\in[K]}$
Collect approximate ECTG function $\hat{\psi}_{DA}$
Solve forward problem $\underline{\mathcal{V}}(\hat{\psi}_{DA})$ and collect solution $\tilde{x}_{DA}$
for $\ell = 1, \ldots, N_{DA}$ **do**
    Run standard SDDP given DA solution $\tilde{x}_{DA}$, DA price P_{DA}^ℓ, and ID scenario tree $\mathcal{T}^\ell$.
    Collect approximate ECTG functions $\{\psi_t^\ell\}_{t\in\bar{\mathbb{T}}}$
end for

Larger values of K provide a richer representation of the DA–ID dependency, as they allow for more detailed modeling of ID scenario trees conditional on the DA price profile, thereby improving the quality of the DA-stage ECTG approximation. However, they also significantly increase the computational burden, which may limit the ability to incorporate updated information for predicting uncertain parameters. Conversely, smaller values of K reduce computational cost at the expense of approximation accuracy. This trade-off motivates the numerical investigation of different values of K in the computational experiments.

Remark 1. *When hybrid NBD/SDDP is approximated with $K = 1$, it is equivalent to SDDP.*

This equivalence arises because, when $K = 1$, there is only one ID scenario tree and a single set of ECTG functions to be approximated. In this case, the approximated hybrid NBD/SDDP behaves as if the ID scenario parameters are independent of the DA price profile, thereby satisfying the stage-wise independence assumption. Consequently, the number of ECTG functions reduces to $N_{DA} + TN_{ID}$. However, this simplification may lead to a significant loss in approximation accuracy.

For illustration, see Figure 4, which compares the approximations under different values of the approximation size K.

6. Numerical Experiments

In this section, we conduct three experiments to evaluate our proposed algorithm. First, we compare its performance with existing methods in terms of the mean profit and examine the factors underlying the performance differences. Second, we investigate the effectiveness of the approximation introduced in Section 5.2 by varying the approximation level K and examining the trade-off

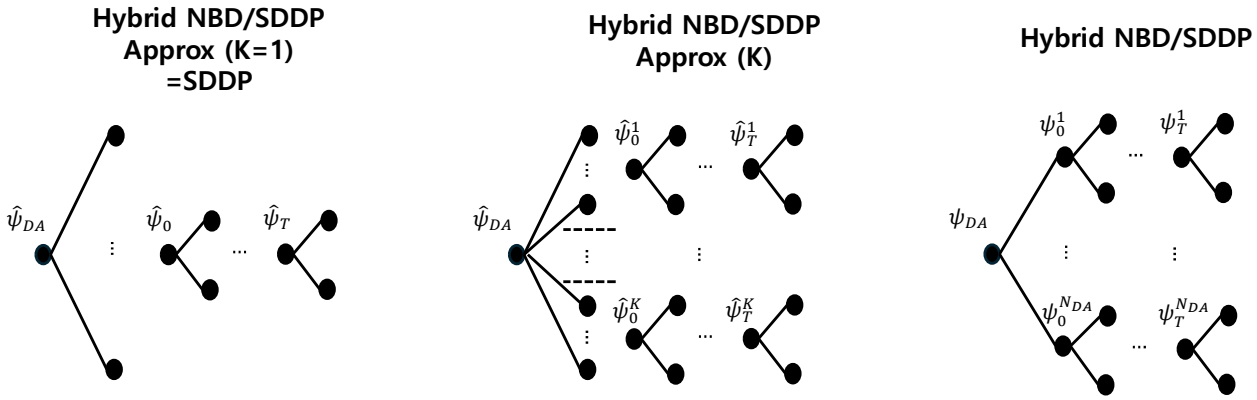

Figure 4: Illustration of hybrid NBD/SDDP with $K = 1$ (equivalent to SDDP), hybrid NBD/SDDP with approximation size K, and hybrid NBD/SDDP without approximation.

between performance and computational time. Finally, we analyze the performance gap between stage-wise hybrid NBD/SDDP and SDDP under different levels of DA–ID dependency by adjusting W. Throughout the experiments, performance is measured as the mean profit over 200 sampled scenarios. We describe the experimental settings in Section 6.1 and analyze the numerical results in Section 6.2.

6.1. Experimental settings

We set the total capacity of PV generators to 21,022 kW, which a domestic VPP company is considering. The maximum and minimum ESS storage capacities are set to 18,920 kWh and 2,100 kWh, which are approximately 0.9 and 0.1 of the PV capacity, respectively. The PCS capacity of the ESS is set to 7,000 kW, which is about one-third of the PV capacity. The bound on the deviation of ID bidding amount from DA bidding amount is set to 6,300 kWh, which is about 30% of the PV capacity.

The data on electricity prices and PV generation forecast profiles are obtained from the DA and ID electricity markets on Jeju Island, South Korea, over the period from March 2024 to March 2025. During this period, negative prices are frequently observed, and both DA and ID prices range from –80 KRW/kWh to 200 KRW/kWh. The over and under delivery penalty coefficient is set equal to the maximum market price of 200 KRW/kWh.

6.1.1. Scenario tree

We consider an empirical distribution of DA price profiles with sample size $N_{DA} = 50$. For each DA price profile $P^{\ell}_{DA} = \left\{P^{\ell}_{DA,1}, \ldots, P^{\ell}_{DA,24}\right\}$, each DA price $P^{\ell}_{DA,t}$ is sampled from the conditional distribution of DA price given the DA generation forecast $E_{DA,t}$, which is estimated using paired

data of generation forecasts and DA prices. A higher forecast scale results in a larger proportion of negative DA prices during daytime hours.

For each DA price P_{DA}^{ℓ}, an ID recombining scenario tree $\mathcal{T}^{\ell}$ is generated based on the conditional distributions of the ID-stage parameters $\xi_t^{\ell} = \left(P_{ID,t}^{\ell},\, \delta_{E,t}^{\ell},\, \delta_{C,t}^{\ell}\right)$. First, the conditional distribution of ID prices P_{ID} given the corresponding DA price P_{DA} is modeled using a truncated Gaussian mixture model (TGMM), fitted via the Expectation–Maximization algorithm on historical DA–ID price pairs from the Jeju Island market. To construct the conditional distribution of ID prices given DA prices and to control the level of DA–ID dependency, we partition the DA price interval $[-80, 200]$ into W subintervals. Then, a separate TGMM is then trained for each subinterval using ID price data whose corresponding DA prices fall within that subinterval. Thus, the level of DA–ID dependency can be controlled through the choice of W, with larger values of W inducing stronger dependency.

The curtailment ratio δ_C is modeled as the product of two independent random variables: a uniform distribution $\mathcal{U}(-1, 0)$ and a Bernoulli distribution $\mathit{Bernoulli}(p)$. To reflect the higher likelihood of curtailment when ID prices are low or negative, the parameter p is specified as a decreasing step function of the ID price, and hence depends indirectly on the DA price. For example, when the ID price is below -60, $p = 0.75$, whereas when the ID price exceeds 20, $p = 0.05$. Lastly, the PV electricity generation is modeled as the product of the fixed generation forecast value $E_{DA,t}$ and the generation-to-forecast ratio δ_E, where δ_E follows a truncated normal distribution $\mathcal{N}(1, \sigma_E^2)$. The variance σ_E^2 is estimated from historical pairs of PV electricity generation forecasts and realizations.

Each ID recombining scenario tree $\mathcal{T}^{\ell}$ is constructed by first generating 300 samples at each stage, and then applying the scenario reduction algorithm of Heitsch and Römisch (2003) to reduce the number of child nodes to $N_{ID} = 30$ at each stage. Figure 5 illustrates the DA price distributions and the corresponding ID price distributions under different levels of DA–ID dependency, controlled by W. When $W = 1$, no DA–ID dependency is imposed, and the ID price distribution remains nearly identical across time. When $W = 2$, where the DA price is partitioned into negative and positive regions, the model captures only coarse dependency (e.g., the tendency for ID prices to share the same sign as DA prices), but fails to represent finer dependencies. In contrast, when $W = 5$, the ID price distributions capture more detailed dependencies, as evidenced by the increased variability during the periods 6h–9h and 15h–18h. However, using larger values of W (e.g., $W \geq 6$) leads to distortions in the estimated conditional distributions due to insufficient data within each

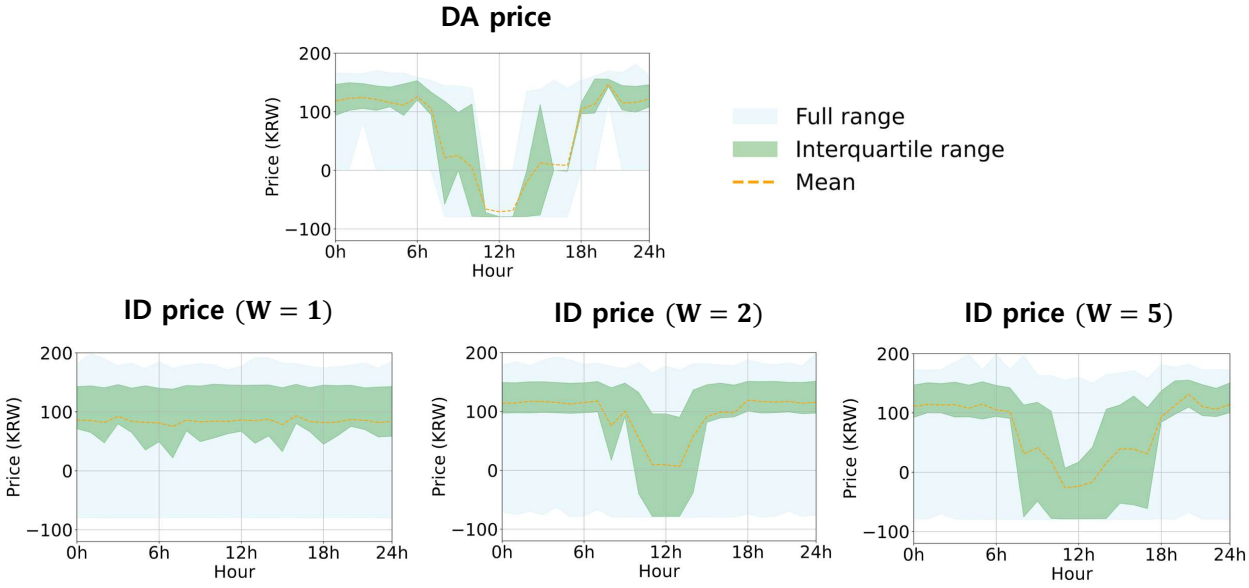

Figure 5: Comparison between different distributions.

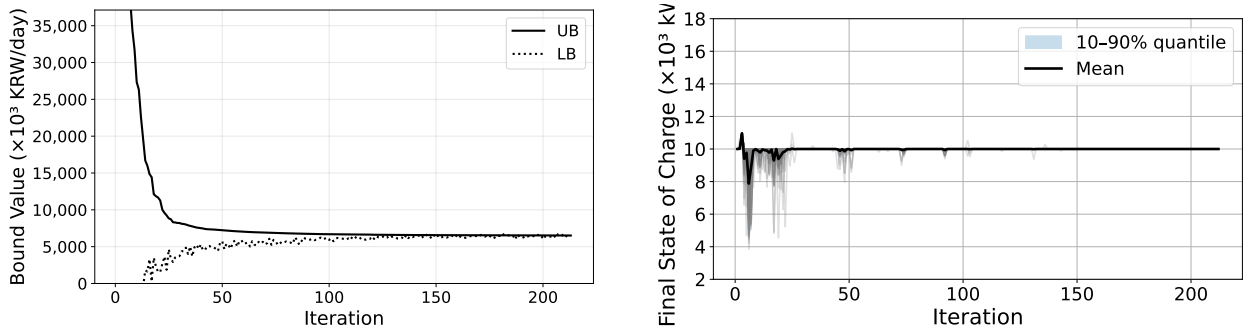

(a) Convergence behavior of the upper and lower bounds.

(b) Convergence of the terminal state of charge (SoC).

Figure 6: Convergence behavior of the proposed method in terms of optimality gap and terminal SoC.

subinterval. Therefore, we adopt $W = 5$ as the baseline setting of ID stage distributions.

6.1.2. Algorithmic settings of the stage-wise hybrid NBD/SDDP

The stopping criterion of the stage-wise hybrid NBD/SDDP is based on both the optimality gap and the convergence of the terminal SoC. Specifically, the algorithm terminates when $(\text{UB} - \text{LB})/\text{UB} \leq 0.1\%$ and the SoC at the last stage attains its target value of $0.5S$ for all sampled scenario paths. Figure 6 illustrates the convergence behavior. The optimality gap $(\text{UB} - \text{LB})/\text{UB}$ approaches zero, supporting Proposition 1. In addition, the terminal SoC converges to its target value, indicating that the penalty for deviation from $0.5S$ is effectively propagated across stages during the backward pass. Notably, this convergence is typically achieved prior to the optimality gap reaching zero.

For other implementation details, DA price profiles are reduced using the scenario reduction algorithm of Heitsch and Römisch (2003) when applying the approximation described in Section 5.2. To improve computational efficiency, subproblems sharing a common parent node in the backward step are solved in parallel using Python's multiprocessing module. In addition, cut selection is applied across all stages to enhance performance. The effectiveness of cut selection in both single-cut and multi-cut settings within SDDP-type algorithms has been demonstrated in De Matos et al. (2015).

The algorithm is implemented in Python 3.12, and all MIP and LP subproblems are solved using Gurobi 12.0.1. All experiments are conducted on a workstation equipped with a 13th Gen Intel® Core™ i7-13700KF CPU (12 cores, 24 threads) at 2.40 GHz and 128 GB of RAM.

6.2. Results

6.2.1. Performance of stage-wise hybrid NBD/SDDP

Table 3 reports the overall performance results for our proposed algorithm and three benchmarks, Deterministic rolling horizon method (Det), two-stage Rolling horizon method (2SP), and SDDP. SDDP corresponds to the approximated hybrid NBD/SDDP with $K = 1$. The relative performance of benchmarks is computed with respect to the hybrid NBD/SDDP. The average profit can be decomposed into two components: DA and ID settlement profit and the imbalance penalty. The DA and ID settlement profit reflects the quality of bidding decisions in terms of maximizing profit with respect to DA and ID prices, while the imbalance penalty captures the algorithm's ability to manage deviations caused by PV generation variability and curtailment events in the ID stages. The results indicate that the hybrid NBD/SDDP not only achieves higher settlement profits but also more effectively mitigates imbalance penalties, demonstrating superior handling of DA and ID price uncertainties and ID-stage deviations caused by PV generation variability and curtailments.

Table 3: Performance of our proposed algorithm and benchmarks

Solution algorithm	Avg. profit (10^3 KRW/day)	Relative performance (%)	DA and ID settlement (10^3 KRW/day)	Imbalance penalty (10^3 KRW/day)
Det	5,671.88	-13.26	5,864.16	-192.28
2SP	5,976.75	-8.59	6,095.45	-118.71
SDDP	5,379.53	-17.73	5,680.89	-220.97
Hybrid NBD/SDDP	**6,538.73**	**0.00**	**6,629.58**	**-20.18**

To identify the underlying causes of these performance differences, Figure 7 illustrates the op-

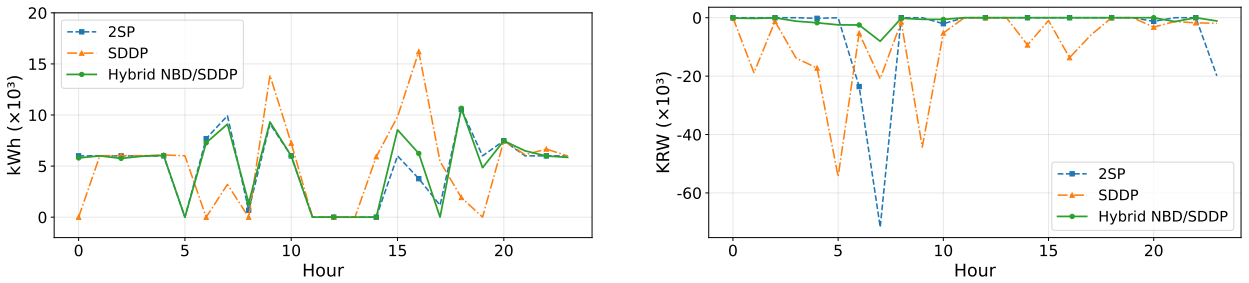

(a) Optimal DA bidding decisions across hours.

(b) Mean imbalance penalty across hours.

Figure 7: Optimal DA bidding decisions and imbalance penalties across hours.

timal DA bidding decisions obtained from the 2SP, SDDP, and hybrid NBD/SDDP algorithms, together with the mean imbalance penalty across hours within a day, averaged over sampled scenarios. From Figure 7(a), we observe that the DA bidding decisions of 2SP and hybrid NBD/SDDP exhibit similar patterns, as both algorithms capture the DA–ID dependency. In contrast, SDDP yields DA bidding decisions that are less consistent with those obtained from hybrid NBD/SDDP, since SDDP fails to incorporate DA–ID dependency. This is because SDDP assumes that all DA price profiles share a common ID recombining scenario tree; as a result, it implicitly considers combinations of DA price profiles and ID scenario realizations that are unlikely under the true joint distribution of DA prices and the corresponding ID-stage parameters.

Figure 7(b) highlights the resulting operational consequences. The 2SP approach, despite producing reasonable DA decisions, fails to fully account for future recourse actions in the ID stage, which leads to occasional but significant penalty spikes at specific hours. On the other hand, SDDP incorporates recourse decisions, but due to its suboptimal DA bidding decisions, it results in frequent imbalance penalties across the horizon. In contrast, the hybrid NBD/SDDP method effectively captures both DA–ID dependency and full recourse in the ID stage, leading to consistently higher DA and ID settlement profits and lower imbalance penalties throughout the day. These results demonstrate that jointly modeling DA–ID dependency and recourse decisions is crucial for achieving superior overall performance across DA and ID stages, thereby highlighting the advantage of the proposed stage-wise hybrid NBD/SDDP method.

6.2.2. Effects of approximation

The running time of the stage-wise hybrid NBD/SDDP until reaching the stopping criterion is 15,182 seconds, which is computationally expensive and may be impractical in real-world settings where the distributions of prices and PV generation forecasts can be updated over time. This

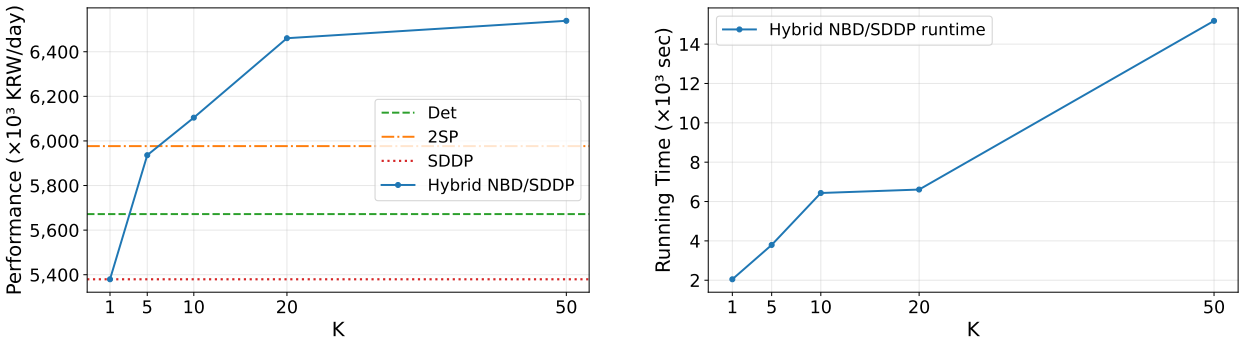

(a) Performance as a function of the approximation parameter K.

(b) Running time as a function of the approximation parameter K.

Figure 8: Trade-off between performance and running time for different values of the approximation parameter K.

motivates the use of the approximation scheme introduced in Section 5.2. In this experiment, we investigate the trade-off between evaluation performance and computational time by varying the approximation parameter K. Figure 8 illustrates the trade-off between evaluation performance and running time for $K \in \{1, 5, 10, 20, 50\}$. Here, $K = 1$ corresponds to the SDDP benchmark, while $K = 50$ represents the full stage-wise hybrid NBD/SDDP method.

Figure 8(a) compares the evaluation performance across different values of K, together with the performance of other benchmark methods (Det, 2SP, and SDDP), shown as horizontal reference lines. As K increases, the evaluation value improves monotonically, indicating that a larger approximation size better captures the underlying DA–ID dependency structure. In particular, $K \geq 10$ already achieves superior performance compared to 2SP. Figure 8(b) shows the corresponding running time of the hybrid NBD/SDDP algorithm. As expected, the computational time increases with K, but remains significantly lower than that of the full model ($K = 50$).

These results demonstrate that the proposed approximation method provides an effective trade-off between solution quality and computational efficiency. For instance, when $K = 20$, the method achieves performance close to the full hybrid NBD/SDDP, while reducing the running time substantially (from 15,182 seconds to 6,608 seconds).

6.2.3. Performance across different DA–ID dependency levels

We investigate how the effectiveness of the stage-wise hybrid NBD/SDDP relative to SDDP varies with the level of DA–ID dependency. Figure 9 compares the performance of SDDP, approximated hybrid NBD/SDDP with $K \in \{5, 10, 20\}$, and the full hybrid NBD/SDDP across different dependency levels.

As shown in Figure 9(b), the relative performance gap increases as the level of DA–ID de-

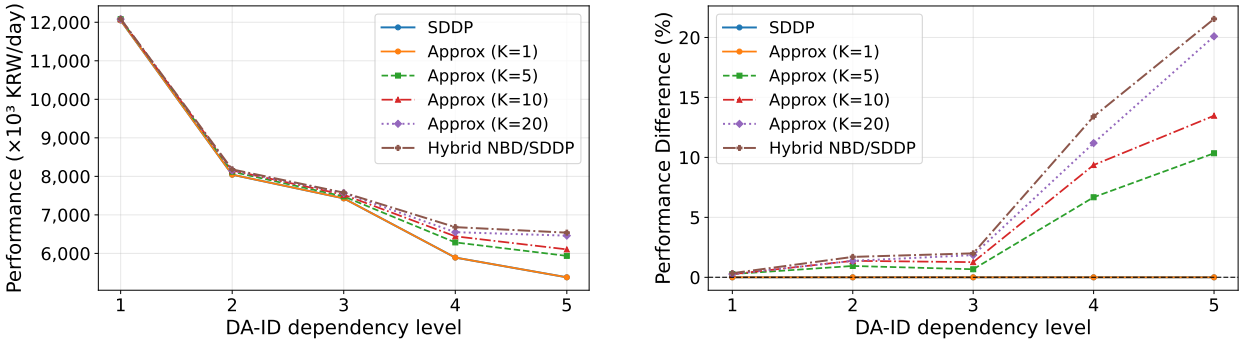

(a) Performance comparison across different DA–ID dependency levels.

(b) Relative performance differences across dependency levels.

Figure 9: Performance comparison among SDDP, approximated hybrid NBD/SDDP, and full hybrid NBD/SDDP under varying DA–ID dependency levels.

pendency becomes stronger. When there is no DA–ID dependency ($W = 1$), i.e., when the ID parameters follow the same distribution across all DA price profile nodes, SDDP and the approximated hybrid NBD/SDDP perform nearly identically to the full hybrid NBD/SDDP. In this case, nodalizing the DA stage provides little to no additional value.

However, as the DA–ID dependency increases, the relative performance difference enlarges substantially. In particular, when $W = 5$, both SDDP and the approximated variants exhibit a clear performance loss compared to the full hybrid NBD/SDDP. This result highlights the importance of explicitly capturing strong nested dependencies, such as DA–ID dependency, and demonstrates the advantage of the proposed stage-wise hybrid NBD/SDDP in such settings.

7. Conclusion

In this paper, we develop a MSLP model tailored to the joint DA–ID bidding and ESS operation problem of a VPP, together with a stage-wise hybrid NBD/SDDP algorithm designed to address the nested DA–ID dependency between DA prices and the corresponding ID-stage parameters. To mitigate the computational burden of the proposed method, we further introduce an approximation scheme.

Numerical experiments demonstrate that the proposed algorithm outperforms benchmarks, including SDDP, due to its ability to capture both nested DA–ID dependency and full recourse in the ID stage. In addition, we show that the proposed approximation scheme provides an effective trade-off between solution quality and computational efficiency. We also demonstrate that the relative advantage of the proposed method over SDDP increases as the level of DA–ID dependency

becomes stronger, highlighting the importance of partially nodalizing specific stages when nested dependencies exist in the underlying stochastic processes.

We acknowledge several limitations of our study and suggest potential extensions for future research. First, this work focuses on DA–ID dependency by nodalizing the DA stage within a single-day framework. However, more complex nested dependency structures may arise in medium- and long-term planning horizons, such as weekday-specific patterns or cross-month dependencies between identical calendar days. Thus, a natural extension would be to generalize the stage-wise hybrid NBD/SDDP framework to selectively nodalize additional stages while decomposing the remaining ones. Second, we do not consider temporal dependencies across stages in this study. However, in practice, ID-stage parameters may exhibit temporal correlations. Thus, the proposed formulation could be extended using TS-SDDP or MC-SDDP (Löhndorf and Shapiro (2019)) to capture such temporal dependencies. Finally, the numerical experiments are conducted under a fixed DA generation forecast and a corresponding discretized scenario tree. However, real-world applications involve varying forecast profiles and richer sources of uncertainty. Thus, incorporating contextual stochastic programming (see the survey in Sadana et al. (2025)) would enable more realistic modeling and out-of-sample evaluation using real-world data.

References

Abgottspon, H., Njálsson, K., Bucher, M.A., Andersson, G., 2014. Risk-averse medium-term hydro optimization considering provision of spinning reserves, in: 2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS), pp. 1–6.

Cordera, F., Moreno, R., Ordoñez, F., 2023. Unit commitment problem with energy storage under correlated renewables uncertainty. Operations Research 71, 1960–1977.

De Matos, V.L., Philpott, A.B., Finardi, E.C., 2015. Improving the performance of stochastic dual dynamic programming. Journal of Computational and Applied Mathematics 290, 196–208.

Fazlalipour, P., Ehsan, M., Mohammadi-Ivatloo, B., 2019. Risk-aware stochastic bidding strategy of renewable micro-grids in day-ahead and real-time markets. Energy 171, 689–700.

Finnah, B., Gönsch, J., Ziel, F., 2022. Integrated day-ahead and intraday self-schedule bidding for energy storage systems using approximate dynamic programming. European Journal of Operational Research 301, 726–746.

Füllner, C., Rebennack, S., 2025. Stochastic dual dynamic programming and its variants: A review. SIAM Review 67, 415–539.

Heitsch, H., Römisch, W., 2003. Scenario reduction algorithms in stochastic programming. Computational Optimization and Applications 24, 187–206.

Heredia, F.J., Cuadrado, M.D., Corchero, C., 2018. On optimal participation in the electricity markets of wind power plants with battery energy storage systems. Computers & Operations Research 96, 316–329.

Hjelmeland, M.N., Zou, J., Helseth, A., Ahmed, S., 2018. Nonconvex medium-term hydropower scheduling by stochastic dual dynamic integer programming. IEEE Transactions on Sustainable Energy 10, 481–490.

Kim, D., Cheon, H., Choi, D.G., Im, S., 2022. Operations research helps the optimal bidding of virtual power plants. INFORMS Journal on Applied Analytics 52, 344–362.

Kim, S., Choi, D.G., 2024. A sample robust optimal bidding model for a virtual power plant. European Journal of Operational Research 316, 1101–1113.

Korea Power Exchange (KPX), 2025. Annual operating performance of the Jeju power system, 2024. Technical Report. Korea Power Exchange (KPX). Republic of Korea. URL: https://www.kpx.or.kr/board.es?mid=a10102000000&bid=0159&tag=&act=view&list_no=74566. in Korean. Accessed May 7, 2026.

Löhndorf, N., Shapiro, A., 2019. Modeling time-dependent randomness in stochastic dual dynamic programming. European Journal of Operational Research 273, 650–661.

Mashhour, E., Moghaddas-Tafreshi, S.M., 2011. Bidding strategy of virtual power plant for participating in energy and spinning reserve markets—part i: Problem formulation. IEEE Transactions on Power Systems 26, 949–956.

Morales, J.M., Conejo, A.J., Pérez-Ruiz, J., 2010. Short-term trading for a wind power producer. IEEE Transactions on Power Systems 25, 554–564.

Neuhoff, K., Ritter, N., Salah-Abou-El-Enien, A., Vassilopoulos, P., 2016. Intraday markets for power: Discretizing the continuous trading. Technical Report EPRG1609. Energy Policy Research Group, University of Cambridge. URL: https://www.jstor.org/stable/resrep30358. accessed May 7, 2026.

Pan, K., Guan, Y., 2022. Integrated stochastic optimal self-scheduling for two-settlement electricity markets. INFORMS Journal on Computing 34, 1819–1840.

Pandžić, H., Kuzle, I., Capuder, T., 2013a. Virtual power plant mid-term dispatch optimization. Applied Energy 101, 134–141.

Pandžić, H., Morales, J.M., Conejo, A.J., Kuzle, I., 2013b. Offering model for a virtual power plant

based on stochastic programming. Applied Energy 105, 282–292.

Pereira, M.V.F., Pinto, L.M.V.G., 1985. Stochastic optimization of a multireservoir hydroelectric system: A decomposition approach. Water Resources Research 21, 779–792.

Pereira, M.V.F., Pinto, L.M.V.G., 1991. Multi-stage stochastic optimization applied to energy planning. Mathematical Programming 52, 359–375.

Rebennack, S., 2016. Combining sampling-based and scenario-based nested benders decomposition methods: Application to stochastic dual dynamic programming. Mathematical Programming 156, 343–389.

Sadana, U., Chenreddy, A., Delage, E., Forel, A., Frejinger, E., Vidal, T., 2025. A survey of contextual optimization methods for decision-making under uncertainty. European Journal of Operational Research 320, 271–289.

Shinde, P., Kouveliotis-Lysikatos, I., Amelin, M., 2022. Multistage stochastic programming for vpp trading in continuous intraday electricity markets. IEEE Transactions on Sustainable Energy 13, 1037–1048.

Wozabal, D., Rameseder, G., 2020. Optimal bidding of a virtual power plant on the spanish day-ahead and intraday market for electricity. European Journal of Operational Research 280, 639–655.

Zou, J., Ahmed, S., Sun, X.A., 2018. Multistage stochastic unit commitment using stochastic dual dynamic integer programming. IEEE Transactions on Power Systems 34, 1814–1823.

Zou, J., Ahmed, S., Sun, X.A., 2019. Stochastic dual dynamic integer programming. Mathematical Programming 175, 461–502.