Analytical Framework for Forecast Horizon Determination of Real-Time Energy Storage System Scheduling in Smart Grids

Nicholas Tetteh Ofoe, *Student Member, IEEE,* Zexin Liu, *Member, IEEE,* Lei Wu, *Fellow, IEEE*
Department of Electrical and Computer Engineering, Stevens Institute of Technology, Hoboken, NJ, USA
Emails: ntofoe@gmail.com; zliu136@stevens.edu; lei.wu@stevens.edu

***Abstract*—Energy Storage Systems (ESSs), which can store electricity when supply is cheap and/or renewable generation is high and release it when supply is expensive and/or demand rises, are essential assets for efficiently and cost-effectively maintaining supply-demand balance in modern power grids. However, a major challenge in operating these systems is deciding how far into the future, known as the forecast horizon, to plan their charging and discharging actions. If the horizon is too short, the ESSs may act myopically and miss profitable opportunities; if too long, the optimization becomes computationally expensive and impractical for real-time applications. This paper presents an analytical framework for determining the minimum forecast horizon that guarantees near-optimal ESS scheduling performance. By reformulating the problem as a deterministic dynamic programming (DP) model and applying Bellman's contraction principle, we derive a closed-form bound that relates the minimum forecast horizon to the discount factor, the reward bound, and an acceptable error tolerance. Major assumptions are explicitly stated and justified, linking modeling parameters such as discount factor and bounded rewards to physical system characteristics like leakage and power limits. Numerical validation using historical electricity price data confirms that the analytical bounds closely match the benchmark brute-force simulation results. The proposed framework replaces expensive brute-force simulation-based tuning with a principled approach that improves efficiency and provides actionable insights for ESS operations in real-time electricity markets.**

***Index Terms*—Forecast horizon, energy storage, rolling horizon, scheduling, state of charge, optimal control.**

NOMENCLATURE

Indices and Sets

H	Decision horizon (hours).
t	Time index (hours).
T	Forecast horizon length (hours).

Parameters

C_t	Electricity price at time t (\$/MWh).
$E_{\min}$	Min energy capacity of the ESS (MWh).
$E_{\max}$	Max energy capacity of the ESS (MWh).
$\underline{P}^c, \overline{P}^c$	Min and max charging power limits of ESS (MW).
$\underline{P}^d, \overline{P}^d$	Min and max discharging power limits of ESS (MW).
$R_{\max}$	Bounded stage reward of $V^T(s)$.
s_0	Initial state-of-charge.
γ	Contraction factor of the Bellman operator, $\gamma \in (0, 1)$.
ε	Allowable sub-optimality tolerance.
η^c, η^d	Charging and discharging efficiencies.
τ	Duration to fully discharge the ESS (hours).
ρ	Leakage factor.

Functions and Variables

$p_t^{\mathrm{c}}, p_t^{\mathrm{d}}$	Charging/discharging power at time t (MW).
$r_\tau(s, u)$	Cumulative reward at the τ step (\$).
r_t	Stage reward of the DP problem at time t, $r_t = (s_t, u_t)$ (\$).
s_t	State of charge at time t.
$T^\star$	Min forecast horizon ensuring ε-optimality (hours).
u_t	ESS control action at time t, $u_t = (p_t^{\mathrm{c}}, p_t^{\mathrm{d}})$ (MW).
$\mathcal{B}$	The Bellman operator.
$f(s_t, u_t)$	State transition equation, $s_{t+1} = f(s_t, u_t)$.
$V^T(s)$	Finite-horizon value function with horizon T (\$).
$V^\infty(s)$	Infinite-horizon value function (\$).

I. INTRODUCTION

The integration of renewable energy into modern power systems has made energy storage systems (ESSs) essential for balancing supply and demand, reducing price volatility, and improving overall reliability. Yet, designing optimal ESS schedules remains challenging as real-world operation is shaped by efficiency losses, state-of-charge (SoC) limits, and degradation effects [1]. One important question is how far ahead to look when making charging and discharging decisions. Although infinite-horizon dynamic programming provides the theoretically ideal answer, it is computationally demanding for practical use, prompting operators to rely instead on finite- or rolling-horizon schemes [2], [3]. This naturally leads to the core problem: *what forecast horizon is long enough to achieve near-optimal performance?*

Earlier works on forecast and solution horizons in nonstationary Markov decision processes [4], [5] showed that, under certain structural properties, a finite planning horizon can approximate the infinite-horizon solution. These ideas have been applied in fields ranging from inventory management and asset control [6] to energy scheduling and ramping services [7], [8], but many ESS applications still rely on heuristic horizons (e.g., 24 or 48 hours) without theoretical justifications.

A recent study [9] addressed this problem for deterministic ESS scheduling by proposing a sufficient condition based on SoC endpoint alignment and an algorithm to identify minimal forecast horizons. While useful, their method cannot effectively identify minimal forecast horizons for cases when such a condition is invalid. The presence of multiple optimal terminal states may also cause the condition to miss valid horizons or overestimate their required length.

This paper takes a different approach by developing an analytical framework that provides a closed-form expression for the minimum forecast horizon using Bellman's contraction

979-8-3315-8138-1

principle [10]. The resulting bound explicitly links T^* to discount factors, bounded rewards, and the allowed deviation in the value function, offering a clear and provable way to determine what finite horizon is adequate. The assumptions behind the model are based on discounting, reward bounds, and decision equivalence, motivated directly by physical ESS operations. We validate the framework using the benchmark brute-force simulations with historical DK1 electricity prices and show that the resulting system-aware horizon estimates are more accurate than those with traditional static discounting.

The rest of the paper is organized as follows. Section II formulates the ESS scheduling model, Section III introduces the proposed analytical framework, Section IV presents numerical results, and Section V discusses broader implications and future extensions.

II. The Real-Time ESS Scheduling Problem

A. Problem Formulation and Minimum Forecast Horizon

We consider the problem of optimally scheduling a price-taking ESS for a large-enough finite planning horizon $\mathcal{T} = \{1, 2, \ldots, T\}$ based on given price forecasts, to mimic the infinite-horizon optimization problem. The ESS is characterized by energy capacity limit, charging and discharging power boundaries, conversion efficiencies, and storage deterioration through leakage. We aim to identify the minimum planning horizon $T^* \leq T$ for which the first time step operation decision under rolling optimization is consistent with the global T-horizon optimal decision, a definition which we pose as the minimum forecast horizon.

The model aims to maximize profit based on given electricity prices $\mathcal{C} = \{C_1, C_2, ..., C_T\} \in \mathbb{R}_+^T$. The standard deterministic ESS model [9] is used to build the optimization model $\mathcal{S}(\mathcal{T}, \mathcal{C})$ as in (1). Constraint (1b) enforces the SoC dynamics, incorporating leakage as well as charging and discharging efficiencies; Constraints (1c) – (1e) impose feasible operation limits on the SoC and power levels; and constraint (1f) prevents simultaneous charging and discharging.

$$\mathcal{S}(\mathcal{T}, \mathcal{C}) : \max_{p_t^c,\, p_t^d,\, s_t} \sum\nolimits_{t\in\mathcal{T}} \Delta t \cdot C_t \cdot (p_t^d - p_t^c) \tag{1a}$$

$$\text{s.t. } s_t = \rho \cdot s_{t-1} + \Delta t \cdot \left(\eta^c \cdot p_t^c - \frac{p_1^d}{\eta^d}\right), \forall t \in \mathcal{T} \tag{1b}$$

$$\underline{s} \leq s_t \leq \overline{s}, \qquad \forall t \in \mathcal{T} \tag{1c}$$

$$\underline{P}^c \leq p_t^c \leq \overline{P}^c, \qquad \forall t \in \mathcal{T} \tag{1d}$$

$$\underline{P}^d \leq p_t^d \leq \overline{P}^d, \qquad \forall t \in \mathcal{T} \tag{1e}$$

$$p_t^c \cdot p_t^d = 0, \qquad \forall t \in \mathcal{T} \tag{1f}$$

B. Reformulation as a Dynamic Programming Problem

Using the control vector u_t to denote the transitions, the problem can be formulated as a deterministic dynamic programming model with a stage reward $r_t(s_t, u_t) = \Delta t \cdot C_t \cdot (p_t^d - p_t^c)$. The *infinite-horizon* value function is $V^\infty(s) = \max_{u_t} \sum_{t=1}^{\infty} \gamma^t \cdot r_t(s_t, u_t)$ and the *finite-horizon* counterpart is $V^T(s) = \max_{u_t} \sum_{t=1}^{T-1} \gamma^t \cdot r_t(s_t, u_t)$. The forecast horizon problem seeks the smallest T^* satisfying $u_1^{T^*}(s_0) = u_1^\infty(s_0)$.

The no-simultaneous-charge-and-discharge constraint restricts the feasible action set evaluated by the Bellman operator. Since the analysis relies on contraction and boundedness properties rather than convexity, this constraint does not affect the validity of the results.

III. Analytical Framework to Determine Minimum Forecast Horizon

A. Bellman Operator and Contraction Property

The Bellman operator $\mathcal{B}$ is defined as $(\mathcal{B}V)(s) = \max_{u_t} \{r_t(s_t, u_t) + \gamma \cdot V(f(s_t, u_t))\}$, where $r_t(s_t, u_t)$ is the stage reward and $f(s_t, u_t) = s_{t+1}$ describes the next state. We first introduce the two theorems below. The detailed principles and proofs can be referred to from [10].

Theorem 0.1 (Contraction Mapping) [10]: *The Bellman operator $\mathcal{B}$ is a γ-contraction under the sup norm: $||\mathcal{B}V_1 - \mathcal{B}V_2||_\infty \leq \gamma \cdot ||V_1 - V_2||_\infty$ for all bounded value functions V_1 and V_2, and discount factor $\gamma \in (0, 1)$.*

Theorem 0.2 (Value Function Approximation) [10]: *If $|r_t(s_t, u_t)| \leq R_{\max}$, then the difference between the infinite- and finite-horizon value functions satisfies: $||V^T - V^\infty||_\infty \leq \frac{\gamma^T}{1-\gamma} \cdot R_{\max}$.*

B. Closed-Form Bound for the Minimum Forecast Horizon

We seek the minimum planning horizon T^* such that the deviation in the value function between the finite- and infinite-horizon formulations is below a given tolerance ε, i.e., $\frac{\gamma^T}{1-\gamma} \cdot R_{\max} \leq \varepsilon$. Rearranging it and taking the logarithm of both sides yields $T \geq \frac{\log(\varepsilon\cdot(1-\gamma)/R_{\max})}{\log\gamma}$. Therefore, we express the closed-form bound for the minimum forecast horizon as:

$$T^* = \left\lceil \frac{\log(\varepsilon \cdot (1-\gamma)/R_{\max})}{\log \gamma} \right\rceil. \tag{2}$$

Equation (2) provides analytical bounds for the minimum horizon length that guarantees a bounded deviation in the value function under the contraction property of the Bellman operator. With this, we can formulate our main theorem.

Theorem 1 (Minimum Forecast Horizon via Contraction Bound) Assume that $\mathcal{B}$ is a γ-contraction under the sup norm (i.e., $\|\mathcal{B}V_1 - \mathcal{B}V_2\|_\infty \leq \gamma \cdot \|V_1 - V_2\|_\infty$), $|r_t(s_t, u_t)| \leq R_{\max}$ for all (s_t, u_t), the infinite-horizon greedy policy $\pi^\infty(s) = \arg\max_{u_t}\{r_t(s_t, u_t) + \gamma \cdot V^\infty(f(s_t, u_t))\}$ is unique with respect to the initial state s_0, and a value function deviation $\|V^T - V^\infty\|_\infty \leq \epsilon$ is sufficient to preserve the optimal first time step action. Then the minimum forecast horizon T^* that ensures $\|V^T - V^\infty\|_\infty \leq \epsilon$ and $u_0^T(s_0) = u_0^\infty(s_0)$ is $T^* = \left\lceil \frac{\log(\epsilon\cdot(1-\gamma)/R_{\max})}{\log\gamma} \right\rceil$.

C. Justifications on the Key Model Assumptions

To establish analytical guarantees for the minimum forecast horizon, several structural assumptions are imposed on the ESS scheduling problem. These assumptions ensure that the Bellman operator satisfies the contraction property, the reward remains bounded, and the equivalence between finite- and infinite-horizon policies can be meaningfully defined.

979-8-3315-8138-1

Assumption 1: The problem is discounted with the discount factor $\gamma \in (0,1)$.

We introduce the discount factor $\gamma \in (0,1)$, which ensures that the Bellman operator is a γ-contraction and guarantees that $\|V^T - V^\infty\|_\infty$ is bounded above by $\frac{\gamma^T}{1-\gamma} R_{\max}$. Discounting is also physically meaningful for ESSs, since forecast uncertainty, degradation, and limited controllability naturally reduce the influence of future rewards. While we could set $\gamma = \rho \cdot \min(\eta^c, \eta^d)$ by linking discounting to efficiency losses, it can misrepresent systems with small energy capacity or large power rating where long-term rewards matter less.

To address this, we propose a more comprehensive system-aware discount factor, $\tilde{\gamma} := \rho \cdot \min(\eta^c, \eta^d) \cdot \left(\frac{E_{\max}}{E_{\max} + \max\{\overline{P}^c, \overline{P}^d\}}\right)$, which incorporates the ESS's maximum continuous full-discharging duration $E_{\max}/\max\{\overline{P}^c, \overline{P}^d\}$. This representation ensures $\tilde{\gamma} \to 0$ for ESSs of smaller energy capacity and $\tilde{\gamma} \to \gamma$ for long-duration ESSs, yielding a discount factor that more accurately reflects ESS operating characteristics and avoids excessively long forecast horizons.

Assumption 2: The stage reward is bounded by $|r_t(s_t, u_t)| \leq R_{\max}$.

This assumption requires the stage reward to be bounded to ensure that the finite-horizon value function remains well-behaved and that the truncation error satisfies $\|V^T - V^\infty\|_\infty \leq \frac{\gamma^T}{1-\gamma} R_{\max}$. Bounding the reward is essential for numerical stability, as it prevents unrealistically large profits caused by extreme prices or infeasible control actions, and guarantees that value iteration converges under the contraction property of the Bellman operator.

To obtain a practical and physically meaningful bound $R_{\max}$, we estimate it using a data-driven rolling-window procedure. Let $\tau = \frac{E_{\max} - E_{\min}}{\overline{P}^d}$ denote the time required to fully discharge the ESS. Sliding a window of length τ across the price vector $\mathcal{C}$, each segment $C^{(i)}$ yields a realizable arbitrage reward $R^{(i)} = \overline{P}^d \cdot \Delta t \cdot \sum_{j=i}^{i+\tau-1} C_j$ and the maximum attainable reward is $R_{\max} = \overline{P}^d \cdot \Delta t \cdot \max_i \left(\sum_{j=i}^{i+\tau-1} C_j\right)$. This construction identifies the most profitable τ-length interval in the historical price profile, representing the best achievable full-discharge arbitrage opportunity under the ESS's physical constraints. Because it derives directly from real price data and operation limits, it avoids artificial scaling and yields a realistic bound on achievable profits.

To extend the one-step reward to a τ-stage cumulative reward, we define $r_\tau(s,u) = \sum_{k=0}^{\tau-1} r_k(s_k, u_k) \leq R_{\max}$, which captures the maximum possible reward over the longest continuous full-discharge cycle. This cumulative formulation provides a more accurate representation of the ESS arbitrage potential and directly links reward bounds to cycle duration, energy capacity, and power constraints. By anchoring the reward limits in physically realizable operation, Assumption 2 ensures that the analytical minimum forecast horizon remains both theoretically valid and operationally meaningful.

Assumption 3: The value function deviation ε preserves optimality of the first time step decisions.

We require that a finite forecast horizon T^* yields the same first-step decision as the infinite-horizon optimal policy, i.e., the maximizing action under V^T matches that under V^∞. Since comparing policies directly is impractical, we instead bound the value-function error and assume that if $\|V^T - V^\infty\|_\infty \leq \varepsilon$, then the action rankings and thus the first-step decision remain unchanged. This assumption is justified when the reward gap between the optimal first action and its competing actions exceeds the deviation ε, guaranteeing decision consistency under finite-horizon approximation.

D. Sensitivity and Convexity Analysis

To interpret the closed-form expression for the minimum forecast horizon, we examine how T^* varies with the tolerance parameter ε. Differentiating (2) yields $\frac{dT}{d\varepsilon} = \frac{1}{\varepsilon \log \gamma}$, and since $\gamma \in (0,1)$ implies $\log \gamma < 0$, we obtain $\frac{dT}{d\varepsilon} < 0$. Thus, T^* is strictly decreasing in ε: larger tolerances lead to shorter minimum forecast horizons.

The sensitivity coefficient is $\left|\frac{dT}{d\varepsilon}\right| = \frac{1}{|\log \gamma|\, \varepsilon}$, which decreases as ε increases, indicating that T^* becomes less responsive to perturbations in ε for larger tolerances. Insignificant sensitivity occurs when $\left|\frac{dT}{d\varepsilon}\right| \ll 1$, or equivalently $\varepsilon \gg -\frac{1}{\log \gamma}$. More generally, introducing a sensitivity threshold δ leads to $\left|\frac{dT}{d\varepsilon}\right| \leq \delta \quad \Rightarrow \quad \varepsilon \geq \frac{1}{\delta\, |\log \gamma|}$, providing a quantitative condition under which changes in ε have minimal influence on the forecast horizon.

Corollary 1 (Sensitivity Range of ε for Fixed T^*) Given the closed form bound in equation (2). For a fixed discount factor $\gamma \in (0,1)$ and known $R_{\max}$, the forecast horizon T^* remains unchanged as long as $\varepsilon \in \left[\frac{R_{\max}}{1-\gamma} \cdot \gamma^{T^*}, \quad \frac{R_{\max}}{1-\gamma} \cdot \gamma^{T^*-1}\right)$.

Proof: By defining $x = \frac{\varepsilon \cdot (1-\gamma)}{R_{\max}}$, the closed-form relation in equation (2) becomes $T^* = \left\lceil \frac{\log x}{\log \gamma} \right\rceil$. By the definition of the ceiling function, we have $T^* - 1 < \frac{\log x}{\log \gamma} \leq T^*$. Multiplying through by $\log \gamma < 0$ and exponentiating gives $\gamma^{T^*} \leq x < \gamma^{T^*-1}$. Substituting back for x, we obtain $\gamma^{T^*} \leq \frac{\varepsilon \cdot (1-\gamma)}{R_{\max}} < \gamma^{T^*-1}$. It follows that $\varepsilon \in \left[\frac{R_{\max}}{1-\gamma} \cdot \gamma^{T^*}, \frac{R_{\max}}{1-\gamma} \cdot \gamma^{T^*-1}\right)$, which completes the proof. ∎

Moreover, taking the second derivative of equation (2), we get $\frac{d^2T}{d\varepsilon^2} = -\frac{1}{\varepsilon^2 \cdot \log \gamma}$. Since $\log \gamma < 0$, we have $\frac{d^2T}{d\varepsilon^2} > 0$. Hence, $T(\varepsilon)$ is *strictly convex* in ε, implying that as ε decreases, the increase in T accelerates. Conversely, as ε grows large, $\frac{d^2T}{d\varepsilon^2} \to 0$, indicating diminishing rewards from further increasing tolerance. In addition, the tolerance ε represents an acceptable deviation in realized profit relative to the infinite-horizon benchmark, reflecting a trade-off between optimality and computational efficiency.

IV. Numerical Results

We present the experimental designs and simulation results for our analytical framework; a closed-form bound for the minimum forecast horizon to compare to the benchmark brute-force simulations in our previous work [11]. We explain the test system setup, data employed, and measured control performance over different deviation ε values.

979-8-3315-8138-1

A. Simulation Setup

We take into account a single ESS over 91 days (from January 1 through March 31, 2024) at an hourly resolution, resulting in a total planning horizon of $T = 2,184$ hours. The electricity price vector $\mathcal{C} \in \mathbb{R}^{2184}$ was extracted from the ENTSO-E transparency platform day-ahead market for the bidding zone DK1 in Denmark [12]. The price data present typical market volatility with periodic cycles as well as sudden spikes. The simulation uses $\eta^c = 0.85$, $\eta^d = 0.85$, $E_{\min} = 0\,\text{MWh}$, $s_0 = 0.5$, $\rho = 0.99$, $\Delta t = 1\,\text{hour}$, and $H = 1$ hour. In the numerical simulation, the decision horizon is set to $H = 1$ hour to reflect the effect of the forecast horizon. Nevertheless, for any fixed H, the analytical bound in (2) remains valid, with changes in H absorbed into the effective reward bound and discount factor.

TABLE I: Profits for Varying Energy and Power Capacities

Energy Capacity ($\overline{P}^d = \overline{P}^c = 1$ MW)		**Power Capacity** ($E_{\max} = 10$ MWh)	
$E_{\max}$ **(MWh)**	**Profit ($)**	$\overline{P}^d, \overline{P}^c$ **(MW)**	**Profit ($)**
10	7,714.77	2	13,320.45
50	8,701.19	4	18,460.98
100	9,512.36	6	20,380.23
200	10,977.06	8	21,680.68
500	14,588.31	10	22,146.27

B. Results and Discussions

We now present numerical evaluations of the analytical bound developed for the minimum forecast horizon T^*, comparing it with the benchmark brute-force simulations outlined in our earlier work [11]. Our results span a range of energy and power capacities while holding the other dimension fixed in each case. We also analyze the behavior of both standard and system-aware discount factors across a range of deviation thresholds $\varepsilon \in \{0.01\%, 0.26\%, 0.51\%, 0.75\%, 1.00\%\}$.

1) Varying E_{max} with Fixed $\overline{P}^d = \overline{P}^c = 1$ MW: Increasing energy capacity while keeping the power limits fixed leads to higher total profits as shown in Table I, since larger storage enables the ESS to exploit longer arbitrage windows. Both the best-case cumulative reward $r_\tau(s, u)$ and the cycle duration τ grow with capacity, and this increases the system-aware discount factor $\tilde{\gamma}$. Thus, reflecting the greater value of future rewards as shown in Table II.

As shown in Fig. 1, the minimum forecast horizon T^* required to meet a target deviation ε rises with energy capacity under standard discounting, but is substantially reduced when using the system-aware discount factor. For example, at a 0.51% deviation, T^* drops from 40 to 25 hours for a 10 MWh system, whereas the reduction is smaller (77 to 76 hours) for a 500 MWh system. In all cases, the resulting profits stay within 1.20% of the desired deviation ε, demonstrating that system-aware discounting consistently yields smaller valid horizons.

2) Varying $\overline{P}^d$ and $\overline{P}^c$ with Fixed $E_{max} = 10$ MWh: Analogously, when power capacity increases while holding $E_{\max} = 10$ MWh fixed, the total profit rises substantially from $13,320.45 at 2 MW to $22,146.27 at 10 MW as shown in Table I, because higher power enables faster cycling and quicker capture of price spreads. As shown in Table II, greater power shortens the arbitrage duration τ but increases the per-step reward, lowering the system-aware discount factor and emphasizing short-term decision values.

Fig. 2 further illustrates that standard discounting yields nearly constant (and often overly conservative) minimum forecast horizons T^* across different power levels. In contrast, system-aware discounting produces much smaller and more responsive horizons. For example, at $\varepsilon = 0.51\%$, the required T^* falls from 45 hours under standard discounting to only 8 hours for a 10 MW system, with a resulting deviation of 2.22%. Although the standard discount factor can recover the exact global optimum, its required horizon is significantly longer, thus exposing decisions to greater future price volatility.

C. Theoretical Insights

Our numerical results demonstrate the practical efficacy of the proposed analytical framework in determining minimum forecast horizons for ESS scheduling. By using the τ-stage reward function $r_\tau(s, u)$, the method captures cumulative arbitrage potential over continuous full-discharge cycles, enabling a tractable approximation of the Bellman value function without fully solving the dynamic programming model.

A key insight is the important influence of ESS parameters on τ-stage rewards and the resulting discount factors. Larger energy capacities lengthen arbitrage windows and increase cumulative rewards, while higher power capacities, on the contrary, shorten the cycle and reduce the effective discount rate. The system-aware discount factor $\tilde{\gamma}$ adapts to these dynamics and produces more accurate estimates of the minimum horizon T^*, substantially reducing under- and overestimation compared to static discounting. Overall, the framework provides a sufficient and system-sensitive bound for T^*, with its advantages becoming more pronounced in the emerging smart grid with proliferated ESSs of heterogeneous sizes.

V. Conclusion

This work presents a rigorous analytical framework for determining the minimum forecast horizon T^* required for real-time ESS scheduling under deterministic electricity prices. By combining Bellman contraction principles with a data-driven τ-stage reward formulation, we derive a closed-form and physically interpretable bound that links theoretical convergence to ESS parameters such as energy capacity, power rating, efficiencies, and leakage. The introduction of a system-aware discount factor further improves accuracy by accounting for how physical constraints could shape reward accumulation, producing forecast horizons that closely align with real behavior. Numerical experiments using DK1 day-ahead price data confirm that truncated-horizon profits remain near-optimal, demonstrating the practical value of the approach for computationally efficient rolling-horizon control.

979-8-3315-8138-1

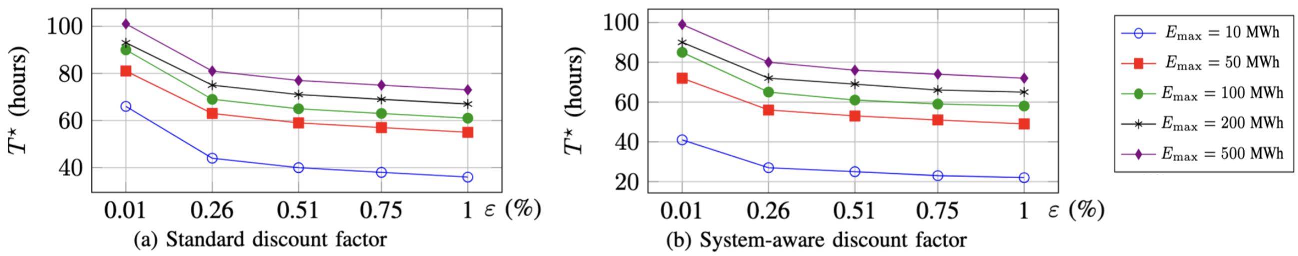

Fig. 1: Minimum forecast horizon $T^\star$ versus deviation tolerance ε for varying energy capacities ($\bar{P}^d = \bar{P}^c = 1$ MW).

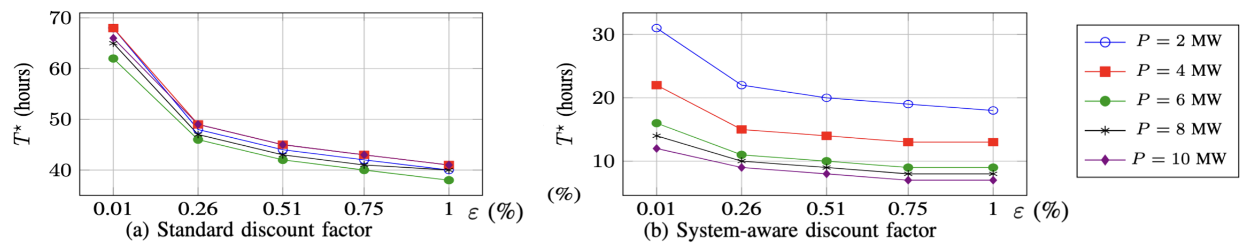

Fig. 2: Minimum forecast horizon $T^\star$ versus deviation tolerance ε for varying power capacities ($E_{\max} = 10$ MWh).

TABLE II: Rewards and System-Aware Discount Factors Against Varying Energy and Power Capacities

Energy Capacity ($\overline{P}^d = \overline{P}^c = 1$ **MW**)					**Power Capacity** ($E_{\max} = 10$ **MWh**)				
$E_{\max}$ **(MWh)**	τ **(h)**	$\mathbf{r}_\tau(\mathbf{s},\mathbf{u})$ (\$)	φ	$\tilde{\gamma}$	$\overline{P}^d$, $\overline{P}^c$ **(MW)**	τ **(h)**	$\mathbf{r}_\tau(\mathbf{s},\mathbf{u})$ (\$)	φ	$\tilde{\gamma}$
10	10	75.89	0.9091	0.7650	2	5	151.80	0.8333	0.7013
50	50	2,077.63	0.9804	0.8250	4	2	180.52	0.7143	0.6011
100	100	5,546.41	0.9900	0.8332	6	1	101.94	0.6250	0.5259
200	200	14,858.29	0.9950	0.8373	8	1	135.92	0.5556	0.4675
500	500	41,814.72	0.9980	0.8398	10	1	169.90	0.5000	0.4208

Overall, the framework establishes a unified analytical and computational basis for understanding forecast horizon requirements in ESS operation and suggests several promising directions for future research:

- Extending the framework to stochastic price environments since this work assumed deterministic prices to enable a tractable analytical bound.
- Generalizing the framework to aggregated or networked ESSs, and incorporating other factors such as degradation, and ramp-rate limits to enhance system-level realism.
- Embedding the analytical insights into an online control framework for real-time, adaptive horizon selection.

References

[1] C. Schubert *et al.*, "Hybrid energy storage systems based on redox-flow batteries: Recent developments, challenges, and future perspectives," *Batteries*, vol. 9, no. 4, 2023. [Online]. Available: https://www.mdpi.com/2313-0105/9/4/211

[2] R. N. Jesudasan and L. L. H. Andrew, "Scheduling long term energy storage," in *Proc. IEEE Conf. Comput. Commun. Workshops*, 2014, pp. 634–639. [Online]. Available: https://doi.org/10.1109/INFCOMW.2014.6849305

[3] E. Cuisinier *et al.*, "New rolling horizon optimization approaches to balance short-term and long-term decisions: An application to energy planning," *Energy*, vol. 245, 2022. [Online]. Available: https://doi.org/10.1016/j.energy.2021.122773

[4] A. Garcia and R. L. Smith, "Solving nonstationary infinite horizon dynamic optimization problems," *J. Math. Anal. Appl.*, vol. 244, no. 2, pp. 304–317, 2000. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0022247X99966945

[5] T. Cheevaprawatdomrong *et al.*, "Solution and forecast horizons for infinite-horizon nonhomogeneous markov decision processes," *Math. Oper. Res.*, vol. 32, pp. 51–72, 2007. [Online]. Available: https://doi.org/10.1287/moor.1060.0224

[6] A. Bardhan *et al.*, "Forecast and rolling horizons under demand substitution and production changeovers: Analysis and insights," *IIE Trans.*, vol. 45, 2010. [Online]. Available: https://doi.org/10.2139/ssrn.1624008

[7] B. Finnah *et al.*, "Integrated day-ahead and intraday self-schedule bidding for energy storage systems using approximate dynamic programming," *Eur. J. Oper. Res.*, vol. 301, no. 3, pp. 726–746, 2022. [Online]. Available: https://doi.org/10.1016/j.ejor.2021.11.010

[8] E. Ela and M. O'Malley, "Scheduling and pricing for expected ramp capability in real-time power markets," *IEEE Trans. Power Syst.*, vol. 31, pp. 1–11, 2015. [Online]. Available: https://doi.org/10.1109/TPWRS.2015.2461535

[9] E. Prat *et al.*, "How long is long enough? finite-horizon approximation of energy storage scheduling problems," 2024. [Online]. Available: https://arxiv.org/abs/2411.17463

[10] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, USA: John Wiley & Sons, Ltd., 1994. [Online]. Available: https://doi.org/10.1002/9780470316887

[11] N. T. Ofoe *et al.*, "On the detection of minimum forecast horizon for real-time scheduling of energy storage systems in smart grid," 2025. [Online]. Available: https://arxiv.org/abs/2511.12029

[12] ENTSO-E Transparency Platform, "Transparency platform," 2024, accessed: 2025-03-22. [Online]. Available: https://newtransparency.entsoe.eu/market/energyPrices

979-8-3315-8138-1