Rolling Horizon Unit Commitment via Polyak-Based Lagrangian Relaxation

Veronica Diaz Pacheco[1,2], Dilip Pandit[2], Cody J. Newlun[3], Atri Bera[2], Tu Nguyen[2], Jordan Kern[4]
[1]Operations Research Program, NC State University, Raleigh, NC, USA
[2]Energy Storage Technology & Systems, Sandia National Laboratories, Albuquerque, NM, USA
[3]Electric Power Systems Research Department, Sandia National Laboratories, Albuquerque, NM, USA
[4]Department of Industrial and Systems Engineering, NC State University, Raleigh, NC, USA

***Abstract*—This paper presents a rolling-horizon (RH) framework for long-horizon unit commitment with storage and companion certification method via time-decomposed Lagrangian relaxation (LR). The approach proposes formulations that allow for time decomposition: a counter-state minimum up/down-time representation and a penalty-based state-of-charge neutrality that preserves feasibility while enabling windowed solves. To obtain certification, we dualize inter-window constraints and use Polyak stepsize with dynamic level adjustment to maximize the resulting dual problem. This results in valid lower bounds by weak duality. On a large regional system and small test system, RH costs are near monolithic solutions while runtime advantages increase with horizon length, and RH–LR certification gaps are between the 2–4% range.**

***Index Terms*—Lagrangian Relaxation, Rolling Horizon, Temporal Decomposition, Unit Commitment.**

I. Introduction

The unit commitment (UC) problem optimizes generation schedules over hourly intervals to meet system demand at minimum operating cost [1]. Large-scale UC problems, especially with long planning horizons, are computationally challenging due to the large number of binary variables and the system-wide coupling they must represent. The growing integration of energy storage systems (ESSs) further adds complexity, since intertemporal state-of-charge (SoC) constraints must be accurately modeled.

UC decomposition is commonly used to improve tractability. Spatial decompositions depend on breaking the system network into smaller spatial zones. The authors in [2], [3] presented approaches for geographic decomposition of scheduling problems. Similarly, the authors in [4] propose a spatial decomposition algorithm for security-constrained UC that uses the alternating direction method of multipliers to coordinate subproblems and find the optimal solution.

Temporal decompositions partition UC planning horizon into smaller time windows. Augmented Lagrangian relaxations (LRs) have been used to coordinate time windows by updating dual variables to improve convergence [5]–[7]. The rolling horizon (RH) heuristic can handle long horizons effectively [8], [9], albeit without optimality guarantees. Additionally, certification via classical subgradient schemes can lead to fluctuations around the optimal solution and are highly sensitive to the choice of step size. Many temporal approaches fail to address ESS end-of-day energy neutrality in a way that remains consistent across decomposed subproblems; recent spatiotemporal methods emphasize representation of storage dynamics but still do not account for additional coupling from neutrality constraints [10].

The main contribution of this work is an implementable RH-LR certification workflow for ISO-scale UC facilitated by a decomposition-friendly UC formulation. Specifically, we present (i) a counter-state representation of minimum up/down time (MUT/MDT) that is particularly useful for temporal decomposition, and (ii) a penalty-based SoC neutrality treatment that is consistent across RH and LR subproblems. Using this structure, time-decomposed LR is used to compute valid lower bounds (LBs) that provide conservative, relative (tolerance-dependent) certificates for the RH schedules via primal-dual upper bound (UB)-LB gaps and we evaluate the Polyak stepsize with dynamic level adjustment (PSADLA) [11] as a reduced-tuning dual update on large UC instances.

The major contributions of this work are the following:

1) A heuristic RH algorithm with freeze/look-ahead windows to efficiently solve long-horizon UC problems while preserving feasibility.
2) Decomposition-ready intertemporal modeling via counter-state MUT/MDT representation and an SoC-neutrality penalty shared by RH and LR.
3) A time-decomposed LR construction (via interface-copy equalities) solved in parallel across windows; while LR is classical, its use to certify RH schedules is novel.
4) Evaluation of PSADLA for dual optimization as a low-tuning alternative to standard subgradient stepsize rules.

II. Methodology

We model temporal coupling via boundary states at window boundaries. RH advances through overlapping time windows

This article has been authored by employees of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan https://www.energy.gov/downloads/doe-public-access-plan. This material is based upon work supported by the U.S. Department of Energy, Office of Electricity (OE), Energy Storage Division. The authors gratefully acknowledge Bryan Arguello for the technical review.

979-8-3315-8138-1

with freeze/look-ahead, while LR uses non-overlapping windows coordinated through dual multipliers.

A. Time Partitioning and Intertemporal Constraint Modeling

We adopt the network-constrained UC model in [1] and partition the horizon, $T = \{1, 2, \ldots, |T|\}$, into k subproblems, each defined over its corresponding time window:

$$T_k = \{\tau_k, \tau_k + 1, \ldots, \tau_{k+1} - 1\} \qquad k = 1, \ldots, K \tag{1}$$

Where τ_k is the start index of window k and τ_{k+1} is the start index of the subsequent window. The single period shared between consecutive windows k and $k+1$ is denoted as the 'interface', and it is expressed as follows:

$$t_k^{\text{int}} = \tau_{k+1} - 1 \tag{2}$$

For each generator $g \in G$, and period $t \in T_k$, let $u_{g,t}, v_{g,t}, w_{g,t} \in \{0, 1\}$ denote the commitment, startup and shutdown status, and $p_{g,t} \geq 0$ the real power dispatch. Let $r^U_{g,t}, r^D_{g,t} \geq 0$ denote the remaining MUT/MDT obligations at the end of period t. All intra-window constraints (power balance, ramping within T_k, generation limits, etc.) are enforced and denoted compactly as $x_k \in X_k$ with local objective $c_k^T x_k$.

Decomposing the horizon into time windows creates inter-window coupling through ramping, storage balance and MUT/MDT. While ramping and SoC dynamics couple only adjacent periods, MUT/MDT can introduce multi-period dependence. Thus, for completeness, we augment the referenced formulation with MUT/MDT logic encoded via state counters, yielding a compact interface state that limits ω , the maximum number of consecutive time periods linked by an intertemporal constraint, to 2. In contrast, standard MUT/MDT formulations can couple decisions across up to $\max\{T_g^U,\ T_g^D\}$ periods, where T_g^U and T_g^D, denote the minimum up/down times.

1) State Representation of MUT/MDT Constraints: This counter-state representation of MUT/MDT is equivalent to classical formulations but tailored for temporal decomposition since it minimizes ω. To the best of our knowledge, an explicit counter-state representation for MUT/MDT for temporal decomposition of UC has not been documented in literature. Let $T_0 = \{\tau_k\}_{k=1}^K$ denote the set of interfaces, and O_g^U, O_g^D, denote number of time periods the unit must be on/off at the start of the planning horizon, such that:

$$\mathrm{O}_g^U = \begin{cases} \max\{0,\, T_g^U - \mathrm{St}_g\}, & \text{if } \mathrm{St}_g > 0, \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

where St_g represents the number of periods the unit has been on (> 0) or off (< 0) at the start of period 1. O_g^D is defined analogously using $St_g < 0$ and T_g^D. The following state-based MUT constraints are employed:

$$r^U_{g,t} \geq\ O_g^U + T_g^U v_{g,t} - u_{g,t}, \quad \forall g \in G,\ \forall t \in T_0,\ t = 1 \tag{4}$$
$$r^U_{g,t} \geq\ \tilde{r}^U_{g,t} + T_g^U v_{g,t} - u_{g,t}, \quad \forall g \in G,\ \forall t \in T_0,\ t \neq 1 \tag{5}$$
$$r^U_{g,t} \geq\ r^U_{g,t-1} + T_g^U v_{g,t} - u_{g,t}, \quad \forall g \in G,\ \forall t \in T_0 \tag{6}$$
$$r^U_{g,t} \leq T_g^U u_{g,t}, \quad \forall g \in G, \forall t \in T \tag{7}$$

A symmetric set of constraints is used for the remaining downtime state variable $r^D_{g,t}$, using $w_{g,t}$ and T_g^D.

2) Modeling ESS State-of-Charge Constraints: ESS introduce intertemporal coupling through SOC dynamics. For storage units $s \in S$, the SoC evolves as:

$$SoC_{s,t} = SoC_{s,t-1} + \eta_c\, C_{s,t} - \eta_d^{-1}\, D_{s,t},\ \forall s \in S,\ t \in T_k \tag{8}$$

where η_c, η_d denote the charge and discharge efficiency, and $C_{s,t}, D_{s,t}$ denotes the charging and discharging power. Additionally, ESS are subject to SoC limits, charging and discharging power limits, and the complimentary charge/discharge constraint. An energy neutrality condition is enforced at the end of each day for accurate ESS operation:

$$SoC_{s,\tau_{k+1}} \geq SoC_{s,\tau_k}, \quad \forall s \in S \tag{9}$$

Since neutrality is enforced at fixed checkpoints, it is local only when a window ends at such a checkpoint. We therefore enforce a soft neutrality penalty for time-decomposed subproblems encouraging end-of-day SoC neutrality behavior that is similar to the full-horizon model with proper window size selection. The soft neutrality is implemented via a penalized slack to avoid introducing additional inter-window coupling:

$$SF_s \geq SoC_{s,\tau_k} - SoC_{s,\tau_k+1}, \quad \forall s \in S \tag{10}$$

where SF_s is the storage shortfall below the initial level. This SoC-neutrality design allows both the RH and LR to operate on identical formulation.

B. Rolling Horizon Heuristic

We refer the reader to Algorithm 1. We use a freeze/look-ahead RH heuristic with fixed freeze F and look-ahead L lengths, and define the window length $W = F + L$. At iteration k, with current start time t_0, we solve a UC subproblem over the window $T_k = \{t_0, \ldots, t_0 + H_k - 1\}$, where $H_k = \min\{W, T - t_0 + 1\}$. After solving, we commit (freeze) decisions for the first F_k periods, where $F_k \leq \min\{F, H_k\}$ is chosen to ensure the final window begins at $t^\star = T - W + 1$ and the horizon is covered without gaps; once $t_0 \geq t^\star$, we freeze the full window ($F_k = H_k$). The boundary state (generator and storage carryover variables) at $t_0 + F_k - 1$ is then passed to the next subproblem, and we advance $t_0 \leftarrow t_0 + F_k$.

III. Certifying the Rolling Horizon Heuristic

This section presents the LR and subgradient methods used to dualize time-coupling constraints, decouple UC temporally in parallel subproblems, and certify RH solutions.

A. Parallelization via Boundary Variables

At each interface t^{int} we introduce local proxy variables in subproblem k+1 that represent the state of the system at t^{int}. We denote these copies with a tilde as follows: $\tilde{u}_{g,t^{\text{int}}}, \tilde{p}_{g,t^{\text{int}}}, \tilde{r}^U_{g,t^{\text{int}}}, \tilde{r}^D_{g,t^{\text{int}}}, \widetilde{SoC}_{s,t^{\text{int}}}$. These mirror their counterparts in window k, enforcing temporal consistency:

$$u_{g,t^{\text{int}}} = \widetilde{u}_{g,t^{\text{int}}}, \quad p_{g,t^{\text{int}}} = \widetilde{p}_{g,t^{\text{int}}},$$
$$r^U_{g,t^{\text{int}}} = \widetilde{r}^U_{g,t^{\text{int}}}, \quad r^D_{g,t^{\text{int}}} = \widetilde{r}^D_{g,t^{\text{int}}}, \quad SoC_{s,t^{\text{int}}} = \widetilde{SoC}_{s,t^{\text{int}}} \tag{11}$$

979-8-3315-8138-1

Algorithm 1 RH Heuristic used in this work.

Set the input parameters H, F, L, W, T
Initialize: $t_0 \leftarrow 1$, $t^* \leftarrow \max(1, T - W + 1)$
while $t_0 < T$ **do**
  $r \leftarrow T - t_0 + 1$
  $H \leftarrow \min\{W, r\}$
  $w_o \leftarrow (t_0 + W - 1 > T)$
  $w_s \leftarrow (t_0 < t^*) \wedge (t_0 + \min(F, H) > t^*)$
  if $t_0 < t^* \wedge (w_o \vee w_s)$ **then**
    $F_k \leftarrow \min\{F, H, t^* - t_0\}$
  else
    $F_k \leftarrow \min\{F, H\}$
  end if
  Build subproblems for $t \in [t_0, t_0 + H - 1]$
  while solving the subproblem **do**
    if $t_0 \geq t^*$ **then**
      $F_{freeze} \leftarrow H$
      Freeze decisions for $t \in [t_0, t_0 + F_{freeze} - 1]$
    else
      $F_{freeze} \leftarrow F_k$
      Freeze decisions for $t \in [t_0, t_0 + F_{freeze} - 1]$
    end if
  end while
  Compute new residuals at $t_0 + F_{freeze}$
  Advance $t_0 \leftarrow t_0 + F_{freeze}$
end while

Inside each window, the standard UC constraints are preserved. However, wherever the first period $t = \tau_{k+1}$ references $t-1$, we add auxiliary constraints using the copied variables to the local feasible set X_k. For example, the first state transition constraint in window $k+1$ becomes:

$$u_{g,\tau_{k+1}} - \widetilde{u}_{g,t_{int}} = v_{g,\tau_{k+1}} - w_{g,\tau_{k+1}} \tag{12}$$

This reformulation keeps the logical structure of the original problem provided interface equalities (11) hold while allowing independent subproblem solves. With this method, we indirectly dualize constraints that couple adjacent time windows.

B. Dual Problem

The original problem can be expressed as a minimization over all subproblems linked through the interface equalities (11). By dualizing these, we obtain the Lagrangian as follows:

Let $\mathcal{I} = \{t_k^{\text{int}} \mid k = 1, \ldots, K-1\}$, and introduce dual multipliers $\lambda_{g,t}^{u}$, $\lambda_{g,t}^{p}$, $\lambda_{g,t}^{up}$, $\lambda_{g,t}^{down}$, $\lambda_{s,t}^{SoC}$ for (11) at each interface $k = 1, \ldots, K-1$. The Lagrangian is defined as

$$L(x, \tilde{x}, \lambda) = \sum_{k=1}^{K} c_k^\top x_k + \sum_{t\in\mathcal{I}} \sum_{g\in G} \Big[\lambda_{g,t}^{u} \left(u_{g,t} - \widetilde{u}_{g,t}\right) + \lambda_{g,t}^{p} \left(p_{g,t} - \widetilde{p}_{g,t}\right) + \lambda_{g,t}^{up} \left(r_{g,t}^{U} - \widetilde{r}_{g,t}^{U}\right) + \lambda_{g,t}^{down} \left(r_{g,t}^{D} - \widetilde{r}_{g,t}^{D}\right)\Big] + \sum_{t\in\mathcal{I}} \sum_{s\in S} \lambda_{s,t}^{SoC} \left(SoC_{s,t} - \widetilde{SoC}_{g,t}\right) \tag{13}$$

Given a fixed value $\bar{\lambda}$, the Lagrangian decouples over windows because, for $t \in \mathcal{I}$, the objective and constraint terms with $u_{g,t}, p_{g,t}, r_{g,t}^U, r_{g,t}^D, SoC_{s,t}$ belong to window $k-1$ while the terms with tilde belong to window k. For a fixed $\bar{\lambda}$, the relaxed problem is expressed as follows:

$$f(\bar{\lambda}) = \sum_{k=1}^{K} \min_{x_k, \tilde{x}_k \in X_k} L_k(x_k, \tilde{x}_k, \bar{\lambda}) \tag{14}$$

By weak duality theorem, $f(\bar{\lambda})$ will provide a valid LB to the original UC cost. Since the dual is concerned with finding the best LB, its defined as:

$$v^* = \max_{\lambda} f(\lambda) \tag{15}$$

which is concave in λ but non smooth. Hence, (15) is solved using the subgradient optimization in the following section.

C. Subgradient and certificates

The subgradient method iteratively updates the multipliers λ based on mismatches at the subproblem interfaces. At iteration m, non-overlapping sequential subproblems are solved in parallel using the current multipliers λ^m and the resulting interface residuals form the subgradient:

$$g^m = \begin{bmatrix} u_{g,t^{\text{int}}} - \widetilde{u}_{g,t^{\text{int}}} \\ p_{g,t^{\text{int}}} - \widetilde{p}_{g,t^{\text{int}}} \\ r_{g,t^{\text{int}}}^U - \widetilde{r}_{g,t^{\text{int}}}^U \\ r_{g,t^{\text{int}}}^D - \widetilde{r}_{g,t^{\text{int}}}^D \\ SoC_{g,t^{\text{int}}} - \widetilde{SoC}_{g,t^{\text{int}}} \end{bmatrix} \tag{16}$$

For dual maximization, multipliers are updated as:

$$\lambda^{m+1} = \lambda^m + \theta^m g^m \tag{17}$$

where θ^m is the step size at m, which is scaled and can be tuned using parameter γ. We test the PSADLA [11] as update rule for dual optimization. PSADLA uses information from the current dual and level values to adapt the step size. The goal is to evaluate the PSADLA convergence behavior. RH yields a feasible UB z^{RH}, which we use for optimality certification and as the PSADLA level/Polyak UB parameter in the stepsize computation. To obtain certified LR lower bounds under a nonzero MIPGap, at iteration m we form:

$$\widehat{f}(\lambda^m) = \sum_{w\in\mathcal{W}} \text{LB}_w(\lambda^m), \tag{18}$$

where $\text{LB}_w(\lambda^m)$ denotes the solver-reported MIP objective lower bound for window subproblem w (e.g., Gurobi `ObjBound`) solved with multipliers λ^m. We report

$$f_{\text{best}} = \max_{j\leq m} \widehat{f}(\lambda^j), \tag{19}$$

which is a valid (possibly conservative) lower bound on the monolithic UC objective by weak duality. The resulting dual values and associated primal-dual gaps are used as certificates of solution quality for the RH UC schedules. The RH–LR decomposition/certification is system and formulation agnostic: it only requires a UC model and interface-copy equalities for

979-8-3315-8138-1

TABLE I: RH performance vs. monolithic UC. Costs reported in M\$. Cost diff. (%) $= 100 \times \frac{z_{\mathrm{RH}} - z_{\mathrm{Mono}}}{z_{\mathrm{Mono}}}$.

Horizon (h)	Regional system					RTS-GMLC system				
	RH Cost	Mono Cost	Cost diff. (%)*	RH (s)	Mono (s)	RH Cost	Mono Cost	Cost diff. (%)*	RH (s)	Mono (s)
72	1154.6	1155.3	-0.06	55.1	65.93	2.1760	2.1685	0.34	9.8	7.7
168	2800.8	2798.3	0.08	133.2	303.35	4.4075	4.4064	0.02	21.1	82.4
336	5934.2	5941.3	-0.12	222.8	903.38	9.2411	9.2322	0.09	41.2	381.6

**Note:* Both runs end at a set MIPGap, so RH may occasionally yield a slightly lower cost.

the chosen linking variables. Other UC variants are handled by adding their constraints within each window and introducing any additional boundary states (and multipliers) they create.

IV. Results

A. Experimental Setup

Experiments were conducted on a regional transmission system in North and South Carolina with 997 buses, 1221 lines, and 795 generators (392 thermal, 403 variable generation sources), along with 5,100 MWh storage capacity and 920 MW discharge capability. We also perform tests on the RTS-GMLC test system (73 buses, 120 lines, 73 thermal, 81 variable generation, and 1 ESS) [12]. Three hourly planning horizons are considered, $T \in \{72, 168, 336\}$. All solves use Gurobi 11.0 with 1% MIPGap; only LR subproblems use a non-persistent solver interface. Lagrange multipliers are initialized uniformly (across runs) using the fixed values $(35, 30, 40, 32, 25, 25)$ for (power generated, commitment, UT, DT, charge, discharge).

B. Rolling Horizon Performance

We use RH as a fast alternative to monolithic UC: the horizon is split into windows of F length with look-ahead L. The resulting RH schedule is feasible by construction and is evaluated in the monolithic problem to achieve a comparable objective value z^{RH} UB. Table I shows findings for a RH run with a fixed window of 24 hours and a look-ahead of 12 hours ($F = 24$ h, $L = 12$ h).

From Table I, the RH approach matches the monolithic very closely, with absolute cost differences below 0.35% across horizons and systems. For the regional system, runtimes improved from about 1.2x faster at 72 h to about 4x faster at 336. For RTS-GMLC, RH is a little slower at 72 hr, but becomes 4x faster at 168 hr and 9x faster at 336 hr.

Table II summarizes the RH (F, L) sensitivity. Larger F and/or L improves the accuracy of costs, possibly removing boundary effects, but with higher runtimes as a result. The chosen configuration offers a reasonable tradeoff between accuracy and runtime across all horizons. Since RH schedules are feasible by construction and evaluated in the monolithic objective function, the reported UB reflects operating cost, while the LR LB quantifies how far the implementable schedule is from optimal.

To further evaluate temporal consistency, we analyze the SoC across the planning horizon. The same operational constraints are enforced in both monolithic and RH models, with a soft SoC-neutrality penalty (\$5000 per MWh) applied at the end of each fixed window for RH. Figure 1 compares SoC for a ($F = 24$ h, $L = 12$ h) configuration, showing smooth daily cycles of charging and discharging. The results show close SoC profiles, confirming that the penalty-based neutrality approach can keep intertemporal feasibility without compromising economic performance.

TABLE II: RH sensitivity to freeze (F) and look-ahead (L) windows. $F \in \{6, 12, 24, 48\}$, $L \in \{6, 12, 24\}$.

System	Horizon (h)	Cost diff. (%)		RH Runtime (s)	
		$\min_{(F,L)}$	$\max_{(F,L)}$	$\min_{(F,L)}$	$\max_{(F,L)}$
Regional	72	$0.00_{(48,24)}$	$1.51_{(6,12)}$	$40.7_{(24,6)}$	$89.2_{(12,24)}$
Regional	168	$0.03_{(48,24)}$	$1.61_{(12,24)}$	$99.3_{(12,6)}$	$205.7_{(6,24)}$
Regional	336	$0.00_{(24,24)}$	$1.59_{(12,24)}$	$204.1_{(48,6)}$	$470.4_{(6,24)}$
RTS-GMLC	72	$0.03_{(24,6)}$	$3.36_{(48,24)}$	$6.0_{(48,12)}$	$12.9_{(48,24)}$
RTS-GMLC	168	$0.03_{(24,12)}$	$1.7_{(48,6)}$	$15.1_{(24,6)}$	$38.6_{(6,24)}$
RTS-GMLC	336	$0.09_{(24,12)}$	$1.6_{(48,6)}$	$31.5_{(24,6)}$	$86.9_{(6,24)}$

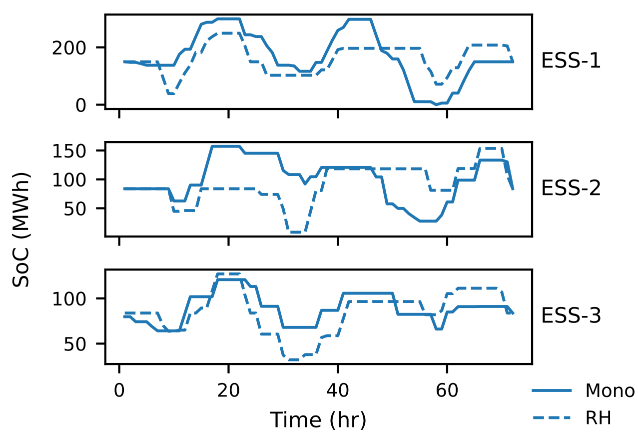

Fig. 1: SoC trajectories for the three most active storage units (ranked by SoC swing in monolithic UC solution).

C. Lagrangian Relaxation Convergence Behavior

The convergence of the PSADLA method is shown in Figure 2, which plots the RH-LR certification gap over iterations for $T = 72$ and LR window $W = 24$ for Polyak stepsizes $\gamma \in \{0.2, 0.4, 0.6\}$. The gap is calculated as $100 \times (UB - LB)/UB$ where UB is the RH cost and LB is the running-best LR dual bound. Across all settings, the running-best bound improves fast in the first 10-15 iterations and then stabilizes near a plateau, indicating diminishing returns from additional iterations. Larger γ reduces the gap more aggressively early on, while smaller γ yields slower improvement; by 30 iterations, the achieved gaps are similar across γ values.

979-8-3315-8138-1

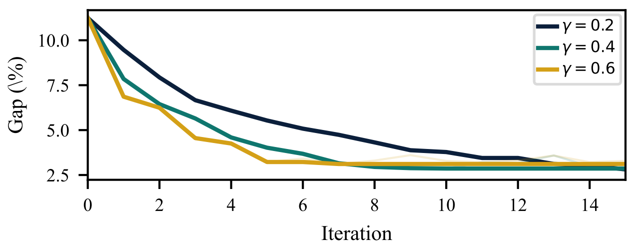

Fig. 2: LR certificate gap vs iteration for T = 72 and W = 24 under Polyak stepsizes $\gamma \in \{0.2, 0.4, 0.6\}$.

D. Bounds and Optimality Gap

Table III summarizes RH UBs, the LR LBs, and the resulting certification gaps after 30 dual iterations. For all cases shown, the RH schedules are certified within a few percent of optimality, with gaps ranging from 1.8-2.6% for the regional system and 2.2-3.6% for RTS-GMLC under $W = 24$.

TABLE III: Solution certification via RH upper bounds and LR lower bounds (1% MIPGap). Costs in M$.

System	Horizon (h)	RH UB	24 h LR window		48 h LR window	
			LR LB	Gap (%)	LR LB	Gap (%)
Regional	72	1154.6	1133.7	1.80	1139.7	1.29
Regional	168	2800.8	2734.3	2.37	2746.8	1.93
Regional	336	5934.2	5777.7	2.63	5812.9	2.04
RTS-GMLC	72	2.1760	2.1278	2.21	2.1345	1.91
RTS-GMLC	168	4.4075	4.2691	3.14	4.3206	1.97
RTS-GMLC	336	9.2411	8.9088	3.59	9.0567	1.99

The length of the LR window is a key parameter. Increasing the LR window from 24 to 48 makes the LB tighter and reduces the certification gap. For RTS-GMLC, the gap decreases from 2.21%, 3.14%, 3.60% (for increasing T) to 1.91%, 1.97%, 1.99% respectively. The larger LR window also reduces the horizon dependence: for $W = 24$, the gap grows with T. For $W = 48$, it remains approximately constant around 2%. For the regional system, the runs show the same effect.

These trends are in line with the role of windowing in time-decomposed LR: larger windows solve more intertemporal structure within each subproblem and reduce the number of interfaces being dualized which can lead to stronger bounds. Finally, because LR subproblems are solved with nonzero MIPGap, and we use solver-reported objective bounds, the reported LBs are conservative certificates; the actual optimality gap may be smaller.

Results indicate that the LR bounds are strong enough to provide informative optimality certificates: the RH schedules are, across horizons, within a few % optimal (2-4% depending on the system). Since the bounds are formed from solver-reported subproblem objective bounds under a nonzero fixed MIPGap, the certificates are conservative, suggesting the true optimality gap is likely smaller.

V. Conclusion

This paper presents a RH and LR framework for solving long-horizon UC with storage. Decomposition-friendly modeling - counter-state MUT/MDT treatment and soft neutrality - allows windowed RH solves and parallel LR subproblems coordinated through interface dual variables. On a large regional system and RTS-GMLC, RH and monolithic costs are very similar, with increasing runtime advantages at longer horizons. Shorter runtimes enable more frequent re-optimization and larger what-if studies (commonly used for planning), and the obtained schedules are directly implementable, suggesting that many long-horizon UC studies do not need to be solved monolithically to obtain high quality solutions/schedules.

The LR provides quality certificates using PSADLA. Certificates are valid by weak duality; yet since both RH and LR terminate at a nonzero MIPGap, the reported bounds may be conservative. Consequently, the reported UB-LB gap is a conservative certificate that depends on windowing choices, system coupling and finite accuracy subproblem solves (can be tightened with stricter MIPGap tolerances). A remaining limitation is the lack of direct comparison to other certification approaches. Future work will address these comparisons, evaluate stress operating conditions and investigate tighter dual algorithms.

References

[1] A. J. Conejo and L. Baringo, *Power system operations*, vol. 11. Springer, 2018.
[2] P. Li, Q. Wu, M. Yang, Z. Li, and N. D. Hatziargyriou, “Distributed distributionally robust dispatch for integrated transmission-distribution systems,” *IEEE Trans. Power Syst.*, vol. 36, no. 2, pp. 1193–1205, 2020.
[3] M. Wang, Y. Wu, M. Yang, M. Wang, and L. Jing, “Dynamic economic dispatch considering transmission–distribution coordination and automatic regulation effect,” *IEEE Trans. Ind. Appl.*, vol. 58, no. 3, pp. 3164–3174, 2022.
[4] A. K. Madhusoodhanan, M. Giuntoli, S. Schmitt, I. Harjunkoski, J. Poland, and T. Leibfried, “A spatial decomposition method for solving the security-constrained unit commitment problem,” in *2024 IEEE Int. Humanitarian Technol. Conf. (IHTC)*, pp. 1–7, IEEE, 2024.
[5] K. Kim, A. Botterud, and F. Qiu, “Temporal decomposition for improved unit commitment in power system production cost modeling,” *IEEE Trans. Power Syst.*, vol. 33, no. 5, pp. 5276–5287, 2018.
[6] F. Safdarian, A. Mohammadi, and A. Kargarian, “Temporal decomposition for security-constrained unit commitment,” *IEEE Trans. Power Syst.*, vol. 35, no. 3, pp. 1834–1845, 2019.
[7] J. Wu, W. He, Y. Sun, X. Wu, M. He, and B. Cao, “Temporal decomposition strategy for scuc with cascaded hydro-station systems,” in *2023 IEEE 7th Conf. Energy Internet Energy Syst. Integration (EI2)*, pp. 1833–1839, 2023.
[8] L. Tziovani, M. Asprou, I. Ciornei, P. Kolios, L. Hadjidemetriou, A. Lazari, R. Tapakis, and S. Timotheou, “Long-term unit commitment with combined-cycle units,” in *2023 IEEE Belgrade PowerTech*, pp. 01–06, IEEE, 2023.
[9] S. Riaz, G. Verbič, and A. C. Chapman, “Computationally efficient market simulation tool for future grid scenario analysis,” *IEEE Trans. Smart Grid*, vol. 10, no. 2, pp. 1405–1416, 2017.
[10] Q. Wang, W. Wu, C. Lin, Y. Yang, and B. Wang, “A spatio-temporal decomposition method for the coordinated economic dispatch of integrated transmission and distribution grids,” *IEEE Trans. Power Syst.*, vol. 39, no. 3, pp. 4835–4851, 2023.
[11] A. Liu, M. A. Bragin, X. Chen, and X. Guan, “Accelerating level-value adjustment for the polyak stepsize: A. liu et al.,” *J. Optim. Theory and Appl.*, vol. 206, no. 3, p. 71, 2025.
[12] C. Barrows, A. Bloom, A. Ehlen, J. Ikäheimo, J. Jorgenson, D. Krishnamurthy, J. Lau, B. McBennett, M. O’Connell, E. Preston, A. Staid, G. Stephen, and J.-P. Watson, “The ieee reliability test system: A proposed 2019 update,” *IEEE Trans. Power Syst.*, vol. 35, no. 1, pp. 119–127, 2020.

979-8-3315-8138-1