energies

MDPI

Article

Spatiotemporal Decomposition for Distributed Multi-Time-Scale Supervisory Scheduling of Cascade Pumping Stations

Jun Jia [1], Zhihan Shi [2] and Guangming Zhang [1,*]

[1] College of Electrical Engineering and Control Science, Nanjing Tech University, Nanjing 211816, China; jiajun20222026@126.com
[2] School of Electrical and Energy Engineering, Nantong Institute of Technology, Nantong 226002, China; 20256115@ntit.edu.cn
* Correspondence: zgm@njtech.edu.cn

Abstract

Economical operation of cascade pumping stations requires temporal feedback and spatial coordination, whereas a single day-ahead schedule cannot correct deviations arising during execution. This study develops a distributed multi-time-scale supervisory scheduling framework based on spatiotemporal decomposition. A convex parent model represents interval-average discharges, lumped canal-storage balances, delivery targets, tariff-based pumping cost, and station-level operating limits; within-reach hydraulic dynamics and unit-level electromechanical behavior are not resolved. Hourly day-ahead, 15 min intraday, and 5 min real-time models are linked through reference mapping, measured-state feedback, and remaining-delivery correction. The cascade is partitioned into station–canal subsystems, and the alternating direction method of multipliers coordinates local copies of each boundary flow without imposing equal adjacent-station discharges. In a deterministic five-station representative-day case, the proposed method reduced the absolute delivery error, flow-variation index, terminal canal-state deviation, and specific operating cost by 90.01%, 16.42%, 95.12%, and 3.42%, respectively, relative to decentralized local feedback without consensus exchange. Its delivered volume, specific energy, and specific operating cost differed from centralized multi-time-scale scheduling by only 0.006%, 0.006%, and 0.026%, respectively. All 313 distributed rolling problems converged, with mean intraday and real-time solution times of 6.583 s and 0.461 s. In constructed day-ahead cases containing 5–40 stations, the objective gap remained below 0.012% and the maximum interface mismatch below 1.0×10^{-4} m^3 s^{-1}. These results support the numerical applicability of the coordination method to larger serial cascades sharing the same aggregate convex structure, but do not establish dynamic-hydraulic, unit-level, or field-scale validity.

Keywords: cascade pumping stations; multi-time-scale scheduling; distributed optimization; spatiotemporal decomposition; alternating direction method of multipliers; canal storage; supervisory feedback scheduling

Academic Editor: Kwok Tong Chau

Received: 20 August 2026
Revised: 17 September 2026
Accepted: 18 September 2026
Published: 20 September 2026

1. Introduction

Large interbasin water-transfer projects are essential for alleviating regional water scarcity and supporting water security in densely populated and economically developed areas [1,2]. The Eastern Route Project, part of the South-to-North Water Diversion Project relies on successive pumping stations to convey water through a long open-channel system, making pumping electricity a major component of operating expenditure [1,3,4]. Unlike isolated pumping facilities, cascade stations are coupled through canal reaches that both convey and temporarily store water. At the scheduling scale, a discharge adjustment at one station changes the stored volume of its downstream reach and consequently affects the operating conditions of the adjacent station. Actual open-channel systems also exhibit wave propagation, transport delay, water-level variation, and hydraulic transients [5]. The present study therefore adopts an aggregate balance model rather than a dynamic hydraulic model. Interval-average discharges and lumped canal storage retain the water-volume coupling required at the supervisory scheduling layer, while excluding wave propagation, water-level transients, and transport delay. The model is evaluated only for this limited purpose through feasibility, mass-balance, solution-consistency, disturbance, and parameter-sensitivity tests. Within this scope, the scheduling problem is to coordinate tariff-based pumping, water-delivery completion, canal-state regulation, and interstation flow adjustment.

The scheduling problem is further complicated by the different rates at which operational information becomes available. Electricity tariffs and daily delivery requirements can be incorporated into a day-ahead schedule, whereas executed discharges and aggregate canal states are revealed progressively during operation. A schedule optimized only once before operation cannot compensate for deviations accumulated afterward, while decisions based exclusively on short-term observations may lose the economic and delivery guidance supplied by the daily plan. Multi-time-scale scheduling and water–energy coordination studies have shown that updated forecasts and observed states can be used to revise operating decisions as the execution time approaches [6–8]. Adaptive model predictive control and hierarchical multi-timescale scheduling provide further evidence that intermediate and fast corrective layers can respond to changing system states without abandoning longer-horizon objectives [9,10]. In the present study, feedback operates at the supervisory scheduling level by updating discharge references and remaining-delivery targets. It does not represent continuous pressure, motor-speed, electromagnetic-torque, or converter control.

System scale introduces a second challenge. A centralized formulation must simultaneously process the operating variables of all pumping stations, canal states, temporal coupling constraints, and water-delivery requirements. Its computational and communication burden increases with both the number of subsystems and the scheduling resolution. Nevertheless, most constraints are physically local: a pumping station is governed by its own capacity and discharge limits, while interactions with neighboring stations are transmitted primarily through the boundary flow and storage balance of the connecting canal. Coordinated water–power studies have demonstrated the value of exploiting such subsystem structure while preserving cross-network interactions [11,12]. More broadly, distributed optimization has been adopted in networked infrastructure to improve scalability and retain local decision autonomy, although communication requirements, subsystem partitioning, and convergence behavior remain important implementation considerations [13,14]. Cascade pumping-station scheduling consequently involves two coupled forms of coordination: decisions must be refined over time as aggregate operating information becomes available, and computation must be distributed across space without violating the canal-storage balance and operating constraints of the centralized parent model.

Pump-scheduling research has progressively moved from a heuristic search toward optimization models that represent hydraulic feasibility, electricity tariffs, and equipment operating limits more explicitly. Binary-expansion formulations and polyhedral relaxations have improved the tractability of pump-scheduling problems in large water-supply networks [15,16], while chance-constrained and contingency-aware models have extended the scheduling scope to uncertain demands, renewable generation, and water–power interdependence [17–19]. These developments provide important tools for reducing pumping costs and maintaining network feasibility. Recent work has also decomposed day-ahead pump and valve scheduling into coordinated smaller subproblems, indicating the computational value of distributed solution architectures for large water networks [20]. However, most existing formulations are designed for pressurized distribution networks, tank-assisted water systems, or integrated water–power infrastructures in which pump switching and nodal hydraulic conditions constitute the main decision structure. Long-distance open-channel cascades exhibit a different form of coupling: each canal reach acts as distributed storage, and the discharge difference between adjacent pumping stations is absorbed by the changing canal state. Consequently, equipment-level scheduling or network-level energy coordination alone does not establish the state-dependent inter-station relationships required for cascade pumping-station operation.

A second line of research concerns the temporal organization of scheduling and control. Multi-time-scale scheduling reconciles long-horizon economic objectives with short-term corrective requirements by establishing a coarse-resolution baseline and progressively updating the decisions as forecasts and operating states become available. Representative applications include two-time-scale robust dispatch for active distribution networks, coordinated scheduling and fast control of flexible loads, and the aggregation, disaggregation, and cross-stage coordination of controllable resources between day-ahead and real-time operation [21–24]. These studies indicate that temporal decomposition can reduce the effect of forecast errors while preserving the economic guidance of the preceding scheduling layer. A similar multi-time-scale feedback structure has recently been introduced for cascade pumping stations, where the daily plan is progressively corrected according to updated water-demand information and observed canal states [6]. Nevertheless, existing multi-time-scale formulations predominantly emphasize uncertainty accommodation within a centralized decision structure. For cascade water-transfer systems, passing a schedule from one temporal layer to the next is insufficient unless the inherited discharge references remain compatible with the aggregate canal-storage balance, cumulative delivery requirement, and interstation flow coupling.

Spatial decomposition addresses a complementary scalability issue by distributing computation among locally operated subsystems. The alternating direction method of multipliers (ADMM) is particularly suitable when an optimization problem becomes separable after shared variables are duplicated and coordinated through consensus constraints [25]. In coupled water–power networks, ADMM-based solvers allow individual operators to retain their local models while exchanging only the variables associated with cross-system interactions [12]. Related developments in open-channel systems have used ADMM-based distributed model predictive control to divide a large canal network into coordinated local control problems [26]. Distributed predictive control has also been extended to interconnected subsystems operating at different update rates, showing that multi-rate decisions can be coordinated without assembling all subsystem models in a single controller [27]. Recent developments further indicate that time-coupled water-network models, nonconvex local subproblems, and asynchronous information exchange require explicit treatment when designing an implementable ADMM-based coordination scheme [28–30]. These studies establish a sound basis for distributed scheduling, but the decomposition must remain consistent with the underlying physical topology. In a cascade pumping system, consensus should be imposed between duplicated representations of the same boundary flow rather than directly between the discharges of two adjacent pumping stations; the latter may differ because their imbalance is retained as a change in canal storage. Preserving this distinction is essential for obtaining a distributed formulation that remains consistent with the centralized aggregate balance model.

Existing studies therefore address pump-operation optimization, temporal refinement, and spatial coordination largely as separate problems. Pump scheduling is commonly single-scale or centralized [15–20]; multi-time-scale methods generally retain centralized computation [6,9,10,21–24]; and distributed methods do not necessarily preserve the temporal organization and physical boundary definitions of cascade systems [11–14,25–30]. The remaining challenge is to combine temporal refinement and spatial decomposition while preserving the aggregate canal-storage balance and maintaining schedule feasibility during inter-layer transfer. Moreover, the benefit of coordinated feedback must be distinguished from the more elementary benefit produced by any local feedback controller. Addressing these issues requires a common parent model, explicit state and reference transfer across temporal layers, physically consistent boundary-flow consensus, and comparisons against both open-loop and decentralized-feedback configurations.

To address these requirements, this study develops a spatiotemporal decomposition-based supervisory scheduling framework derived from a centralized convex parent model containing tariff-based electricity cost, flow regulation, canal-storage continuity, water delivery, and station-level operating constraints. Temporally, the hourly day-ahead (DA), 15 min intraday (ID), and 5 min real-time (RT) scheduling layers are connected through reference mapping and feedback of aggregate states, executed discharges, and cumulative delivery deviations. Spatially, each nonterminal station and its downstream canal reach form a subsystem, while the terminal subsystem retains the delivery task. ADMM coordinates duplicated representations of the same boundary flow without imposing equal discharges at adjacent pumping stations. The formulation represents each pumping station as an aggregate continuous-flow subsystem. It does not identify a specific pump type or resolve individual motors, converters, unit commitment, torque–speed characteristics, or detailed hydraulic-network dynamics. Equivalent head and efficiency are treated as fixed coefficients in the numerical case; consequently, the optimization reallocates pumping with respect to the time-of-use tariff rather than optimizing energy consumption as an independent objective.

The main contributions of this study are summarized as follows:

1. A multi-time-scale supervisory feedback structure links DA planning, ID rolling correction, and RT execution adjustment through explicit reference mapping, aggregate-state feedback, and remaining-delivery reconstruction. The DA-only, DA+ID, and DA+ID+RT configurations are evaluated separately to quantify the incremental contribution of each temporal layer;
2. A spatial decomposition is derived directly from the centralized convex parent model. Each subsystem retains its station constraints and downstream canal balance, while ADMM coordinates local copies of the same physical boundary flow rather than imposing an artificial equality between adjacent pumping-station discharges;
3. The proposed method is evaluated against an open-loop DA reference, a decentralized local-feedback controller without consensus exchange, and a centralized multi-time-scale implementation under matched conditions. Its numerical adequacy and applicability are examined using delivered-volume-normalized indicators, feasibility and mass-balance checks, temporal and storage ablations, execution-disturbance stress testing, head–efficiency sensitivity analysis, and constructed DA scalability cases containing 5, 10, 20, and 40 pumping stations.

The remainder of this paper is organized as follows. Section 2 formulates the aggregate centralized scheduling model and defines its physical and operational scope. Section 3 develops the DA, ID, and RT scheduling layers and their inter-layer reference-, state-, and delivery-transfer mechanisms. Section 4 presents the spatial decomposition, boundary-flow consensus formulation, ADMM updates, and nested spatiotemporal solution procedure.

Section 5 presents the numerical case and evaluates parameter selection, supervisory feedback performance, decentralized and centralized benchmarks, temporal and storage ablations, execution-error sensitivity, head–efficiency sensitivity, coordination accuracy, convergence, computational cost, spatial scalability, and the reliability–cost trade-off. Section 6 summarizes the findings, limitations, and directions toward a complete electromechanical–hydraulic closed-loop system.

2. System Description and Centralized Supervisory Scheduling Formulation

Before temporal and spatial decomposition, a common station-level supervisory scheduling model is established. The cascade is represented by aggregate pumping stations connected through canal reaches that convey water and provide short-term storage. The model retains the interval water-volume balance and operating constraints shared by the centralized benchmark, the three temporal layers, and the distributed formulation. This representation is intended to coordinate station-level discharge setpoints over scheduling intervals rather than reproduce within-reach hydraulic transients. Accordingly, wave propagation, transport delay, backwater effects, and individual pump-unit electromechanical dynamics are not resolved and must be examined by a calibrated hydraulic or lower-level control model before field implementation.

2.1. Cascade Pumping-Station System and Scheduling Scope

The Eastern Route Project conveys water through long open channels and successive pumping stations where gravity flow is insufficient. This study considers a representative main-line reach comprising an upstream source, cascade pumping stations, intermediate canal reaches, and a receiving end. Branches, gates, offtakes, and confluences are excluded, yielding the alternating station–canal structure shown in Figure 1.

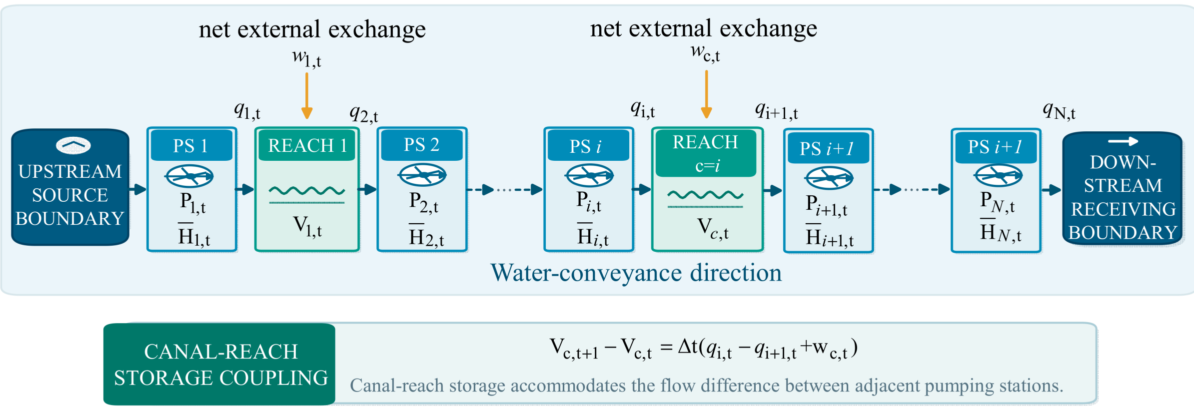

Figure 1. Cascade pumping-station–canal system and principal variables.

The formulation does not distinguish among centrifugal, axial-flow, vortex, or other specific pump types. Instead, each pumping station is represented as an aggregate continuous-flow facility characterized by station-level discharge limits, an aggregate input-power limit, an equivalent head, and a station-level equivalent electrical-to-hydraulic efficiency. The numerical case in Section 5 instantiates five such stations with aggregate input-power limits of 17–20 MW per station. These values define the scheduling scale and should not be interpreted as the ratings of individual pump motors. Accordingly, the intended application is the supervisory scheduling of cascade pumping stations rather than unit-level motor or converter control.

Consider a cascade system containing N pumping stations and $N-1$ intermediate canal reaches. The corresponding index sets are defined as $\mathcal{I} = \{1, 2, \ldots, N\}$ and $\mathcal{C} = \{1, 2, \ldots, N-1\}$. Pumping stations are numbered in the direction of water conveyance, and canal reach $c = i$ connects pumping stations i and $i+1$. The upstream source defines the inflow boundary, whereas the discharge from pumping station N defines the delivery boundary at the receiving end. Each pumping station lifts water into its downstream reach, while each reach conveys water and provides short-term storage.

As shown in Figure 1, adjacent stations are coupled through the storage state of the intervening canal. Their interval discharges may differ: excess upstream discharge stores water, whereas excess downstream discharge releases it. Canal storage therefore transfers flow imbalances across space and time.

Let $q_{i,t}$ denote the mean discharge of pumping station i during interval t, $V_{c,t}$ denote the equivalent stored volume of canal reaching c at the beginning of that interval, and Δt denote the scheduling time step. The water-Balance of the reach connecting pumping stations i and $i+1$ is as follows:

$$V_{c,t+1} = V_{c,t} + \Delta t(q_{i,t} - q_{i+1,t} + w_{c,t}),\ c = i,\ \mathrm{i} \in \mathcal{I} \backslash \{N\} \quad (1)$$

Here, $w_{c,t}$ is the net external exchange rate, positive for lateral inflow and negative for offtakes or losses. Equation (1) allows the discharges of adjacent stations to differ through a corresponding change in canal storage. In the numerical case, $w_{c,t} = 0$; variability in lateral withdrawals is therefore not included in the reported robustness tests. If both external exchange and storage variation vanish, Equation (1) reduces to $q_{i,t} = q_{i+1,t}$.

A discrete quasi-steady, lumped-storage balance is adopted because the decision variables of the present framework are interval-average station discharges and aggregate canal-storage states. This representation preserves the intertemporal water-volume coupling required for supervisory scheduling while maintaining a convex structure that can be repeatedly solved at the DA, ID, and RT layers. A dynamic hydraulic model based, for example, on the Saint-Venant equations would be required to reproduce water-level propagation, transport delay, backwater effects, and transient hydraulic feasibility. Such processes are not represented by Equation (1).

Model adequacy is therefore assessed only with respect to the intended scheduling function. Water-volume continuity, canal-storage bounds, station operating limits, solver residuals, and interface consistency are independently checked in the numerical experiments, while Section 5.7 evaluates the sensitivity of the reported energy and cost indicators to the aggregate head and efficiency coefficients. No calibration against measured water-level trajectories or validation against a dynamic hydraulic simulator was performed. The model should consequently be interpreted as adequate for evaluating the proposed supervisory scheduling and coordination mechanism under the stated aggregate assumptions, rather than as a validated predictor of detailed hydraulic behavior.

2.2. Aggregate Water-Balance, Operational, and Water-Delivery Constraints

The scheduling horizon is discretized into $\mathcal{T} = \{1, \ldots, T\}$, with $\mathcal{T}^+ = \{1, \ldots, T+1\}$ denoting the associated state indices. The duration of each interval is Δt in seconds, and $\Delta h = \Delta t / 3600$ is its equivalent duration in hours. The following constraints define the common physical feasible region used by the centralized and distributed scheduling models.

2.2.1. Canal Water-Balance and Storage Constraints

At the scheduling resolution, each canal reach is represented by an aggregate stored-volume state. Equation (1) updates this state from interval-average boundary flows and should not be interpreted as a dynamic open-channel flow model. In addition to this balance, the stored volume of each canal reach must remain within its admissible operating range:

$$V_c^{\min} \le V_{c,t} \le V_c^{\max}, \quad c \in \mathcal{C}, t = 1, 2, \ldots, T+1 \tag{2}$$

The bounds $V_c^{\min}$ and $V_c^{\max}$ define the allowable operating range, including the reserved safety margin.

For a full-horizon scheduling problem, the initial storage state is prescribed and the terminal state is required to recover to a specified target:

$$V_{c,1} = V_c^{\mathrm{ini}}, \quad V_{c,T+1} = V_c^{\mathrm{tar}}, \quad c \in \mathcal{C} \tag{3}$$

Equation (3) prevents economic gains from persistent storage depletion. It is enforced in the centralized and DA models; the ID and RT layers use updated initial states and horizon-specific recovery treatments defined in Section 3.

2.2.2. Pumping-Station Operating Constraints

The conveyance discharge of each pumping station is bounded by its admissible operating range:

$$q_i^{\min} \le q_{i,t} \le q_i^{\max}, \quad i \in \mathcal{I}, t = 1, 2, \ldots, T \tag{4}$$

Here, $q_i^{\min}$ and $q_i^{\max}$ denote the admissible station-level discharge limits under the prescribed operating conditions. A scheduled station discharge and equivalent head may in practice be realized by different combinations of operating units, rotational speeds, or blade settings. The present aggregate formulation does not select among these alternatives: $q_{i,t} = 0$, represents shutdown at the station level, whereas individual unit commitment, minimum up/down times, switching counts, and adjustable unit parameters are not decision variables. The resulting discharge setpoint must therefore be translated into feasible unit-level commands by a lower-level unit-dispatch or control system.

The equivalent pumping head is evaluated from the difference between the upstream and downstream reference water levels and the associated head loss:

$$\overline{H}_{i,t} = \overline{Z}_{i,t}^{d} - \overline{Z}_{i,t}^{u} + k_i \left(\overline{q}_{i,t}\right)^2, \; i \in \mathcal{I}, \; t = 1, 2, \ldots, T \tag{5}$$

Here, $\overline{Z}_{i,t}^{u}$ and $\overline{Z}_{i,t}^{d}$ are the reference upstream and downstream water levels, respectively; k_i is a lumped hydraulic-resistance coefficient; and $\overline{q}_{i,t}$ is the reference discharge used to evaluate the quadratic loss term before optimization. Thus, $\overline{H}_{i,t}$ combines the reference water-level difference with an aggregate representation of hydraulic resistance. It is evaluated from the reference conditions supplied to each temporal optimization and then held fixed throughout that optimization, including all ADMM iterations. The decision flow therefore does not endogenously update the head within a solve, and changes in the detailed hydraulic network are not separately resolved; operating points outside the prescribed head range are excluded.

Consequently, the resulting schedule does not certify transient hydraulic feasibility or reproduce pressure and water-level propagation within a canal reach. Before practical deployment, the scheduled discharge trajectories should be evaluated using a calibrated dynamic hydraulic model or measured system response.

Abrupt discharge changes between consecutive intervals are limited by the station ramping capability:

$$-\Gamma_i^{\mathrm{down}} \Delta t \leqq_{i,t} - q_{i,t-1} \le \Gamma_i^{\mathrm{up}} \Delta t, \;\; i \in \mathcal{I}, \quad t = 1, \ldots, T \tag{6}$$

where Γ_i^{up} and Γ_i^{down} denote the maximum rates of discharge increase and decrease, respectively. The initial value $q_{i,0}$ is taken from the latest executed operating condition. Expressing the constraint in terms of the product of ramp rate and interval duration preserves its physical meaning across the different temporal resolutions.

2.2.3. Power and Electricity-Cost Accounting and Water-Delivery Constraints

Given the precomputed equivalent head and station-level equivalent efficiency, input power is affine in the scheduled discharge. Energy use and electricity cost are consequently derived accounting quantities, while the time-of-use tariff affects the temporal allocation of discharge. The model does not independently optimize hydraulic energy demand or operating efficiency. The power and cost relationships, together with the aggregate power-capacity constraint, are written as follows:

$$\begin{aligned} &P_{i,t} = \frac{\rho g \overline{H}_{i,t}}{10^6 \eta_i} q_{i,t} = \kappa_{i,t} q_{i,t} \\ &E_{i,t} = P_{i,t} \Delta h \;\; C_{i,t} = \pi_t E_{i,t} \\ &0 \leq P_{i,t} \leq P_i^{\max}, \; i \in \mathcal{I}, \quad t \in \mathcal{T} \end{aligned} \tag{7}$$

Here, ρ is the water density, g is the gravitational acceleration, and $\kappa_{i,t}$ is the discharge-to-power coefficient. The coefficient η_i denotes a station-level equivalent electrical-to-hydraulic efficiency that lumps the conversion losses associated with the pump, motor, transmission, and drive rather than identifying these components separately. It is treated as a station-specific constant in the present case study. Moreover, $P_i^{\max}$ is the aggregate input-power limit of pumping station i, not the rated power of an individual motor.

Accordingly, motor efficiency maps, torque–speed characteristics, and load-torque variations caused by changes in head or hydraulic resistance are not separately represented. The quantities $P_{i,t}$, $E_{i,t}$, and $C_{i,t}$ are expressed in MW, MWh, and CNY, respectively, and π_t is converted from CNY/kWh to CNY/MWh. Because $\overline{H}_{i,t}$ and η_i are fixed during each optimization, $E_{i,t}$ is determined by the discharge schedule and is not an independent optimization variable or objective. The influence of uncertainty in these equivalent parameters is evaluated through the sensitivity analysis in Section 5.7.

Accordingly, the reported tariff-based cost does not include the additional energy-saving potential or switching cost associated with optimizing individual unit combinations and adjustable unit parameters.

The scheduled discharge from the terminal pumping station must satisfy the prescribed delivery requirement over the complete horizon:

$$\sum_{t=1}^{T} q_{N,t} \Delta t = V^{\mathrm{req}} \tag{8}$$

Here, V^{req} is the receiving-end target water volume. Equation (8) applies to the centralized and DA models; Section 3 reformulates it for rolling horizons using the conveyed volume and remaining executable time.

Equations (1)–(8) define the aggregate supervisory-scheduling feasible region shared by all temporal and spatial formulations.

2.3. Centralized Benchmark Scheduling Model

The centralized model coordinates all stations and canal reaches over the full scheduling horizon. Its economic term minimizes electricity expenditure under time-of-use tariffs using the head and efficiency coefficients fixed within each solve; it does not minimize energy use as a separate objective. The remaining terms regulate consecutive discharge changes and deviations from the reference schedule.

2.3.1. Scheduling Objectives

The normalized electricity-cost, discharge-smoothness, and reference-tracking objectives are defined as follows:

$$\begin{aligned}
\hat{J}_\mathrm{e} &= \frac{1}{C_0}\sum_{i\in\mathcal{I}}\sum_{t\in\mathcal{T}}\pi_t P_{i,t}\Delta h,\ C_0=\sum_{i\in\mathcal{I}}\sum_{t\in\mathcal{T}}\pi_t P_i^{\max}\Delta h \\
\hat{J}_\mathrm{s} &= \frac{1}{N(T-1)}\sum_{i\in\mathcal{I}}\sum_{t=2}^{T}\left(\frac{q_{i,t}-q_{i,t-1}}{q_i^{\max}-q_i^{\min}}\right) \\
\hat{J}_\mathrm{r} &= \frac{1}{NT}\sum_{i\in\mathcal{I}}\sum_{t=1}^{T}\left(\frac{q_{i,t}-q_{i,t}^{\mathrm{ref}}}{q_i^{\max}-q_i^{\min}}\right)^2
\end{aligned} \tag{9}$$

The cost scale C_0 represents operation at maximum allowable power under the prevailing tariff. The smoothness and tracking terms penalize consecutive discharge changes and deviation from $q_{i,t}^{\mathrm{ref}}$, respectively. The DA layer sets $\omega_r^{\mathrm{DA}}=0$; finer layers inherit $q_{i,t}^{\mathrm{ref}}$ from the preceding schedule.

The three criteria are combined as follows:

$$J_\mathrm{C}=\omega_\mathrm{e}\hat{J}_\mathrm{e}+\omega_\mathrm{s}\hat{J}_\mathrm{s}+\omega_\mathrm{r}\hat{J}_\mathrm{r},\ \omega_\mathrm{e}+\omega_\mathrm{s}+\omega_\mathrm{r}=1,\ \omega_\mathrm{e},\omega_\mathrm{s},\omega_\mathrm{r}\geq 0 \tag{10}$$

The weights are determined through the parameter analysis in Section 5.2.

2.3.2. Centralized Optimization Problem

Let $\mathbf{x}$ collect the station discharges $q_{i,t}$, aggregate input powers $P_{i,t}$, and canal-storage states $V_{c,t}$. Combining Equations (1)–(10) yields a convex quadratic program: the fixed head and efficiency coefficients make power affine in discharge, the physical constraints are linear, and the smoothness and tracking terms are positive semidefinite. Its solution provides the centralized benchmark and the aggregate supervisory parent model for the DA, ID, and RT formulations, which modify the scheduling horizon, initial state, reference trajectory, and remaining delivery requirement without changing the underlying balance and operating constraints.

3. Hierarchical Multi-Time-Scale Coordinated Scheduling

A day-ahead schedule cannot anticipate all forecast updates, canal-storage deviations, and execution errors arising during conveyance. The scheduling horizon is therefore organized into three temporal layers: DA, ID, and RT. All three layers use the aggregate station–canal model defined in Section 2 but differ in horizon length, temporal resolution, reference trajectory, and available feedback information. Reducing the scheduling interval from 1 h to 15 min and subsequently to 5 min increases the frequency of supervisory decision updates but does not increase the hydraulic fidelity of the underlying canal model. In particular, the RT layer does not resolve wave propagation, transport delay, pressure transients, or water-level dynamics. Feedback in this study refers to periodic schedule re-optimization using aggregate canal-storage states, executed discharges, and cumulative delivery deviations, rather than to a pressure- or level-based electromechanical–hydraulic control loop.

The DA layer establishes an hourly baseline over 24 h. The ID layer performs hourly receding-horizon corrections over a 4 h window discretized at 15 min, whereas the RT layer updates the current discharge command every 5 min. As shown in Figure 2, discharge and canal-state references are transferred from coarse to fine temporal resolutions, while executed discharges and updated storage states are returned to subsequent scheduling updates. Superscripts DA, ID, and RT identify the three temporal scheduling layers, whereas the superscripts $*$, ref, act, and obs denote optimized, reference, executed, and observed quantities, respectively. A vertical bar followed by k identifies a variable defined within the k-th ID rolling window, while m denotes the chronological RT interval.

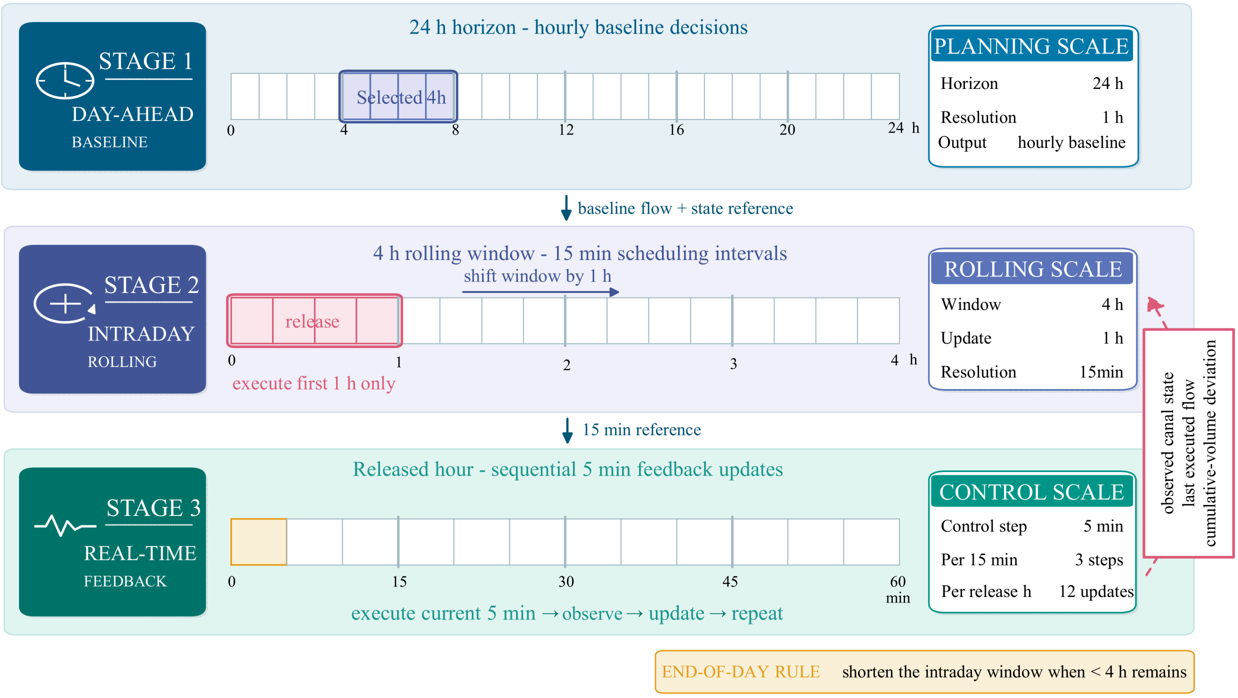

Figure 2. Multi-time-scale nesting and rolling execution mechanism.

3.1. Day-Ahead Baseline Scheduling

The DA layer determines hourly discharge and canal-storage trajectories over the 24 h operating horizon, with the corresponding input-power profile calculated through Equation (7). The interval and state-index sets are $\mathcal{T}^{\text{DA}} = \{1, 2, \ldots, 24\}$ and $\mathcal{T}^{\text{DA}+} = \{1, 2, \ldots, 25\}$, respectively, with $\Delta t^{\text{DA}} = 3600$ s. The daily delivery requirement, time-of-use tariff, initial canal states, and aggregate station parameters are treated as known. The equivalent head and station-level efficiency coefficients remain fixed within this optimization; consequently, the economic term redistributes conveyance among tariff periods rather than optimizing energy-conversion efficiency.

Because no higher-layer reference exists, the DA objective is obtained from Equation (10) by setting $\omega_r^{\text{DA}} = 0$. It therefore contains only the normalized electricity-cost and discharge-smoothness terms defined in Equation (9), with $\omega_{\text{e}}^{\text{DA}} + \omega_{\text{s}}^{\text{DA}} = 1$.

The DA schedule must also satisfy the full-day delivery target and prescribed terminal canal states:

$$\begin{aligned} \sum_{t \in \mathcal{T}^{\text{DA}}} q_{N,t}^{\text{DA}} \Delta t^{\text{DA}} &= V^{\text{req}} \\ V_{c,1}^{\text{DA}} = V_c^{\text{ini}}, \; V_{c,25}^{\text{DA}} &= V_c^{\text{tar}}, \quad \forall c \in \mathcal{C} \end{aligned} \tag{11}$$

Let $\mathcal{X}^{\text{DA}}$ denote the feasible set formed by the hourly constraints in Equations (1), (2) and (4)–(7), together with the boundary conditions in Equation (11). The DA problem minimizes J_{DA} over $\mathcal{X}^{\text{DA}}$. Its solution provides the baseline discharge and canal-storage trajectories and the corresponding derived power profile for the ID layer. Subsequent layers use this schedule as a reference and may redistribute the remaining discharge in response to observed state and delivery deviations without changing the aggregate water-balance constraints or the within-solve head and efficiency representation. The resulting discharge and canal-state trajectories remain supervisory scheduling references and require hydraulic-feasibility and unit-level verification before field execution.

3.2. Intraday Receding-Horizon Correction

At each hourly update, the ID layer reconstructs a 4 h prediction problem using the latest aggregate canal-storage states, executed discharges, time-of-use tariffs, station-level operating limits, and equivalent hydraulic parameters. The horizon is discretized into 15 min intervals, but only the first-hour decisions are released before the window shifts forward. The update therefore corrects the supervisory scheduling trajectory using newly available aggregate information. The observed canal state is treated as an input to the rolling optimization; the method does not reconstruct within-reach hydraulic propagation or pump-unit electromechanical dynamics from that observation.

Let $k \in \{0, \ldots, 23\}$ denote the hourly update. To accommodate the end of the operating day, define the remaining horizon and local index sets as follows:

$$\begin{aligned} &L_k = \min\{4, 24-k\},\ M_k = 4L_k \\ &\mathcal{T}_k^{\mathrm{ID}} = \{0, \ldots, M_k - 1\},\ \mathcal{T}_{k,+}^{\mathrm{ID}} = \{0, \ldots, M_k\},\ \Delta t^{\mathrm{ID}} = 900\ \mathrm{s} \end{aligned} \tag{12}$$

The hourly DA schedule is expanded into a piecewise-constant 15 min discharge reference, while the terminal state of each ID window is inherited from the corresponding DA state:

$$\begin{aligned} &\overline{q}_{i,j|k}^{\mathrm{DA}\to\mathrm{ID}} = q_{i,k+\lfloor j/4 \rfloor+1}^{\mathrm{DA},*}\ ,\ i \in \mathcal{I},\ j \in \mathcal{T}_k^{\mathrm{ID}} \\ &\overline{V}_{c,M_k|k}^{\mathrm{DA}\to\mathrm{ID}} = V_{c,k+L_k+1}^{\mathrm{DA},*}\ ,\ \ c \in \mathcal{C} \end{aligned} \tag{13}$$

These quantities guide the ID correction without imposing hard tracking equalities, allowing discharge to be redistributed when updated states or delivery deviations require adjustment.

Each rolling problem is initialized from the observed canal state and most recently executed discharge:

$$V_{c,0|k}^{\mathrm{ID}} = V_c^{\mathrm{obs}}(k),\ \ q_{i,-1|k}^{\mathrm{ID}} = q_i^{\mathrm{act}}(k^-),\ \ \forall i \in \mathcal{I}, c \in \mathcal{C} \tag{14}$$

The first relation provides the initial condition for the canal water balance, whereas the second maintains the discharge-ramping constraint across consecutive optimization windows. Predicted states from the preceding window are not reused as initial conditions once updated observations are available.

In addition to electricity cost, discharge smoothness, and DA-reference tracking, the ID objective penalizes the normalized terminal-state deviation:

$$\hat{J}_{\mathrm{v},k}^{\mathrm{ID}} = \frac{1}{|\ \mathcal{C}\ |} \sum_{c \in \mathcal{C}} \left(\frac{V_{c,M_k|k}^{\mathrm{ID}} - \overline{V}_{c,M_k|k}^{\mathrm{DA}\to\mathrm{ID}}}{V_c^{\max} - V_c^{\min}} \right)^2 \tag{15}$$

The objective for window k is then as follows:

$$\begin{aligned} J_{\mathrm{ID},k} &= \omega_{\mathrm{e}}^{\mathrm{ID}} \hat{J}_{\mathrm{e},k}^{\mathrm{ID}} + \omega_{\mathrm{s}}^{\mathrm{ID}} \hat{J}_{\mathrm{s},k}^{\mathrm{ID}} + \omega_{\mathrm{r}}^{\mathrm{ID}} \hat{J}_{\mathrm{r},k}^{\mathrm{ID}} + \omega_{\mathrm{v}}^{\mathrm{ID}} \hat{J}_{\mathrm{v},k}^{\mathrm{ID}} \\ &\omega_{\mathrm{e}}^{\mathrm{ID}} + \omega_{\mathrm{s}}^{\mathrm{ID}} + \omega_{\mathrm{r}}^{\mathrm{ID}} + \omega_{\mathrm{v}}^{\mathrm{ID}} = 1,\quad \omega_{\mathrm{e}}^{\mathrm{ID}}, \omega_{\mathrm{s}}^{\mathrm{ID}}, \omega_{\mathrm{r}}^{\mathrm{ID}}, \omega_{\mathrm{v}}^{\mathrm{ID}} \geq 0 \end{aligned} \tag{16}$$

The first three terms follow Equation (9) at the 15 min resolution, using the reference in Equation (13). The terminal-state term remains a soft penalty, preserving sufficient flexibility to correct delivery deviations within the canal-storage bounds.

The water volume assigned to the current window by the DA schedule is as follows:

$$\begin{aligned} \overline{V}^{\mathrm{DA}}_{\mathrm{win},k} &= \sum_{j\in\mathcal{T}_k^{\mathrm{ID}}} \overline{q}^{\mathrm{DA}\to\mathrm{ID}}_{N,j|k}\Delta t^{\mathrm{ID}} \\ \sum_{j\in\mathcal{T}_k^{\mathrm{ID}}} q^{\mathrm{ID}}_{N,j|k}\Delta t^{\mathrm{ID}} &= V^{\mathrm{tar}}_{\mathrm{win},k} \end{aligned} \tag{17}$$

Here, $V^{\mathrm{tar}}_{\mathrm{win},k}$ is the corrected window target. It equals the DA-assigned volume in the absence of accumulated execution error and is otherwise adjusted by the compensation mechanism in Section 3.4.

Let $\mathcal{X}_k^{\mathrm{ID}}$ denote the feasible set formed by the 15 min constraints in Equations (1), (2) and (4)–(7), the initial conditions in Equation (14), and the window-delivery condition in Equation (17). The ID problem minimizes $J_{\mathrm{ID},k}$ over this set. Although the full rolling horizon is optimized, only the first four 15 min decisions and associated canal-state references are released to the RT layer:

$$\begin{aligned} q^{\mathrm{ID,ref}}_{i,4k+r} &= q^{\mathrm{ID},*}_{i,r|k}, \quad r = 0,\ldots,3 \\ V^{\mathrm{ID,ref}}_{c,4k+r} &= V^{\mathrm{ID},*}_{c,r|k}, \quad r = 0,\ldots,4 \end{aligned} \tag{18}$$

Later predictions are retained only to account for the future effects of the released decisions. In the complete DA+ID+RT framework, the next hourly problem is reconstructed from the observed canal-storage state and the final discharge executed by the RT layer.

For the DA+ID ablation evaluated in Section 5.5, the released 15 min ID decisions are mapped to the common 5 min execution grid without solving an RT optimization problem. Execution deviations and canal-storage states are still propagated at the 5 min simulation step, and the subsequent ID window is initialized from the resulting executed discharge and updated storage state. This configuration isolates the incremental contribution of RT feedback while retaining the ID receding-horizon correction.

3.3. Real-Time Supervisory Feedback Adjustment

The RT layer refines the released ID schedule every 5 min using the latest aggregate canal-storage and discharge information. It applies the 300 s versions of the station-level discharge, ramping, aggregate power, and canal-storage constraints. The term "real-time" denotes the update interval of the supervisory scheduling layer and should not be interpreted as continuous hydraulic or electromechanical control. Pressure regulation, variable-speed drive dynamics, motor–pump transients, and within-reach water-level propagation are not represented. The RT problem therefore addresses short-term reference correction and aggregate scheduling-constraint feasibility rather than dynamic hydraulic regulation.

Here, $n(m) \in \{0,\ldots,95\}$ identifies the 15 min ID interval containing RT interval m, and $r(m) \in \{0,1,2\}$ gives its position within that interval.

The discharge released by the ID layer is held constant over the corresponding three RT intervals. The real-time reference and corrected command are written as follows:

$$\overline{q}^{\mathrm{ID}\to\mathrm{RT}}_{i,m} = q^{\mathrm{ID,ref}}_{i,n(m)},\ q^{\mathrm{RT}}_{i,m} = \overline{q}^{\mathrm{ID}\to\mathrm{RT}}_{i,m} + \delta q^{\mathrm{RT}}_{i,m},\ i \in \mathcal{I} \tag{19}$$

The correction $\delta q^{\mathrm{RT}}_{i,m}$ is optimized every 5 min, allowing local feedback adjustment without replacing the economic allocation inherited from the DA and ID layers. Both $q^{RT}_{i,m}$ and $\delta q^{RT}_{i,m}$ are aggregate station-level quantities. They do not determine which pumping units are committed or how the total discharge is distributed among operating units. Unit combinations, adjustable speeds or blade settings, and switching decisions must be determined by a lower-level unit-dispatch controller subject to the supervisory discharge reference.

Each RT problem is initialized using the current observed canal state and the discharge executed during the preceding interval:

$$V_{c,m}^{\mathrm{RT}} = V_c^{\mathrm{obs}}(m),\ q_{i,m-1}^{\mathrm{RT}} = q_i^{\mathrm{act}}(m-1),\ \forall i \in \mathcal{I}, c \in \mathcal{C} \tag{20}$$

At $m = 0$, these values are obtained from the pre-horizon measurements; subsequent problems use feedback from the preceding RT interval.

The RT objective penalizes ID-reference tracking error, abrupt discharge variation, and deviation from the transferred canal-state reference:

$$\begin{aligned}
\hat{J}_{\mathrm{r},m}^{\mathrm{RT}} &= \frac{1}{N}\sum_{i\in\mathcal{I}}\left(\frac{q_{i,m}^{\mathrm{RT}} - \overline{q}_{i,m}^{\mathrm{ID}\to\mathrm{RT}}}{q_i^{\max} - q_i^{\min}}\right)^2 \\
\hat{J}_{\mathrm{s},m}^{\mathrm{RT}} &= \frac{1}{N}\sum_{i\in\mathcal{I}}\left(\frac{q_{i,m}^{\mathrm{RT}} - q_i^{\mathrm{act}}(m-1)}{q_i^{\max} - q_i^{\min}}\right)^2 \\
\hat{J}_{\mathrm{v},m}^{\mathrm{RT}} &= \frac{1}{\mid\mathcal{C}\mid}\sum_{c\in\mathcal{C}}\left(\frac{V_{c,m+1}^{\mathrm{RT}} - \overline{V}_{c,m+1}^{\mathrm{ID}\to\mathrm{RT}}}{V_c^{\max} - V_c^{\min}}\right)^2
\end{aligned} \tag{21}$$

The state reference $\overline{V}_{c,m+1}^{\mathrm{ID}\to\mathrm{RT}}$ is obtained from Section 3.4 and guides, rather than replaces, the water-balance state.

The RT objective is as follows:

$$\begin{aligned}
J_{\mathrm{RT},m} = \quad & \omega_{\mathrm{r}}^{\mathrm{RT}}\hat{J}_{\mathrm{r},m}^{\mathrm{RT}} + \omega_{\mathrm{s}}^{\mathrm{RT}}\hat{J}_{\mathrm{s},m}^{\mathrm{RT}} + \omega_{\mathrm{v}}^{\mathrm{RT}}\hat{J}_{\mathrm{v},m}^{\mathrm{RT}} \\
& \omega_{\mathrm{r}}^{\mathrm{RT}} + \omega_{\mathrm{s}}^{\mathrm{RT}} + \omega_{\mathrm{v}}^{\mathrm{RT}} = 1,\ \ \omega_r^{\mathrm{RT}}, \omega_s^{\mathrm{RT}}, \omega_v^{\mathrm{RT}} \geq 0
\end{aligned} \tag{22}$$

Electricity cost is omitted from the RT objective because tariff-driven discharge allocation is established by the DA layer and updated by the ID layer. The RT problem therefore focuses on reference tracking, discharge regularity, and correction of the aggregate canal-storage state. It does not introduce a separate energy-minimization objective.

Let $\mathcal{X}_m^{\mathrm{RT}}$ denote the single-step feasible set formed by the 300 s constraints in Equations (1), (2) and (4)–(7), together with Equations (19) and (20). Each RT update minimizes $J_{\mathrm{RT},m}$ over $\mathcal{X}_m^{\mathrm{RT}}$.

The global terminal-state condition in Equation (3) and the full-horizon delivery constraint in Equation (8) are not imposed on this single-step problem. Their effects are retained through the ID references, the canal-state target, and the remaining-delivery correction described in Section 3.4.

After the optimized command is issued, the realized discharge may differ from that command because of execution error. The measured discharge and resulting canal state are represented by the following:

$$\begin{aligned}
q_{i,m}^{\mathrm{act}} &= q_{i,m}^{\mathrm{RT},*} + d_{i,m} \\
P_{i,m}^{\mathrm{act}} &= \kappa_{i,m} q_{i,m}^{\mathrm{act}}, \\
V_{c,m+1}^{\mathrm{act}} &= V_{c,m}^{\mathrm{act}} + \Delta t^{\mathrm{RT}}\left(q_{i,m}^{\mathrm{act}} - q_{i+1,m}^{\mathrm{act}} + w_{c,m}\right),\ c = i \\
V_c^{\mathrm{obs}}(m+1) &= V_{\mathrm{c},m+1}^{\mathrm{act}}
\end{aligned} \tag{23}$$

Here, $d_{i,m}$ denotes the synthetic execution deviation applied to pumping station i during RT interval m. It is not included in the deterministic RT prediction but affects the realized discharge and canal-storage state in Equation (23), which are then used to initialize the next scheduling update. Exact command execution corresponds to $d_{i,m} = 0$. The construction, magnitude, dispersion, and temporal autocorrelation of the execution deviations are reported in Section 5.1. The observed canal states in the numerical study are generated from the executed discharges using the same lumped-storage balance adopted by the scheduling model. The closed-loop simulations therefore verify internal consistency, constraint satisfaction, and recovery from controlled execution deviations, but they do not constitute independent validation against field measurements or a dynamic hydraulic plant model.

After every three RT intervals, the executed discharges are aggregated to the corresponding 15 min scale:

$$\tilde{q}_{i,n}^{\mathrm{RT}} = \frac{1}{3}\sum_{r=0}^{2} q_{i,3n+r}^{\mathrm{act}},\ n = 0, 1, \ldots, 95 \tag{24}$$

The aggregated discharge, updated canal-storage states, and cumulative delivered volume are returned to the next ID update, thereby closing the supervisory DA–ID–RT scheduling loop. This loop refers to state-informed schedule correction and should not be interpreted as a complete electromechanical–hydraulic control loop.

3.4. Cross-Scale State and Water-Delivery Consistency

Because the three layers represent the same aggregate conveyance process at different temporal resolutions, their interaction must preserve water volume and transfer consistent state information. Reference trajectories are mapped from coarse to fine resolutions, whereas the realized canal states are updated from executed discharges through the lumped-storage balance. The mapping procedure preserves cross-scale consistency within the adopted supervisory model but does not introduce the transport-delay or wave-propagation relationships of a dynamic hydraulic model.

3.4.1. Cross-Scale Discharge Mapping

Each hourly DA interval corresponds to four 15 min ID intervals, and each ID interval corresponds to three 5 min RT intervals. The mapped reference discharges must preserve the water volume assigned by the upper layer:

$$\begin{array}{l} \sum\limits_{r=0}^{3} \overline{q}_{i,4(h-1)+r}^{\mathrm{DA}\to\mathrm{ID}} \Delta t^{\mathrm{ID}} = q_{i,h}^{\mathrm{DA},*} \Delta t^{\mathrm{DA}},\ \ i \in \mathcal{I},\ \ h = 1, \ldots, 24 \\ \sum\limits_{r=0}^{2} \overline{q}_{i,3n+r}^{\mathrm{ID}\to\mathrm{RT}} \Delta t^{\mathrm{RT}} = q_{i,n}^{\mathrm{ID,ref}} \Delta t^{\mathrm{ID}},\ \ \ i \in \mathcal{I},\ \ n = 0, \ldots, 95 \end{array} \tag{25}$$

Equation (25) preserves the water volume assigned by the upper layer during reference mapping. In the absence of execution deviation and subsequent feedback correction, the mapped schedules therefore retain the same assigned delivery volume. When execution deviations occur, the actual delivered volume is calculated from the realized RT discharges rather than from the mapped references.

3.4.2. Canal-State Mapping and Feedback Transfer

Under the piecewise-constant discharge mapping in Equation (25), canal-state references are interpolated between adjacent coarse-scale states:

$$\begin{array}{l} \overline{V}_{c,4(h-1)+r}^{\mathrm{DA}\to\mathrm{ID}} = V_{c,h}^{\mathrm{DA},*} + \frac{r}{4}\left(V_{c,h+1}^{\mathrm{DA},*} - V_{c,h}^{\mathrm{DA},*}\right),\ \ r = 0, \ldots, 4 \\ \overline{V}_{c,3n+r}^{\mathrm{ID}\to\mathrm{RT}} = V_{c,n}^{\mathrm{ID,ref}} + \frac{r}{3}\left(V_{c,n+1}^{\mathrm{ID,ref}} - V_{c,n}^{\mathrm{ID,ref}}\right),\ \ \ r = 0, \ldots, 3 \end{array} \tag{26}$$

The interpolated values are used only as lower-layer references. Physical canal states remain governed by the water balance and are propagated from the executed quantities in Equation (23).

At the end of each released hour, the latest RT state and executed discharge are transferred directly to the next ID window:

$$\begin{aligned} V_{c,0|k}^{\mathrm{ID}} &= V_{c,12k}^{\mathrm{RT,act}} \\ q_{i,-1|k}^{\mathrm{ID}} &= q_{i,12k-1}^{\mathrm{act}} \end{aligned} \quad k = 1, \ldots, 23 \tag{27}$$

For $k = 0$, the values are obtained from pre-horizon measurements. Equation (27) ensures that each ID window begins from the realized, rather than predicted, system state.

3.4.3. Execution-Deviation Compensation and Remaining Delivery Target

At hourly update k, the cumulative delivery deviation is evaluated as the difference between the water volume actually delivered by the RT layer and the corresponding volume assigned by the DA schedule:

$$\Delta V_k^{\mathrm{cum}} = \sum_{m=0}^{12k-1} q_{N,m}^{\mathrm{act}} \Delta t^{\mathrm{RT}} - \sum_{h=1}^{k} q_{N,h}^{\mathrm{DA},*} \Delta t^{\mathrm{DA}}, \quad k = 0, \ldots, 24 \tag{28}$$

With $\Delta V_0^{\mathrm{cum}} = 0$, positive and negative values denote accumulated over-delivery and under-delivery, respectively. Using the DA-assigned volume in Equation (17), the unprojected window target is as follows:

$$\widetilde{V}_{\mathrm{win},k}^{\mathrm{tar}} = \overline{V}_{\mathrm{win},k}^{\mathrm{DA}} - \Delta V_k^{\mathrm{cum}} \tag{29}$$

Because the target in Equation (29) may exceed the correction capability of the current window, it is projected onto the reachable delivery interval:

$$\begin{aligned} \mathcal{V}_k^{\mathrm{feas}} &= \left\{ \sum_{j \in \mathcal{T}_k^{\mathrm{ID}}} q_{N,j|k}^{\mathrm{ID}} \Delta t^{\mathrm{ID}} \mid \mathbf{x}_k^{\mathrm{ID}} \in \mathcal{X}_k^{\mathrm{ID},0} \right\} = \left[\; V_{\mathrm{win},k}^{\min}, V_{\mathrm{win},k}^{\max} \; \right] \\ V_{\mathrm{win},k}^{\mathrm{tar}} &= \Pi_{\mathcal{V}_k^{\mathrm{feas}}} \left(\hat{V}_{\mathrm{win},k}^{\mathrm{tar}} \right) = \min\left\{ V_{\mathrm{win},k}^{\max}, \max\left[V_{\mathrm{win},k}^{\min}, \hat{V}_{\mathrm{win},k}^{\mathrm{tar}} \right] \right\} \end{aligned} \tag{30}$$

The projection retains a reachable target and otherwise selects the nearest feasible boundary. Any unrecovered deviation is carried to the next update.

At the end of the operating day, delivery performance and terminal canal-state recovery are evaluated using the following:

$$\varepsilon_V^{\mathrm{day}} = \sum_{m=0}^{287} q_{N,m}^{\mathrm{act}} \Delta t^{\mathrm{RT}} - V^{\mathrm{req}}, \quad \varepsilon_c^{\mathrm{term}} = V_{c,288}^{\mathrm{RT,act}} - V_c^{\mathrm{tar}} \tag{31}$$

The two quantities measure the final delivery error and terminal canal-storage deviation from the executed trajectories and are evaluated in Section 5. The resulting delivery and terminal-state errors quantify scheduling performance within the aggregate model; they should not be interpreted as errors relative to measured hydraulic states. Figure 3 summarizes the temporal hierarchy and its connection to the spatial decomposition introduced in Section 4. To quantify the contribution of successive temporal feedback layers, Section 5.5 compares DA-only, DA+ID, and DA+ID+RT configurations under matched inputs, physical constraints, and execution deviations.

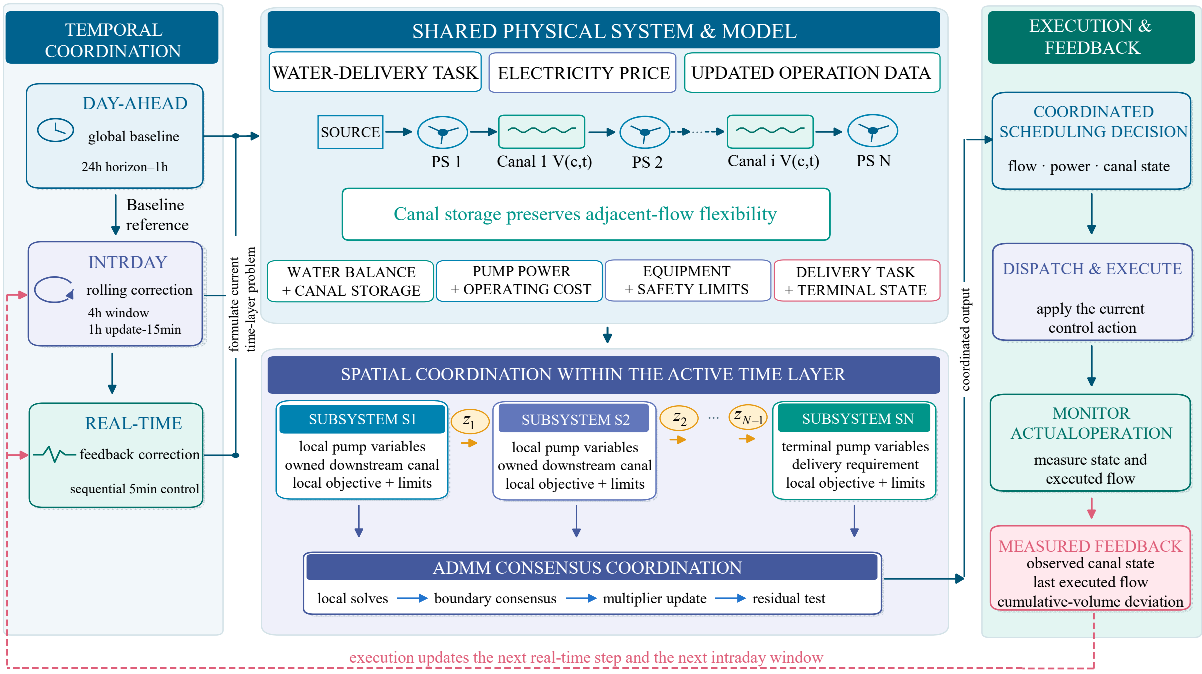

Figure 3. Overall framework of spatiotemporal coordinated scheduling.

4. Distributed Spatiotemporal Solution Method

When solved in centralized form, each temporal layer requires a single coordinator to process all station and canal-state information. To distribute this computation, the scheduling problem is partitioned into pumping-station subsystems coupled through shared boundary-flow variables. The same spatial partition is applied to the DA, ID, and RT formulations, whereas their horizons, temporal resolutions, objective weights, and boundary conditions remain layer-specific. The resulting local scheduling problems are coordinated using the alternating direction method of multipliers (ADMM). This decomposition concerns scheduling variables and interface-flow information; it does not introduce decentralized pump-unit, motor, converter, or pressure-control models.

4.1. Spatial Decomposition and Interface Consensus

Temporal superscripts are omitted in this section, and all time sets, time steps, objectives, and boundary conditions refer to the active scheduling layer. The cascade is partitioned into N subsystems $\mathcal{S}_i$. For $i < N$, subsystem $\mathcal{S}_i$ contains pumping station i and its downstream canal reaches $c = i$; $\mathcal{S}_N$ contains the terminal station and receiving-end delivery requirement. Each canal state and continuity equation therefore belongs to one subsystem.

For subsystem $\mathcal{S}_i$, $i < N$, the local water-balance equation is as follows:

$$V_{i,t+1}^{[i]} = V_{i,t}^{[i]} + \Delta t\left(q_{i,t}^{[i]} - q_{i+1,t}^{[i]} + w_{i,t}\right), \quad i = 1, \ldots, N-1 \tag{32}$$

Superscript $[i]$ denotes ownership by subsystem $\mathcal{S}_i$. In Equation (32), $q_{i,t}^{[i]}$ is the discharge of its own station, whereas $q_{i+1,t}^{[i]}$ is the local copy of the downstream-station discharge required by the canal balance.

The local decision variables are organized as follows:

$$
\begin{aligned}
\boldsymbol{x}_i &= col\left(\left\{q_{i,t}^{[i]}, P_{i,t}^{[i]}, q_{i+1,t}^{[i]}\right\}_{t\in\mathcal{T}}, \left\{V_{i,t}^{[i]}\right\}_{t\in\mathcal{T}_+}\right), \ i = 1,\dots,N-1 \\
\boldsymbol{x}_N &= col\left(\left\{q_{N,t}^{[N]}, P_{N,t}^{[N]}\right\}_{t\in\mathcal{T}}\right) \\
J_C(\boldsymbol{x}) &= \sum_{i=1}^{N} f_i(\boldsymbol{x}_i)
\end{aligned}
\tag{33}
$$

Each feasible set $\mathcal{X}_i$ contains the constraints of the components owned by $\mathcal{S}_i$. Station capacity, power, and ramping constraints are assigned to the station owner, while canal continuity, storage, and state-recovery terms are assigned to the reach owner. The delivery requirement belongs to the terminal subsystem. The local objective f_i is partitioned by the same ownership rule, ensuring that every centralized objective term is included exactly once.

At interface i_t the same physical discharge is represented by $q_{i+1,t}^{[i]}$ in the upstream subsystem and $q_{i+1,t}^{[i+1]}$ in the downstream subsystem. Consensus variable $z_{i,t}$ coordinates these two copies:

$$
q_{i+1,t}^{[i]} = z_{i,t}, \quad q_{i+1,t}^{[i+1]} = z_{i,t}, \quad i = 1,\dots,N-1, \ \ t\in\mathcal{T}
\tag{34}
$$

Equation (34) equates two copies of the same boundary flow, rather than the physical discharges of adjacent pumping stations. The difference between $q_{i,t}^{[i]}$ and $q_{i+1,t}^{[i+1]}$ remains accommodated by the canal-storage variation in Equation (32). Figure 4 summarizes the resulting subsystem ownership, boundary-flow copies, and consensus structure.

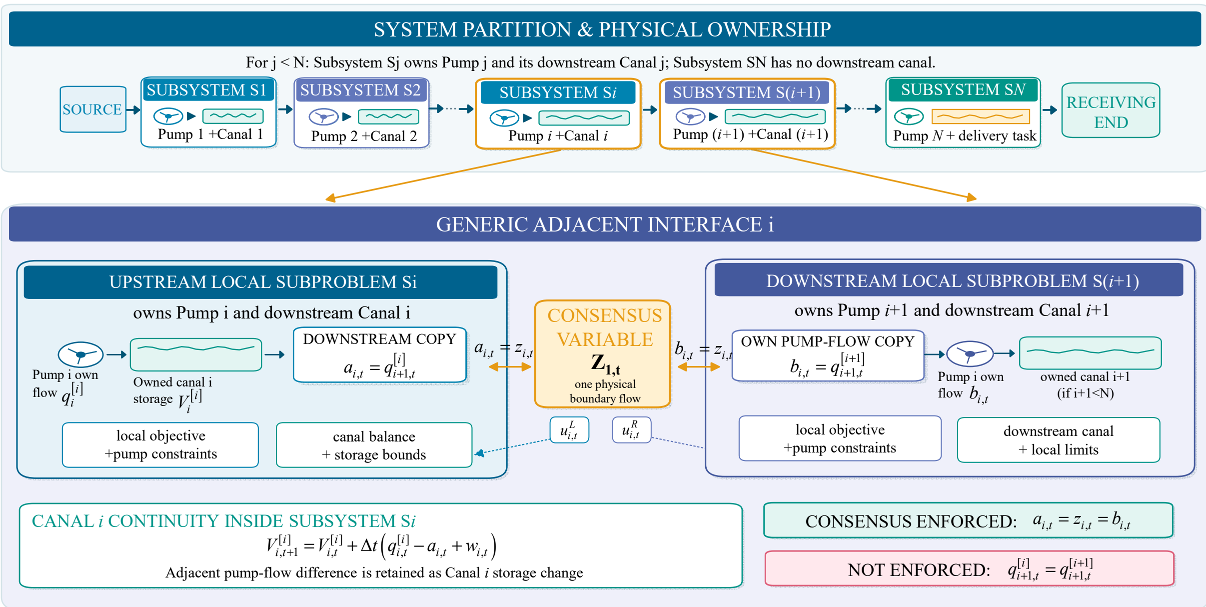

Figure 4. Spatial decomposition, boundary-flow copies, and consensus coordination.

Based on this variable-copy construction, the centralized scheduling problem can be expressed in the following equivalent distributed form:

$$
\begin{aligned}
\mathcal{P}_D: \quad & \underset{\{x_i\}_{i=1}^{N},\{z_i\}_{i=1}^{N-1}}{\text{minimize}} \ \sum_{i=1}^{N} f_i(\boldsymbol{x}_i) \\
& \text{subject to } \ \boldsymbol{x}_i \in \mathcal{X}_i, \quad i = 1,\dots,N \\
& \qquad\qquad q_{i+1,t}^{[i]} = z_{i,t}, \quad q_{i+1,t}^{[i+1]} = z_{i,t} \quad i = 1,\dots,N-1, \quad t\in\mathcal{T}
\end{aligned}
\tag{35}
$$

Equations (34) and (35) provide an exact reformulation of the continuous centralized problem because no physical constraint or objective term is duplicated or omitted. The centralized and distributed formulations therefore have the same feasible decisions and optimal objective value, although non-unique solutions need not coincide pointwise. For the serial cascade topology, each interior subsystem exchanges only two boundary-flow trajectories regardless of the total number of stations. This locality limits the communication burden of each subsystem, but does not imply that the ADMM iteration count or computation time is independent of system size; these properties are evaluated empirically in Section 5.4.

4.2. ADMM-Based Distributed Coordination

In Equation (35), the local feasible sets are coupled only by interface consensus. For each interface, collect the upstream copy, downstream copy, and consensus trajectory into vectors:

$$\boldsymbol{a}_i = \left[q_{i+1,t}^{[i]}\right]_{t\in\mathcal{T}}, \ \ \boldsymbol{b}_i = \left[q_{i+1,t}^{[i+1]}\right]_{t\in\mathcal{T}}, \ \ \mathbf{z}_i = [z_{i,t}]_{t\in\mathcal{T}} \ \boldsymbol{a}_i = \mathbf{z}_i, \ \ \boldsymbol{b}_i = \mathbf{z}_i \tag{36}$$

Here, $\boldsymbol{a}_i$ is maintained by $\mathcal{S}_i$, while $\boldsymbol{b}_i$ is maintained by $\mathcal{S}_{i+1}$, both represent the discharge trajectory of pumping station $i+1$. For compactness, $\mathbf{a}$, $\mathbf{b}$, $\mathbf{z}$, $\mathbf{u}^L$, and $\mathbf{u}^R$ subsequently denote the corresponding trajectories stacked over all interfaces. Their time components are contained in these vectors and are therefore omitted from Equations (37)–(43).

Let λ_i^L and λ_i^R denote the multipliers of the two equalities in Equation (36). Their scaled forms are $\mathbf{u}_i^L = \lambda_i^L/\rho$ and $\mathbf{u}_i^R = \lambda_i^R/\rho$, where $\rho > 0$. The scaled augmented Lagrangian is as follows:

$$\mathcal{L}_\rho = \sum_{i=1}^{N} f_i(\boldsymbol{x}_i) + \frac{\rho}{2}\sum_{i=1}^{N-1}\left(\| \boldsymbol{a}_i - z_i + \boldsymbol{u}_i^{\mathrm{L}} \|_2^2 + \| \boldsymbol{b}_i - z_i + \boldsymbol{u}_i^{\mathrm{R}} \|_2^2 - \| \boldsymbol{u}_i^{\mathrm{L}} \|_2^2 - \| \boldsymbol{u}_i^{\mathrm{R}} \|_2^2\right) \tag{37}$$

The two quadratic terms penalize disagreement between the local copies and the consensus boundary flow. Differences between adjacent physical station discharges remain governed by the canal-storage balance in Equation (32).

Let ν denote the ADMM iteration index. Given the current consensus and scaled dual variables, each subsystem solves the following:

$$\begin{aligned}
\boldsymbol{x}_i^{\nu+1} \quad &= \underset{\boldsymbol{x}_i\in\mathcal{X}_i}{\operatorname{argmin}} f_i(\boldsymbol{x}_i) \\
&+\mathbb{I}_{\{i>1\}}\tfrac{\rho}{2}\left\|\boldsymbol{b}_{i-1}(\boldsymbol{x}_i) - \boldsymbol{z}_{i-1}^{\nu} + \boldsymbol{u}_{i-1}^{\mathrm{R},\nu}\right\|_2^2 \\
&+\mathbb{I}_{\{i<N\}}\tfrac{\rho}{2}\left\|\boldsymbol{a}_i(\boldsymbol{x}_i) - \boldsymbol{z}_i^{\nu} + \boldsymbol{u}_i^{\mathrm{L},\nu}\right\|_2^2
\end{aligned} \tag{38}$$

The indicator functions remove the unavailable interface term for the first and terminal subsystems. With the consensus and dual variables fixed, all local problems can be solved independently.

After completing the local updates, the consensus trajectory at each interface is obtained in closed form:

$$z_i^{\nu+1} = \frac{1}{2}\left(\boldsymbol{a}_i^{\nu+1} + \boldsymbol{b}_i^{\nu+1} + \boldsymbol{u}_i^{\mathrm{L},\nu} + \boldsymbol{u}_i^{\mathrm{R},\nu}\right), \ \ i = 1,\ldots,N-1 \tag{39}$$

With zero-initialized scaled dual variables, Equation (39) reduces to the mean of the two local copies at each interface and time interval.

The scaled dual variables are subsequently updated as follows:

$$\begin{aligned} u_i^{L,\nu+1} &= \boldsymbol{u}_i^{L,\nu} + \boldsymbol{a}_i^{\nu+1} - \boldsymbol{z}_i^{\nu+1} \\ \boldsymbol{u}_i^{R,\nu+1} &= \boldsymbol{u}_i^{R,\nu} + \boldsymbol{b}_i^{\nu+1} - \boldsymbol{z}_i^{\nu+1} \end{aligned} \tag{40}$$

Equations (38)–(40) complete one ADMM iteration through local optimization, interface consensus, and scaled-dual updates. Convergence is monitored using the primal and dual residuals:

$$\begin{aligned} \boldsymbol{r}^{\nu+1} &= col\left(\boldsymbol{a}^{\nu+1} - \boldsymbol{z}^{\nu+1}, \boldsymbol{b}^{\nu+1} - \boldsymbol{z}^{\nu+1}\right) \\ \boldsymbol{s}^{\nu+1} &= \rho_s col\left(\boldsymbol{z}^{\nu+1} - \boldsymbol{z}^{\nu}, \boldsymbol{z}^{\nu+1} - \boldsymbol{z}^{\nu}\right) \end{aligned} \tag{41}$$

The primal residual measures boundary-copy disagreement, whereas the dual residual measures the change in the consensus trajectory between successive iterations.

$$\left\|\boldsymbol{r}^{\nu+1}\right\|_2 \leq \varepsilon_{\text{pri}}, \quad \left\|\boldsymbol{s}^{\nu+1}\right\|_2 \leq \varepsilon_{\text{dual}} \tag{42}$$

or when the maximum iteration count is reached. Section 4.4 specifies the layer-dependent penalties and stopping tolerances.

4.3. Nested Spatiotemporal Solution Procedure

The proposed method combines an outer temporal loop with the inner spatial coordination in Equations (38)–(42). For temporal layer $s \in \{\text{DA}, \text{ID}, \text{RT}\}$, $\mathcal{A}_{\text{ADMM}}$ denotes the solution of active problem $\mathcal{P}_s$ using deterministic input $\boldsymbol{\Theta}_s$, initial consensus $\boldsymbol{z}^{s,0}$, and scaled dual variable $\boldsymbol{u}^{s,0}$.

The DA inputs comprise the initial canal-storage states, initial discharges, delivery requirement, time-of-use tariff, and aggregate station and canal parameters. The ID layer additionally uses DA references, observed storage states, executed discharges, and the corrected window target, whereas the RT layer uses the released ID references and current aggregate measurements. The equivalent heads are evaluated from the reference water levels and lumped resistance representation supplied to each temporal call, while the station-level equivalent efficiencies remain fixed in the present case study. The resulting discharge-to-power coefficients are held fixed throughout all ADMM iterations of that call.

The DA problem is solved once before operation, the ID problem is reconstructed hourly, and the RT problem is updated every 5 min. Algorithm 1 specifies their invocation, reference mapping, decision release, and feedback sequence.

Algorithm 1 defines the executable sequence of the complete DA+ID+RT scheduling framework, whereas Figure 5 visualizes the relationship between the outer rolling updates and inner ADMM coordination. The DA+ID ablation follows the same procedure but omits the RT optimization loop, as specified in Section 3.2; the released ID decisions are instead mapped directly to the common 5 min execution grid.

The outer temporal loop is sequential because each update depends on the previously executed decisions and resulting states. Within each ADMM iteration, however, the local problems in Equation (38) are independent and algorithmically parallelizable. All reported wall-clock times were obtained using serial local-subproblem execution on one computer. For the scalability experiment, an ideal parallel critical-path time is additionally estimated by accumulating the longest local-subproblem time in each iteration; this quantity is an analytical estimate rather than a measured parallel runtime.

Algorithm 1 Spatiotemporal decomposition-based coordinated scheduling of cascade pumping stations

Input: Aggregate station and lumped canal parameters; V^{ini}, V^{tar}, and V^{req}; time-of-use tariff; temporal objective weights; ADMM penalties, tolerances, and iteration limits.

Output: DA, ID, and RT schedules; executed discharges and corresponding derived input-power profiles; realized aggregate canal-storage states; delivery deviation; ADMM residuals and iteration counts.

Construct the DA problem $\mathcal{P}_{\mathrm{DA}}$ using the initial system state, daily delivery requirement, electricity tariff, and physical parameters.
Solve $\mathcal{P}_{\mathrm{DA}}$ using $\mathcal{A}_{\mathrm{ADMM}}$ with the updates in Equations (38)–(42).
Map the optimized hourly discharge and canal-state trajectories to the ID resolution using Equations (13), (25) and (26).
for $k = 0, \ldots, 23$ **do**
    Read the current canal states, most recently executed discharges, and cumulative delivery deviation.
    Compute the corrected and feasible ID window target using Equations (28)–(30).
    Construct the rolling problem $\mathcal{P}_{\mathrm{ID}}(k)$ with the horizon and initial conditions defined in Equations (12) and (14).
    Solve $\mathcal{P}_{\mathrm{ID}}(k)$ using $\mathcal{A}_{\mathrm{ADMM}}$.
    Release only the first four 15 min decisions and associated canal-state references according to Equation (18).
    Map the released ID trajectories to 5 min references using Equations (19), (25) and (26).
    for $m = 12k, \ldots, 12k + 11$ **do**
        Read the current observed canal states and the discharges executed in the preceding RT interval.
        Construct the single-step problem $\mathcal{P}_{\mathrm{RT}}(m)$.
        Solve $\mathcal{P}_{\mathrm{RT}}(m)$ using $\mathcal{A}_{\mathrm{ADMM}}$.
        Issue the current 5 min discharge command, record the realized discharge, and update the realized aggregate canal-storage states using Equation (23).
    end for
    Aggregate the 12 executed RT intervals according to Equation (24).
    Transfer the actual terminal state and final executed discharge to the next ID window using Equation (27).
end for
Evaluate the daily delivery and terminal canal-state errors using Equation (31), and record the residuals, iteration counts, and computation times of all ADMM calls.

4.4. Convergence Scope and Implementation Settings

Under the adopted continuous scheduling model, the fixed within-solve head and efficiency coefficients make the power–discharge relation linear. The local objectives and feasible sets are convex, and the consensus constraints are affine. Provided that the global problem is feasible and each local subproblem is solved exactly, fixed-penalty consensus ADMM converges for any $\rho_s > 0$ the primal and dual residuals approach zero and the distributed objective approaches the corresponding centralized optimum. This result does not require a unique decision vector; centralized and distributed schedules may therefore differ pointwise while remaining equivalent in objective value and feasibility.

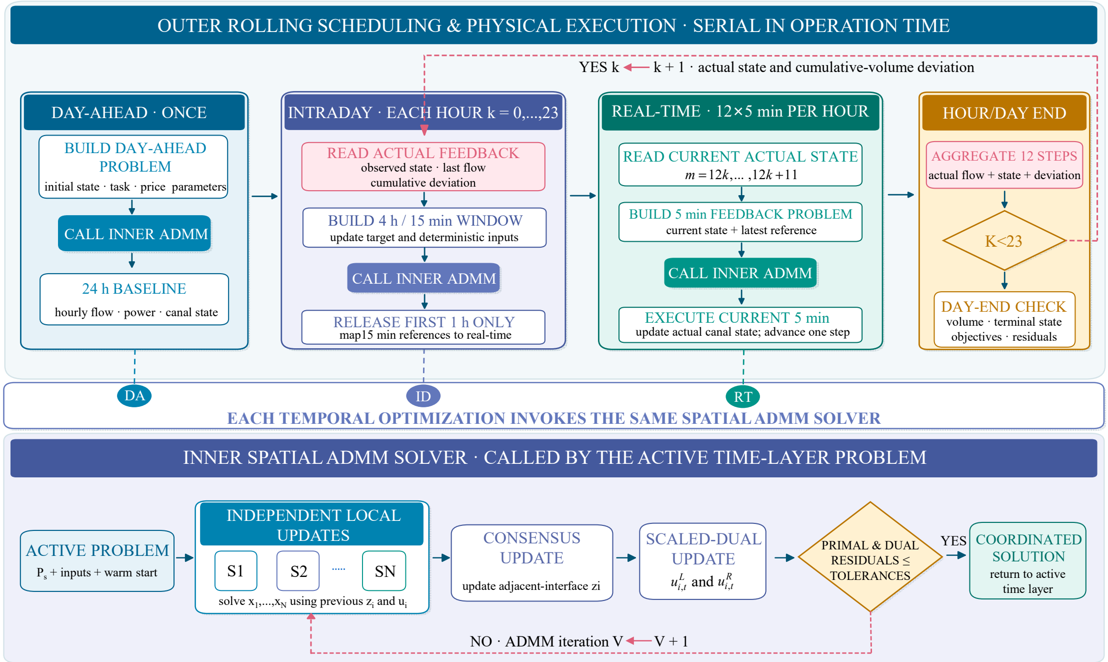

Figure 5. Nested solution procedure of outer rolling scheduling and inner ADMM coordination.

For a layer containing T_s intervals and $C = N - 1$ interfaces, let $\overline{\mathbf{a}}_s = \mathrm{col}(\mathbf{a}_s, \mathbf{b}_s)$, $\overline{\mathbf{z}}_s = \mathrm{col}(\mathbf{z}_s, \mathbf{z}_s)$, and $\overline{\mathbf{u}}_s = \mathrm{col}(\mathbf{u}_s^L, \mathbf{u}_s^R)$.

$$
\begin{aligned}
\varepsilon^{\nu}_{\mathrm{pri},s} &= \sqrt{2CT_s}\,\varepsilon_{\mathrm{abs}} + \varepsilon_{\mathrm{rel}} \max\left\{ \left\| \overline{\mathbf{a}}_s^{\nu} \right\|_2, \left\| \overline{\mathbf{z}}_s^{\nu} \right\|_2 \right\} \\
\varepsilon^{\nu}_{\mathrm{dual},s} &= \sqrt{2CT_s}\,\varepsilon_{\mathrm{abs}} + \varepsilon_{\mathrm{rel}} \rho_s \left\| \overline{\mathbf{u}}_s^{\nu} \right\|_2
\end{aligned}
\tag{43}
$$

An ADMM call is declared converged when Equation (42) and the interface-flow mismatch criterion are satisfied; otherwise, it terminates at the prescribed maximum iteration count.

The absolute and relative tolerances are 10^{-6}, and the interface-flow tolerance is 10^{-4} $\mathrm{m^3 s^{-1}}$. Unless otherwise stated, each ADMM call in the complete distributed multi-time-scale scheduling (DMTS) framework is limited to 2500 iterations. The separately evaluated DA+ID ablation uses a maximum of 10,000 iterations for its distributed ID calls while retaining the same residual and interface-feasibility criteria; no centralized fallback is used.

For the complete DA+ID+RT framework, the fixed layer-specific penalties are $\rho_{\mathrm{DA}} = 10^{-5}$, $\rho_{\mathrm{ID}} = 10^{-4}$, and $\rho_{\mathrm{RT}} = 10^{-3}$. The DA+ID ablation initializes its ID penalty at $\rho_{\mathrm{ID}} = 1$ and applies residual balancing every ten ADMM iterations, using a residual-ratio threshold of 10 and multiplicative increase and decrease factors of 2.

In the complete DMTS framework, the penalties remain fixed within each ADMM call. DA consensus is initialized from the uniform discharge associated with the daily target, whereas ID and RT consensus variables are initialized from the reference trajectory transferred by the preceding layer. The scaled dual variables are initialized to zero, and no ADMM iterates are propagated between successive rolling windows. For the DA+ID ablation, the first ID window uses zero-initialized dual variables, while subsequent ID windows use shifted consensus and scaled-dual warm starts obtained from the preceding window.

To isolate spatial scalability, Section 5.4.3 additionally evaluates the DA problem for cascades containing 5, 10, 20, and 40 pumping stations. All cases use the same fixed DA penalty $\rho_{\mathrm{DA}} = 10^{-5}$, absolute and relative tolerances of 10^{-6}, and interface-flow tolerance of 10^{-4} m^{3}s^{-1}. The maximum iteration count is set to 10,000 to avoid prematurely truncating the larger cases; the convergence tolerances remain unchanged.

Each subsystem exchanges only its boundary-flow copy and receives the updated consensus and scaled dual variables. Local parameters, objectives, states, and constraints remain within the subsystem. Although the local problems are parallelizable, Section 5 reports serial execution times on a single computer.

The preceding convergence statement applies only to the fixed-penalty continuous convex formulation used by the complete DMTS framework. Convergence of the adaptive-penalty DA+ID ablation is assessed from its recorded residuals, interface mismatch, feasibility checks, and solver termination. Introducing individual unit commitment, switching penalties, variable-speed settings, or nonlinear pump-efficiency curves would produce mixed-integer or nonconvex local subproblems; the present convergence result would then no longer apply, and a different distributed solution and convergence analysis would be required.

5. Numerical Case Study and Results

The proposed method was evaluated primarily using a synthetic representative-day case comprising five cascade pumping stations and four canal reaches. This case was used to examine tariff-based scheduling, decentralized-feedback benchmarking, temporal and spatial decomposition, ADMM convergence, execution-error sensitivity, and the influence of fixed equivalent head and efficiency. Additional constructed day-ahead cases containing 10, 20, and 40 pumping stations were used exclusively to assess the spatial scalability of the distributed coordination method. All method comparisons were conducted under matched operating limits, electricity tariffs, delivery requirements, initial canal states, and disturbance realizations. Energy consumption was calculated ex post from the simulated discharge and fixed equivalent coefficients; it was not treated as an independent decision variable or optimization objective. All reported trajectories and performance indicators were independently checked for operating-limit compliance, mass balance, convergence, and metric consistency.

5.1. Case System and Experimental Setup

5.1.1. Case System and Scheduling Data

The case comprised five aggregate pumping stations connected by four canal reaches. Each pumping station was represented as a continuous-flow supervisory subsystem rather than as an individual pump–motor unit. The formulation therefore does not distinguish centrifugal, axial-flow, vortex, or other pump configurations, and the results should be interpreted at the station scheduling level. The equivalent head, efficiency, and input-power limits used in the simulation are listed in Table 1. These parameters remained fixed throughout the representative-day calculation.

PS5, with a maximum discharge of 100 m^{3}s^{-1}, defined the conveyance bottleneck. The daily requirement was set to 90% of its 24 h capacity, equivalent to 7.776×10^6 m^3. The DA, ID, and RT models used 1 h, 15 min, and 5 min resolutions, respectively. The ID horizon was 4 h and updated hourly, while RT comprised 288 execution steps.

Canal storage was expressed relative to its reference state. Each reach had a physical range of $\pm$81,000 m^3, with 5% reserved for execution disturbances, giving a scheduling range of $\pm$76,950 m^3. Initial and terminal target deviations were both zero.

The pre-horizon discharge of each station was initialized to 90 $m^3 s^{-1}$. Symmetric ramping limits allowed station i to traverse its usable range within 1 h, giving an interval limit of $q_i^{\max}\Delta t/3600$. Net external exchange was neglected, with $w_{c,t} = 0$ for all reaches and intervals. Consequently, the uncertainty and stress tests reported below concern pumping-station execution errors only. Variability in lateral withdrawals or external canal exchanges was not included in the present numerical case and is therefore not covered by the resulting robustness claims.

Figure 6 presents the tariff, DA schedules, and execution-error field. The DA solution shifted pumping away from high-price periods while satisfying the delivery and capacity constraints.

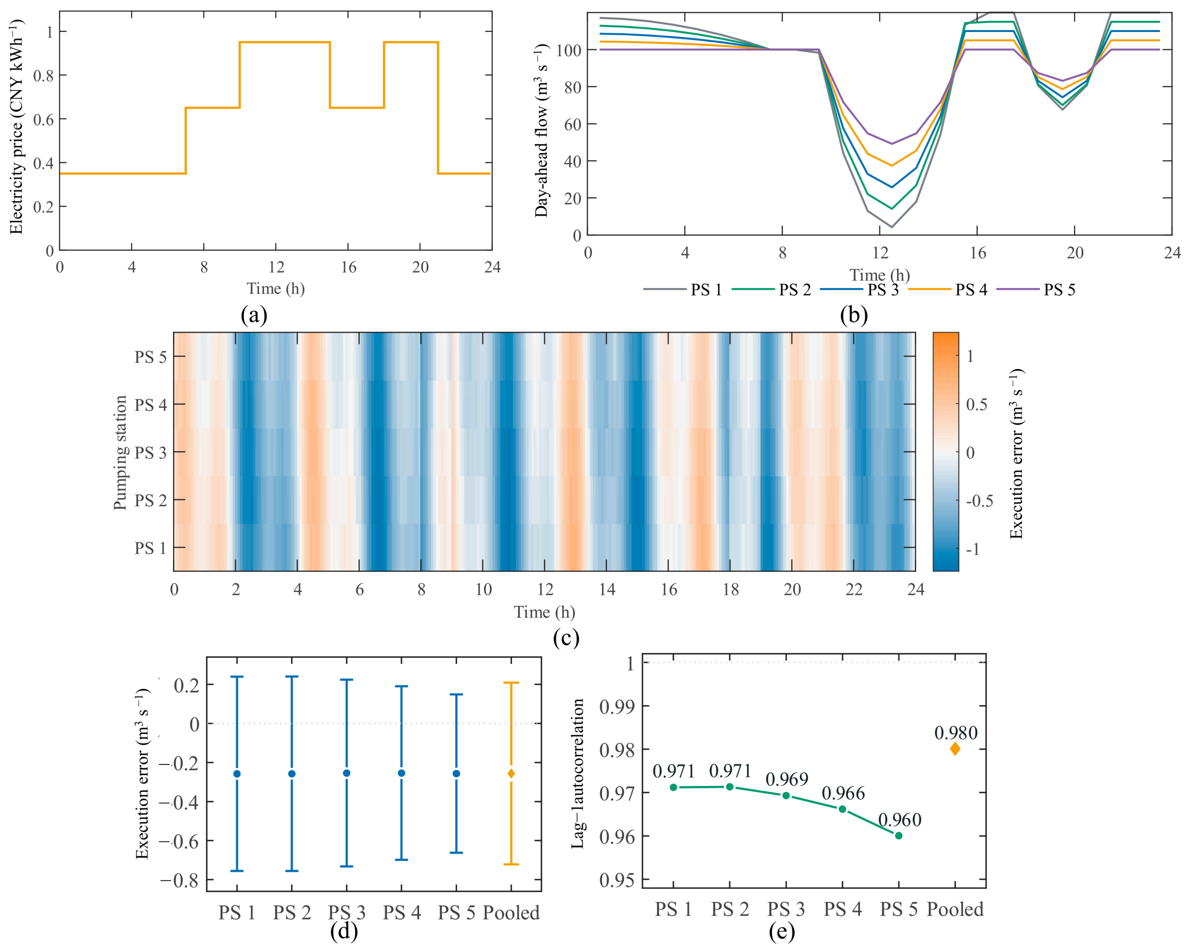

Figure 6. Representative-day scheduling conditions and execution-error characteristics. (**a**) Time-of-use electricity tariff; (**b**) day-ahead discharge schedules of the five pumping stations; (**c**) five-minute execution-error field; (**d**) station-wise mean and standard deviation of the execution error; and (**e**) station-wise and pooled lag-1 autocorrelation coefficients.

The execution-error field combined periodic, station-dependent, degradation–recovery, and random components. Across all stations and time steps, the pooled error had a mean of -0.256 m^3/s, a standard deviation of 0.466 m^3/s, a root-mean-square (RMS) value of 0.531 m^3/s, and a lag-1 autocorrelation coefficient of 0.980. The station-wise standard deviations ranged from 0.406 to 0.498 m^3/s, while the corresponding lag-1 autocorrelation coefficients ranged from 0.960 to 0.971. The maximum absolute error was 1.237 m^3/s. A fixed seed of 20260721 generated the common error realization used in all formal method comparisons. These data constitute a controlled and reproducible synthetic disturbance rather than field measurements.

Table 1. Station-level equivalent operating parameters used in the case study.

Pumping Station	Minimum Discharge (m^3/s)	Maximum Discharge (m^3/s)	Equivalent Head (m)	Overall Efficiency	Maximum Input Power (MW)
PS1	0	120	12	0.85	18
PS2	0	115	14	0.86	19
PS3	0	110	13	0.87	18
PS4	0	105	15	0.86	20
PS5	0	100	11	0.84	17

Note: Each row represents an aggregate pumping station rather than a single pumping unit. The equivalent head combines the reference static lift and lumped hydraulic loss. The equivalent efficiency is interpreted as a station-level electrical-to-hydraulic coefficient that collectively represents pump, motor, transmission, and drive losses without identifying them separately. The values of 17–20 MW are aggregate station input-power limits and should not be interpreted as individual motor ratings.

5.1.2. Benchmark Methods and Evaluation Protocol

Table 2 distinguishes the benchmark, reference, proposed, and ablation configurations. The centralized single-scale scheduling method (CSS) executes the DA schedule without feedback and is retained only as an open-loop reference. The decentralized local-feedback method (DLF) provides the principal feedback benchmark requested for comparison: each nonterminal station corrects its assigned DA reference using only the storage deviation of its own downstream canal, whereas the terminal station uses its locally accumulated delivery deviation. No boundary-flow copy, consensus variable, or interstation coordination message is exchanged during DLF execution. The centralized multi-time-scale scheduling (CMTS) and DMTS methods implement the complete DA–ID–RT hierarchy using centralized and distributed spatial solutions, respectively. The no-feedback (DMTS-NF), DA+ID (DMTS-DAID), and no-storage-flexibility (DMTS-NS) variants isolate the effects of feedback, the RT layer, and usable canal-storage flexibility, respectively.

Table 2. Benchmark, reference, proposed, and ablation methods.

Method	Spatial Implementation	Temporal Organization	Feedback Mechanism	Canal Storage	Experimental Purpose
CSS	Centralized DA solution	DA only	None	Enabled	Open-loop reference
DLF	Local execution; no consensus exchange	DA reference plus 5-min local correction	Local storage/delivery feedback	Enabled	Decentralized-feedback benchmark
CMTS	Centralized	DA–ID–RT	Coordinated rolling feedback	Enabled	Reference for spatial decomposition
DMTS	Distributed ADMM	DA–ID–RT	Coordinated rolling feedback	Enabled	Proposed method
DMTS-NF	Distributed ADMM	DA only	None	Enabled	DA-only temporal ablation
DMTS-DAID	Distributed ADMM	DA–ID	Hourly ID feedback; no RT optimization	Enabled	RT-layer ablation
DMTS-NS	Distributed ADMM	DA–ID–RT	Coordinated rolling feedback	± 1 m^3 state band	Canal-storage-flexibility ablation

The two DLF correction time constants were calibrated independently using seed 20260720. Sixteen combinations were evaluated by selecting each time constant from {1800, 3600, 7200, 14,400} s. The feasible combination minimizing the normalized delivery-error, canal-state, and flow-variation score was $\tau_v = \tau_d = 1800$ s. These values were then frozen and evaluated using the common formal-comparison seed 20260721.

Performance was assessed using specific energy, specific operating cost, delivered volume, absolute delivery error, flow variation, and canal-state deviation. Specific energy and specific cost were calculated by dividing the corresponding aggregate quantities by the actual delivered volume. The CSS comparison quantifies the effect of introducing feedback, whereas DLF provides the more demanding benchmark for assessing the value of coordinated rather than merely local feedback. Centralized–distributed consistency was evaluated separately using flow-trajectory differences, interface mismatch, iteration count, and runtime.

Calculations used MATLAB R2023b and "quadprog" on a 64-bit Windows 11 computer with an Intel Core i9-13900HX processor and 16 GB RAM. ADMM followed Section 4.4, except that the DMTS-NS iteration limit was increased from 2500 to 5000 because of its restrictive state bounds.

Local subproblems were executed sequentially; reported runtimes exclude communication and parallel acceleration. Because the experiments were deterministic single-run calculations with fixed inputs, comparisons are reported as numerical and relative differences rather than statistical tests.

5.2. Multi-Objective Parameter Selection

5.2.1. Candidate Generation and Screening

The DA weights for cost, smoothness, tracking, and canal-state regulation were fixed at [0.75, 0.25, 0, 0]. ID and RT weights were selected using CMTS and then applied unchanged to CMTS and DMTS. For ID, the cost, smoothness, and tracking weights were sampled from {0.15, 0.25, 0.35, 0.45, 0.55}, {0.05, 0.10}, and {0.15, 0.25, 0.35}, respectively. Requiring the residual canal-state weight to be at least 0.10 yielded 28 candidates. For RT, the cost weight was zero, while smoothness and tracking were sampled from {0.05, 0.10, 0.15} and {0.30, 0.40, 0.50, 0.60}. All 12 combinations satisfied the minimum canal-state weight of 0.20.

The 28 ID candidates were first evaluated with a fixed RT setting, and six passed the engineering screening. Combining these six candidates with all 12 RT settings produced 72 joint cases and 100 evaluations in total. A candidate was retained if its absolute delivery error did not exceed 0.05%, terminal state deviation remained below 4.05×10^3 m^3, and flow variation did not exceed 8 $(\mathrm{m^3 s^{-1}})^2$. Feasibility and convergence were also required.

5.2.2. Pareto Front and Final Weight Selection

Among the 72 joint candidates, 36 passed all screening criteria and 21 were Pareto non-dominated in cost, delivery error, terminal-state deviation, and flow variation. Figure 7 also retains the rejected candidates to show how the engineering criteria restricted the admissible region.

Qualified costs ranged from 912.796×10^3 to 913.628×10^3 CNY, with delivery errors below 0.05%. The minimum-cost candidate achieved 912.796×10^3 CNY but had a delivery error of 0.04718%, close to the acceptance limit. Its terminal-state deviation and flow-variation index were 2.313×10^3 m^3 and 5.797 $(\mathrm{m^3 s^{-1}})^2$, respectively.

The final setting was selected among candidates within 0.05% of the minimum cost, with successive priority given to delivery error, terminal-state deviation, flow variation, and cost. It incurred 913.229×10^3 CNY, only 0.0474% above the minimum, while reducing delivery error to 0.02012%. Its terminal-state deviation and flow variation were

1.844×10^3 m^3 and 5.267 (m^{3}s^{-1})^2. The selected ID and RT weight vectors were [0.15, 0.10, 0.15, 0.60] and [0, 0.05, 0.30, 0.65], respectively, and were fixed for all subsequent CMTS, DMTS, and ablation experiments.

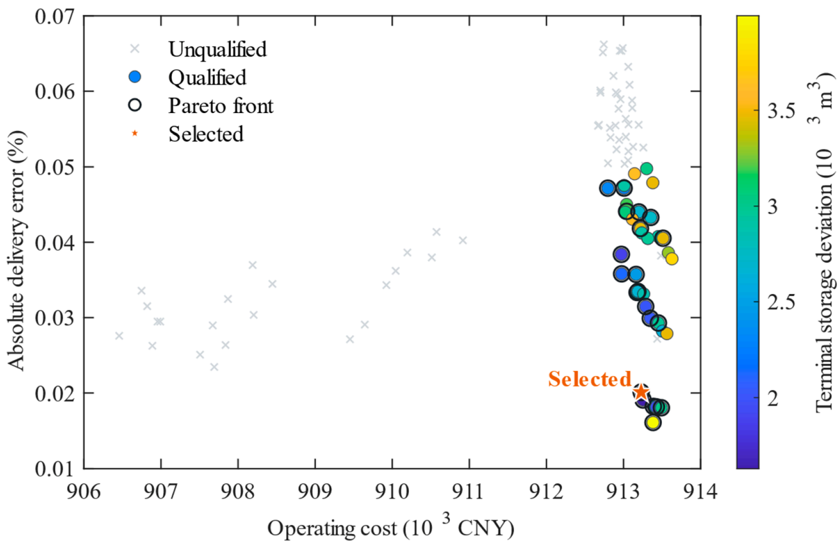

Figure 7. Candidate screening and Pareto-based parameter selection.

5.3. Multi-Time-Scale Supervisory Feedback Performance

5.3.1. Day-Ahead References, Rolling Commands, and Simulated Execution

Figure 8 compares the DA reference, rolling command, and simulated executed discharge at PS1, PS3, and PS5. The rolling supervisory updates retained the main temporal pattern of the DA schedule while smoothing its hourly transitions, particularly during 10–15 h and 18–22 h.

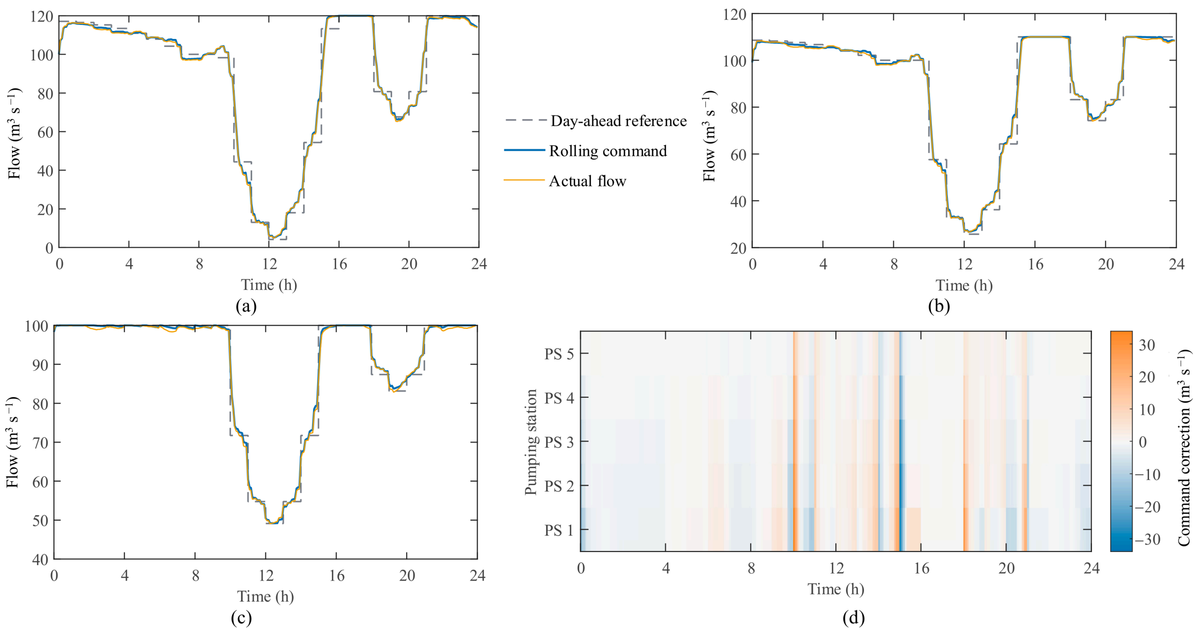

Figure 8. Reference, command, and simulated executed discharge trajectories under DMTS. (**a**) PS1; (**b**) PS3; (**c**) PS5; and (**d**) spatiotemporal corrections of the rolling commands relative to the day-ahead reference.

Corrections concentrated near DA operating transitions and remained small during steady periods. Actual flows closely tracked the commands, with a maximum command–

execution difference of 1.237 m^3s^{-1}. The ID and RT layers therefore corrected the DA reference locally in time without replacing its tariff-driven temporal pattern.

5.3.2. Water-Delivery Performance and Canal-State Evolution

Table 3 compares the principal results on a delivered-volume basis. The aggregate energy and total operating cost previously reported as direct comparative metrics were replaced by specific energy and specific cost because the methods did not transport exactly the same volume. Absolute delivery error is used instead of shortage so that both under-delivery and over-delivery are counted consistently.

Table 3. Delivered-volume-normalized operating performance of the principal scheduling methods.

Method	Specific Energy (kWh/m^3)	Specific Operating (CNY/m^3)	Delivered Volume (10^6 m^3)	Absolute Delivery Error (10^3 m^3)	Absolute Delivery Error (%)	Flow Variation (m^3 s^{-1})^2	Maximum Canal-State Deviation (10^3 m^3)	Terminal Canal-State Deviation (10^3 m^3)
CSS	0.206797	0.116727	7.750060	25.940	0.3336	25.551	79.622	2.087
DLF	0.207669	0.121594	7.796049	20.049	0.2578	6.700	73.367	32.108
DMTS-DAID	0.206052	0.116692	7.770630	5.370	0.0691	6.905	80.410	17.847
CMTS	0.206766	0.117466	7.774435	1.565	0.0201	5.267	74.915	1.844
DMTS	0.206778	0.117435	7.773996	2.004	0.0258	5.600	74.829	1.567

Relative to DLF, DMTS reduced the absolute delivery error by 90.01%, the flow-variation index by 16.42%, and the terminal canal-state deviation by 95.12%. It also reduced specific energy and specific operating cost by 0.43% and 3.42%, respectively. The maximum canal-state excursion increased by 1.99% relative to DLF but remained within the prescribed scheduling range. These results show that the improvement over DLF arose from coordinated feedback rather than from the mere presence of feedback.

CSS is retained to illustrate the accumulation of execution errors under open-loop schedule execution. Relative to CSS, DMTS reduced the absolute delivery error and flow variation by 92.28% and 78.08%, respectively. Its specific cost increased by 0.61%, whereas its specific energy differed by less than 0.01%. This comparison is therefore treated as evidence of the feedback effect rather than as the principal demonstration of coordinated-scheduling superiority.

Figure 9 compares cumulative delivery and tracking deficit. Because the total delivered volumes nearly overlap, Figure 9b provides the clearer comparison.

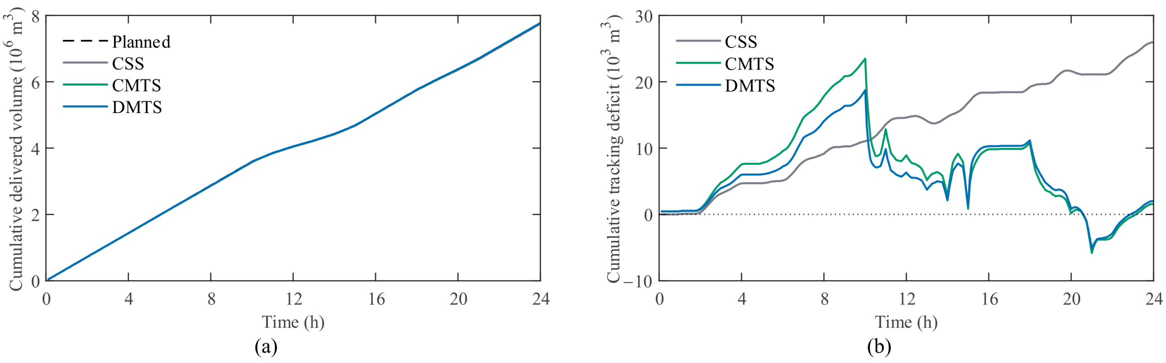

Figure 9. Cumulative water-delivery performance of CSS, CMTS, and DMTS. (**a**) Cumulative delivered volume and (**b**) accumulated deviation from the planned delivery trajectory.

Under CSS, the cumulative delivery deficit increased to 25.940 × 10^3 m^3 because the DA schedule was executed without feedback. CMTS and DMTS repeatedly updated their remaining targets and progressively recovered the accumulated deficit after 20 h. This trajectory comparison isolates the effect of introducing rolling feedback; the additional value of coordinated feedback over local feedback is evaluated separately against DLF in Table 3.

Figure 10 shows alternating storage and release in the four canal reaches, which buffered temporary discharge differences between adjacent stations.

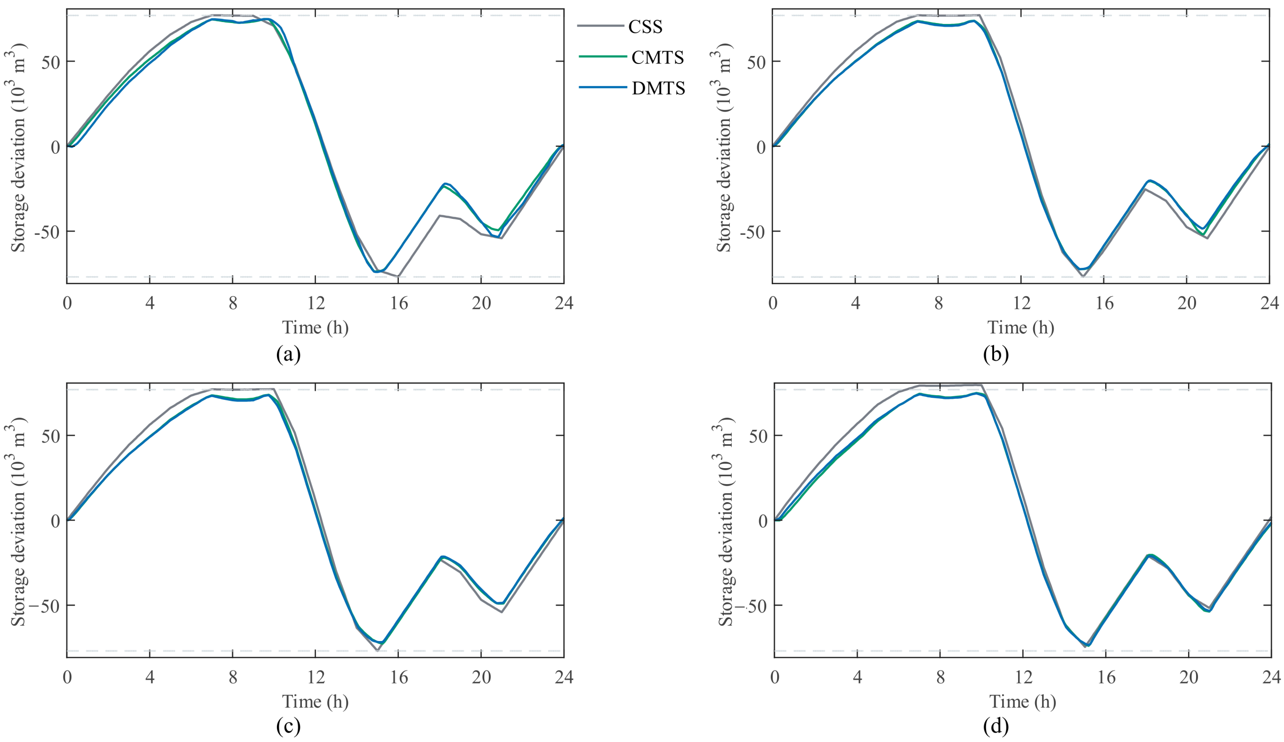

Figure 10. Canal-storage deviation under CSS, CMTS, and DMTS. (**a**) Canal 1; (**b**) Canal 2; (**c**) Canal 3; and (**d**) Canal 4. The horizontal dashed lines denote the scheduling and physical storage bounds.

The maximum CSS state deviation was 79.622 × 10^3 m^3, within the physical limit but beyond the scheduling range. CMTS and DMTS remained within the scheduling range, with maxima of 74.915 × 10^3 and 74.829 × 10^3 m^3, respectively. Their maximum terminal deviations were 1.844 × 10^3 and 1.567 × 10^3 m^3, compared with 2.087 × 10^3 m^3 for CSS. The coordinated methods therefore used temporary canal storage to absorb interstation flow differences while recovering the terminal state, rather than relying on persistent storage depletion.

5.4. Distributed Coordination Accuracy and Computational Performance

5.4.1. Centralized–Distributed Scheduling Consistency

CMTS and DMTS used identical models and rolling rules, differing only in centralized versus ADMM-based spatial solution. Figure 11 compares their system-level outputs.

The terminal-station trajectories largely overlapped at the system scale, although the distributed and centralized solutions were not pointwise identical. The maximum local flow difference was 3.431 m^3/s and occurred during a rolling transition. At the end of the scheduling horizon, the delivered-volume difference was 439.297 m^3, corresponding to 0.006% of the common target. The differences in specific energy and specific operating cost were 0.006% and 0.026%, respectively. These differences are smaller than the resolution justified by the fixed equivalent head and efficiency and the omitted hydraulic dynamics.

The results therefore indicate no discernible system-level performance difference between CMTS and DMTS under the present aggregate model, rather than establishing meaningful superiority of either solution.

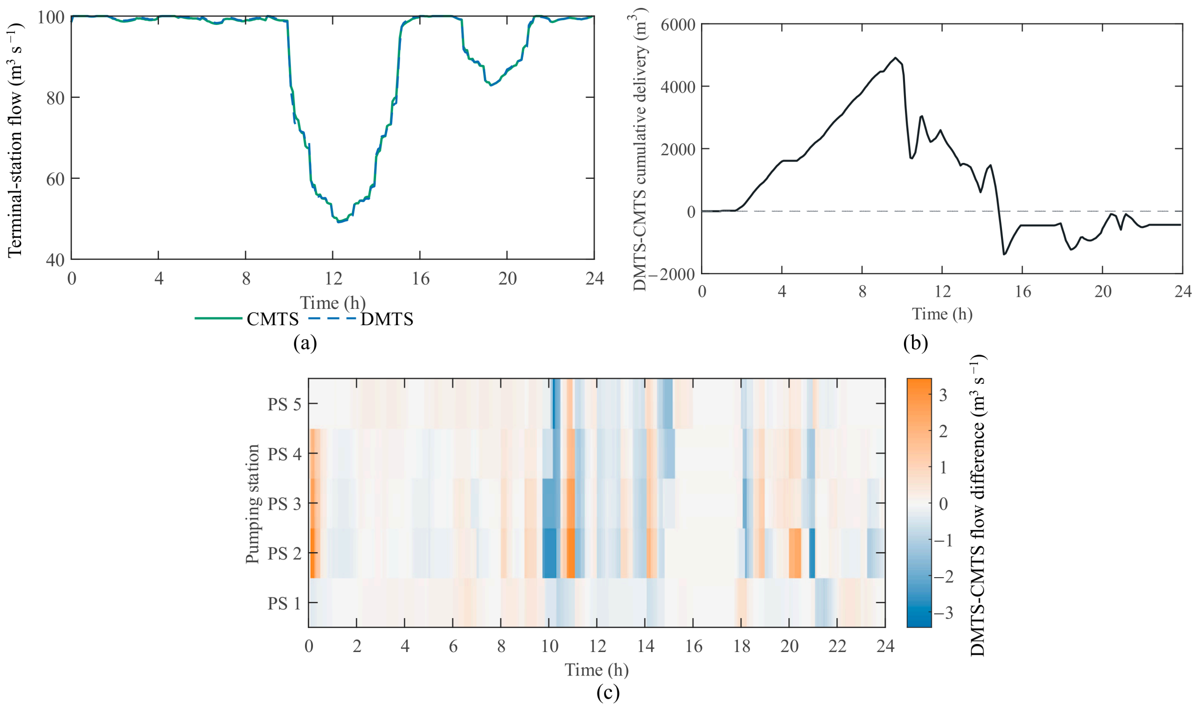

Figure 11. Comparison between centralized and distributed multi-time-scale scheduling. (**a**) Actual discharge of the terminal pumping station; (**b**) cumulative delivered-volume difference between DMTS and CMTS; and (**c**) spatiotemporal differences in actual pumping-station discharges.

5.4.2. ADMM Convergence and Computational Cost

Figure 12 summarizes DMTS convergence and computation across the three layers. The DA residuals reached their stopping tolerances after 1492 iterations.

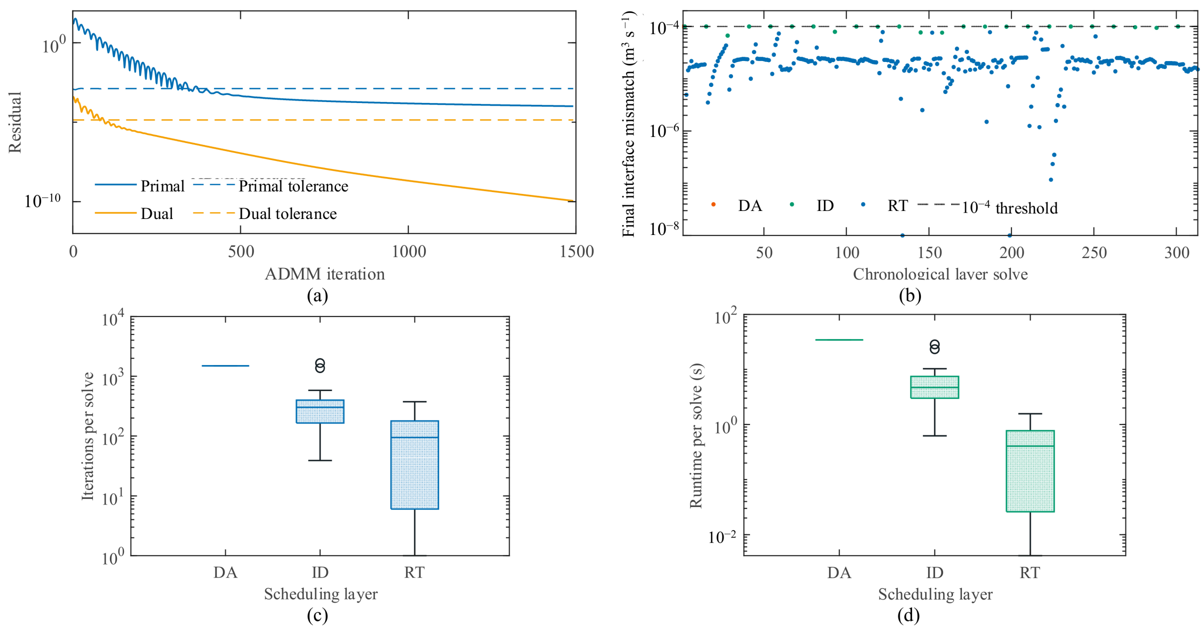

Figure 12. ADMM convergence and computational performance of DMTS. (**a**) Primal and dual residuals of the day-ahead problem; (**b**) final interface mismatch for all chronological layer solves; (**c**) iteration-count distributions; and (**d**) runtime distributions.

Over the complete 24 h scheduling horizon, DMTS performed 313 distributed optimizations, comprising one DA problem, 24 ID problems, and 288 RT problems. Every optimization terminated normally. The maximum final interface mismatch was 9.997×10^{-5} m^3 s^{-1}, which remained below the prescribed threshold of 1.0×10^{-4} m^3 s^{-1}. Figure 12b further shows that this consistency requirement was satisfied at all three temporal layers.

Table 4 reports the iteration and runtime distributions for all three layers.

Table 4. ADMM performance at different temporal scheduling layers.

Scheduling Layer	Penalty Parameter	Number of Solves	Median Iterations	Maximum Iterations	Total Iterations	Cumulative Runtime (s)
DA	1×10^{-5}	1	1492	1492	1492	34.509
ID	1×10^{-4}	24	301.5	1661	9081	157.995
RT	1×10^{-3}	288	95	374	31,071	132.793
Total	—	313	—	1661	41,644	325.297

Mean ID and RT runtimes were 6.583 s and 0.461 s, respectively, well below their 1 h and 5 min update intervals. The full process required 325.297 s and 41,644 ADMM iterations.

These serial single-computer timings exclude communication delays and establish numerical feasibility only for the five-station case. They do not demonstrate parallel acceleration or large-scale scalability.

5.4.3. Spatial Scalability with Increasing Numbers of Subsystems

To examine spatial scalability without confounding it with the rolling ID and RT updates, the 24 h DA coordination problem was further solved for constructed serial cascades containing 5, 10, 20, and 40 pumping stations. For the expanded cases, the five station-parameter profiles listed in Table 1 were repeated cyclically, while the added canal reaches adopted the same storage bounds as the base case. The delivery target, tariff profile, objective weights, ADMM penalty parameter, and stopping tolerances were kept unchanged. These cases were designed as controlled computational tests of the spatial decomposition method rather than representations of specific 10–40-station water-transfer projects.

In Table 5, all four cases satisfied the residual criteria and the prescribed interface-flow tolerance. As the number of stations increased from 5 to 40, the ADMM iteration count increased from 1492 to 7453, whereas the relative objective gap remained below 0.012%. The maximum interface mismatch was below 1.0×10^{-4} m^3 s^{-1} in every case. The distributed solution therefore retained numerical consistency with the corresponding centralized solution over the tested range, although a larger cascade required more coordination iterations.

Table 5. DA ADMM convergence and computation with increasing cascade size.

Number of Stations	Centralized Decision Variables	ADMM Iterations	Centralized Runtime (s)	Serial ADMM Runtime (s)	Estimated Parallel Critical-Path Time (s)	Objective Gap (%)	Maximum Interface Mismatch (m^3 s^{-1})
5	220	1492	0.023	35.320	8.676	1.57×10^{-5}	9.997×10^{-5}
10	465	2741	0.032	136.263	16.446	1.09×10^{-3}	9.997×10^{-5}
20	955	3730	0.059	400.685	25.361	7.09×10^{-3}	9.958×10^{-5}
40	1935	7453	0.137	1596.805	54.087	1.11×10^{-2}	9.725×10^{-5}

The measured serial ADMM runtime increased from 35.320 s for the five-station case to 1596.805 s for the 40-station case and remained substantially longer than the centralized runtime. Thus, the present serial MATLAB implementation does not provide a computational advantage over the monolithic quadratic program. By contrast, the estimated parallel critical-path time increased from 8.676 to 54.087 s, indicating that concurrent local-subproblem solution could moderate the growth in wall-clock time. This observation reflects algorithmic parallelizability rather than demonstrated parallel acceleration; verification on a parallel computing architecture with explicit communication overhead remains necessary.

These results support the numerical applicability of the proposed coordination structure to larger serial cascades governed by the same aggregate convex model. They do not establish equivalent performance for branching hydraulic networks, systems with material transport delays, nonlinear hydraulic transients, or unit-level commitment and switching decisions. Such applications require system-specific model calibration and additional validation.

5.5. Ablation and Execution-Disturbance Analyses

DMTS-NF, DMTS-DAID, and DMTS-NS were used to identify the separate contributions of the ID layer, the RT layer, and usable canal-storage flexibility. Figure 13 combines the original mechanism ablations with the new temporal-layer comparison.

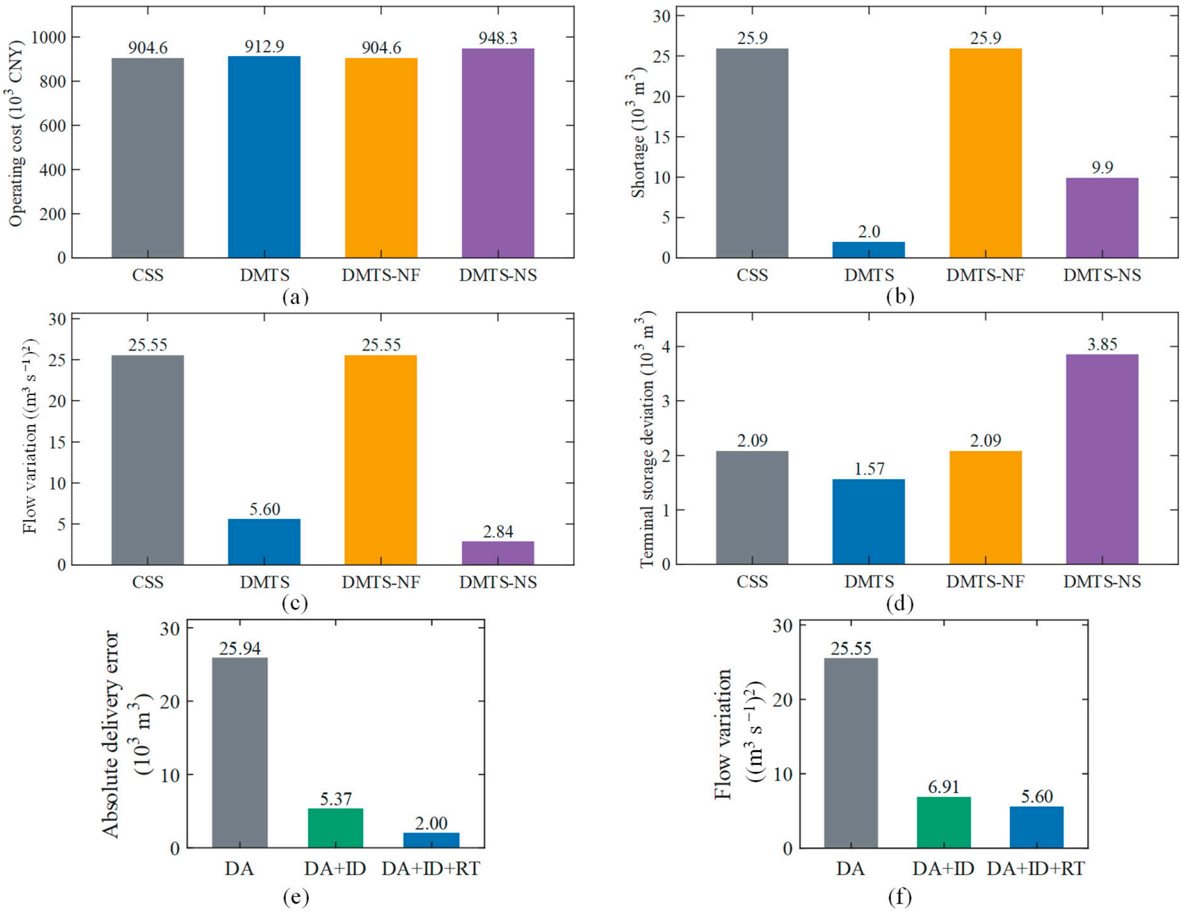

Figure 13. Ablation results for temporal feedback and canal-storage flexibility. (**a**) Specific operating cost; (**b**) absolute delivery error; (**c**) flow variation; (**d**) maximum terminal canal-state deviation; (**e**) absolute delivery error under DA-only, DA+ID, and DA+ID+RT scheduling; and (**f**) the corresponding flow-variation indices.

5.5.1. Incremental Contributions of the ID and RT Layers

The temporal ablation compares DMTS-NF (DA only), DMTS-DAID (DA+ID), and complete DMTS (DA+ID+RT). Adding the ID layer reduced the absolute delivery error from 25.940×10^3 to 5.370×10^3 m^3 and the flow-variation index from 25.552 to 6.905 (m^3/s)^2, corresponding to reductions of 79.30% and 72.98%, respectively. This confirms that hourly target correction accounts for the largest part of the improvement over open-loop DA execution.

Adding the RT layer further reduced the absolute delivery error by 62.68% and the flow-variation index by 18.91% relative to DA+ID. More importantly, the maximum canal-state excursion decreased from 80.410×10^3 to 74.829×10^3 m^3, returning the trajectory from the physical reserve region to the prescribed scheduling range, while the terminal deviation decreased from 17.847×10^3 to 1.567×10^3 m^3. These improvements were accompanied by a 0.64% increase in specific cost relative to DA+ID. The RT layer therefore contributes primarily to short-horizon feasibility recovery and terminal-state regulation rather than to the main economic schedule.

5.5.2. Effect of Reduced Canal-Storage Flexibility

Restricting canal-state variation in DMTS-NS increased the absolute delivery error to 9.934×10^3 m^3, the specific cost to 0.122107 CNY/m^3, and the terminal canal-state deviation to 3.855×10^3 m^3. Relative to DMTS-NS, complete DMTS reduced these quantities by 79.83%, 3.83%, and 59.34%, respectively. DMTS-NS produced the lowest flow-variation index, 2.843 (m^3/s)^2, but this apparent smoothness was obtained by suppressing the storage flexibility needed for delivery correction and terminal-state recovery.

The two ablations identify complementary roles: the ID and RT layers correct accumulated delivery and state deviations at different temporal resolutions, while canal storage provides the physical flexibility required to implement those corrections without excessive interstation flow changes.

5.5.3. Response to Increased Execution-Error Amplitude

To examine whether the nominal execution error was too small to challenge the feedback mechanism, the complete error sequence was multiplied by two while its temporal pattern and autocorrelation were retained. The pooled standard deviation and RMS value consequently increased from 0.466 and 0.531 m^3/s to 0.931 and 1.063 m^3/s, respectively. No model parameter or controller setting was retuned for this test.

In Table 6, under the doubled disturbance, all DMTS layer solves converged and the maximum canal-state deviation remained below the scheduling bound of 76.950×10^3 m^3. The absolute delivery error increased from 0.0258% to 0.1012%, while specific cost and flow variation increased by only 0.24% and 1.31%, respectively. The method therefore retained feasibility and bounded performance degradation within the tested twofold disturbance range. This result should not be interpreted as robustness to arbitrary disturbances or to untested lateral-withdrawal variability.

Table 6. DMTS performance under nominal and doubled execution-error amplitudes.

Execution-Error Amplitude	Pooled Standard Deviation (m^3/s)	Pooled RMS (m^3/s)	Absolute Delivery Error (10^3 m^3)	Absolute Delivery Error (%)	Specific Energy (kWh/m^3)	Specific Cost (CNY/m^3)	Flow Variation ($(m^3/s)^2$)	Maximum Canal-State Deviation (10^3 m^3)	ADMM Convergence
Nominal	0.466	0.531	2.004	0.0258	0.206778	0.117435	5.600	74.829	All layer solves converged
2× nominal	0.931	1.063	7.871	0.1012	0.206865	0.117719	5.673	75.050	All layer solves converged

5.6. Reliability–Cost Trade-Off Analysis

5.6.1. Ex-Post Reliability-Adjusted Cost

To compare delivery reliability and electricity expenditure on a common per-unit basis, an ex-post reliability-adjusted specific cost was calculated by adding the evaluated terminal-shortage cost to the electricity cost and dividing the result by the actual delivered volume. The shortage-evaluation coefficient was used only in post-processing and does not represent an electricity tariff, contractual penalty, or optimization parameter.

As shown in Figure 14, the reliability-adjusted specific cost of CSS increased most rapidly because its terminal shortage reached 25.940×10^3 m^3. The corresponding slopes for CMTS and DMTS were substantially smaller because their shortages were limited to 1.565×10^3 and 2.004×10^3 m^3, respectively. The CMTS and DMTS curves consequently remained close over the evaluated coefficient range.

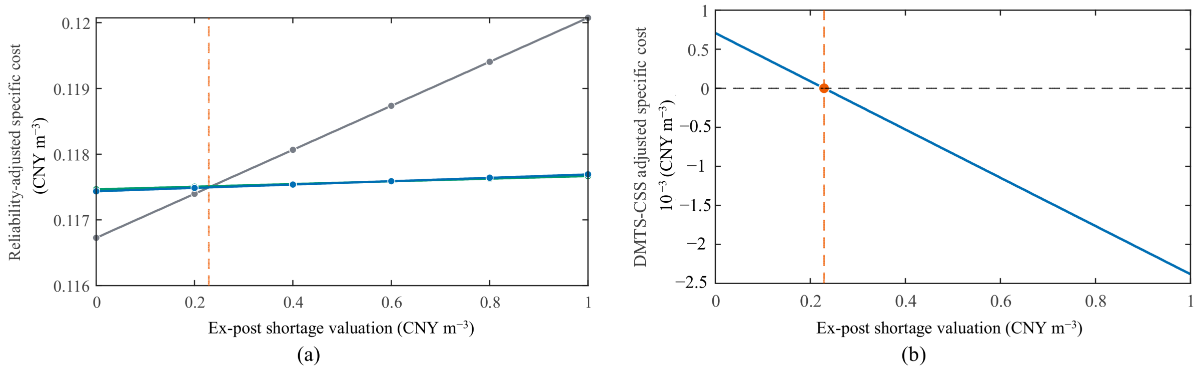

Figure 14. Reliability–cost trade-off expressed per unit of delivered water. (**a**) Ex-post reliability-adjusted specific costs of CSS, CMTS, and DMTS; and (**b**) the difference in reliability-adjusted specific cost between DMTS and CSS. The vertical dashed line denotes the break-even shortage-evaluation coefficient.

5.6.2. Break-Even Relationship

On the delivered-volume-normalized basis, DMTS had a direct specific-cost increment of 0.000707 CNY/m^3 relative to CSS but a substantially smaller shortage contribution. The resulting break-even shortage-evaluation coefficient was 0.229 CNY/m^3. At coefficients of 0.4 and 1.0 CNY/m^3, the reliability-adjusted specific cost of DMTS was lower than that of CSS by 5.28×10^{-4} and 2.38×10^{-3} CNY/m^3, respectively.

CMTS and DMTS remained nearly indistinguishable on this metric. DMTS had a slightly lower adjusted specific cost at small coefficients, whereas the smaller CMTS shortage became dominant above approximately 0.55 CNY/m^3. These small differences are consistent with the absence of a discernible system-level performance difference in Section 5.4.

The break-even coefficient is specific to the common tariff, delivery target, and synthetic disturbance used in this case. It is an ex-post interpretation of the trade-off and should not be treated as a universal shortage valuation or as part of the optimization objective.

5.7. Sensitivity to Equivalent Head and Efficiency

Because the present scheduling model uses fixed station-level equivalent head and efficiency, a deterministic sensitivity analysis was performed to quantify the effect of these assumptions. The equivalent heads or efficiencies of all five stations were uniformly varied by ±5%, while the tariff, delivery target, execution-error realization, constraints, and temporal scheduling settings were held unchanged. Two combined cases were also considered: a favorable case with a 5% head reduction and a 5% efficiency increase, and an adverse case with a 5% head increase and a 5% efficiency reduction. CMTS was used for the primary sensitivity comparison to isolate parameter effects from distributed-solver differences.

As shown in Figure 15, a 5% reduction and increase in equivalent head changed the specific energy by −5.00% and +4.99%, respectively, and changed the specific cost by −4.98% and +5.06%. A 5% reduction in equivalent efficiency increased specific energy and specific cost by 5.26% and 5.34%, whereas a 5% efficiency increase reduced them by 4.76% and 4.74%. The favorable combined case reduced specific energy and specific cost by 9.52% and 9.49%, while the adverse case increased them by 10.48% and 11.15%. All tested cases remained feasible and converged.

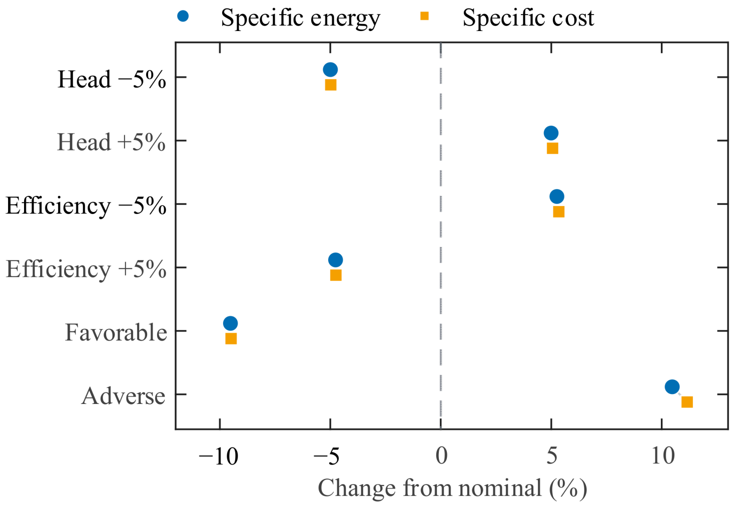

Figure 15. Sensitivity of ex-post specific energy and specific operating cost to station-level equivalent head and efficiency.

The combined ±5% perturbations changed the absolute energy and cost estimates by approximately 10–11%, which is substantially larger than the numerical difference between the centralized and distributed solutions. Together with the feasibility, mass-balance, interface-consistency, and residual checks, this analysis supports the numerical adequacy of the aggregate model for comparative supervisory scheduling within the tested parameter range. It does not validate within-reach hydraulic dynamics or unit-level converter–motor–pump behavior, nor does it quantify energy savings or switching reductions attributable to unit-combination optimization.

6. Conclusions

This study developed a spatiotemporal decomposition-based supervisory scheduling method for cascade pumping stations. A centralized convex model provides the aggregate water-balance and operating constraints shared by the DA, ID, and RT layers. Reference mapping, aggregate-state feedback, and remaining-delivery correction connect the three temporal resolutions, while ADMM coordinates station–canal subsystems through local copies of the same boundary flow. This construction preserves canal storage as the mechanism accommodating unequal adjacent-station discharges. The method addresses tariff-based scheduling and delivery regulation at the station supervisory level rather than unit-level energy or electromechanical control.

In the representative five-station case, DMTS reduced the absolute delivery error, flow-variation index, and terminal canal-state deviation by 90.01%, 16.42%, and 95.12%, respectively, relative to decentralized local feedback without consensus exchange. Its specific operating cost was also 3.42% lower. The temporal ablation showed that adding the ID layer reduced the absolute delivery error by 79.30% relative to DA-only scheduling, while the RT layer produced a further reduction of 62.68%. The RT layer also returned the maximum canal-state excursion to the prescribed scheduling range and reduced the terminal deviation from 17.847×10^3 to 1.567×10^3 m^3.

No discernible system-level performance difference was resolved between centralized and distributed multi-time-scale scheduling under the aggregate model. Their delivered-volume difference was 0.006% of the target, while the differences in specific energy and specific operating cost were below 0.03%. All 313 DMTS problems satisfied the prescribed convergence and interface-consistency criteria, with mean ID and RT solution times of 6.583 s and 0.461 s, respectively, under serial execution. In the constructed DA scalability cases, all 5–40-station problems converged, the objective gap remained below 0.012%, and the maximum interface mismatch remained below 1.0×10^{-4} m^3 s^{-1}. However, the serial ADMM runtime increased from 35.320 to 1596.805 s and remained longer than the centralized runtime. The scalability results therefore demonstrate retained numerical consistency, not measured parallel acceleration or a computational advantage over the centralized solver.

When the execution-error amplitude was doubled without parameter retuning, all temporal-layer problems remained convergent and the absolute delivery error was limited to 0.1012%. The head–efficiency sensitivity analysis showed that combined favorable and adverse perturbations changed the specific energy and operating cost by approximately 10–11%. These changes substantially exceed the numerical differences between CMTS and DMTS, confirming that the absolute energy and cost estimates depend strongly on the adopted equivalent coefficients. The ex-post reliability analysis produced a break-even shortage-evaluation coefficient of 0.229 CNY/m^3 under the tested tariff, delivery target, and disturbance realization.

The results support the use of the proposed coordination structure for larger serial cascades governed by the same aggregate convex model, but they do not establish applicability to arbitrary large-scale or branching hydraulic networks. The present formulation omits Saint-Venant channel dynamics, transport delay, water-level transients, variable lateral withdrawals, communication delay, individual pump curves, motor and converter dynamics, unit commitment, and switching decisions. Future work should integrate the supervisory scheduling framework with a calibrated dynamic hydraulic model and unit-level electromechanical models. The resulting formulation should consider operating-unit combinations, adjustable pump parameters, efficiency variation, and switching minimization. Validation using measured data, field-derived larger cascade systems, and an actual parallel computing architecture is required before operational deployment.

Author Contributions: Conceptualization, J.J. and Z.S.; methodology, J.J. and G.Z.; software, J.J. and Z.S.; validation, G.Z., Z.S. and J.J.; formal analysis, J.J. and Z.S.; investigation, Z.S. and J.J.; resources, G.Z. and J.J.; data curation, G.Z., J.J. and Z.S.; writing—original draft preparation, J.J. and Z.S.; writing—review and editing, J.J., Z.S. and G.Z.; visualization, J.J. and G.Z.; supervision, G.Z.; project

administration, G.Z.; funding acquisition, G.Z. All authors have read and agreed to the published version of the manuscript.

Funding: The authors declare that this study received funding from the China Civil Engineering Society Research Project (CCES2024KT10), the China State Construction Engineering Corporation Technology R&D Program (CSCEC-2024-Q-74), and the Chongqing Construction Science and Technology Program (CQCS2024-8-5). The funder was not involved in the study design, collection, analysis, interpretation of data, the writing of this article or the decision to submit it for publication.

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

Acknowledgments: The authors would like to express their gratitude to all those who helped them during the writing of this paper. The authors would like to thank the reviewers for their valuable comments and suggestions.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Liu, Y.; Zheng, H.; Wan, W.; Zhao, J. Optimal operation toward energy efficiency of the long-distance water transfer project. *J. Hydrol.* **2023**, *618*, 129152. [CrossRef]
2. Long, D.; Yang, W.; Scanlon, B.R.; Zhao, J.; Liu, D.; Burek, P.; Pan, Y.; You, L.; Wada, Y. South-to-North Water Diversion stabilizing Beijing's groundwater levels. *Nat. Commun.* **2020**, *11*, 3665. [CrossRef] [PubMed]
3. Feng, X.; Qiu, B.; Wang, Y. Optimizing parallel pumping station operations in an open-channel water transfer system using an efficient hybrid algorithm. *Energies* **2020**, *13*, 4626. [CrossRef]
4. Jowitt, P.W.; Germanopoulos, G. Optimal pump scheduling in water-supply networks. *J. Water Resour. Plan. Manag.* **1992**, *118*, 406–422. [CrossRef]
5. Schuurmans, J.; Clemmens, A.J.; Dijkstra, S.; Hof, A.; Brouwer, R. Modeling of irrigation and drainage canals for controller design. *J. Irrig. Drain. Eng.* **1999**, *125*, 338–344. [CrossRef]
6. Xia, H.; Wang, C.; Sun, J.; Lei, X.; Wang, H. Multi-scale closed-loop coupled real-time water quantity optimization scheduling of cascade pumping station in water supply canal systems. *J. Hydrol.* **2024**, *641*, 131802. [CrossRef]
7. Zhao, P.; Gu, C.; Cao, Z.; Ai, Q.; Xiang, Y.; Ding, T.; Lu, X.; Chen, X.; Li, S. Water-energy nexus management for power systems. *IEEE Trans. Power Syst.* **2021**, *36*, 2542–2554. [CrossRef]
8. Wang, C.; Gao, N.; Wang, J.; Jia, N.; Bi, T.; Martin, K. Robust operation of a water-energy nexus: A multi-energy perspective. *IEEE Trans. Sustain. Energy* **2020**, *11*, 2698–2712. [CrossRef]
9. Ulusoy, A.-J.; Nerantzis, D.; Stoianov, I. Adaptive MPC for burst incident management in water distribution networks. *IEEE Trans. Control Netw. Syst.* **2023**, *10*, 2060–2071. [CrossRef]
10. Shirsat, A.; Muthukaruppan, V.; Hu, R.; Paduani, V.D.; Xu, B.; Song, L.; Li, Y.; Lu, N.; Baran, M.; Lubkeman, D.; et al. A secure and adaptive hierarchical multi-timescale framework for resilient load restoration using a community microgrid. *IEEE Trans. Sustain. Energy* **2023**, *14*, 1057–1075. [CrossRef]
11. Oikonomou, K.; Parvania, M. Optimal coordinated operation of interdependent power and water distribution systems. *IEEE Trans. Smart Grid* **2020**, *11*, 4784–4794. [CrossRef]
12. Zamzam, A.S.; Dall'Anese, E.; Zhao, C.; Taylor, J.A.; Sidiropoulos, N.D. Optimal water–power flow-problem: Formulation and distributed optimal solution. *IEEE Trans. Control Netw. Syst.* **2019**, *6*, 37–47. [CrossRef]
13. Patari, N.; Venkataramanan, V.; Srivastava, A.; Molzahn, D.K.; Li, N.; Annaswamy, A. Distributed optimization in distribution systems: Use cases, limitations, and research needs. *IEEE Trans. Power Syst.* **2022**, *37*, 3469–3481. [CrossRef]
14. Kou, X.; Li, F.; Dong, J.; Starke, M.; Munk, J.; Xue, Y.; Olama, M.; Zandi, H. A scalable and distributed algorithm for managing residential demand response programs using alternating direction method of multipliers (ADMM). *IEEE Trans. Smart Grid* **2020**, *11*, 4871–4882. [CrossRef]
15. Cariaga, D.; Lorca, Á.; Anjos, M.F. A binary expansion approach for the water pump scheduling problem in large and high-altitude water supply systems. *Energies* **2024**, *17*, 4107. [CrossRef]
16. Tasseff, B.; Bent, R.; Coffrin, C.; Barrows, C.; Sigler, D.; Stickel, J.; Zamzam, A.S.; Liu, Y.; Van Hentenryck, P. Polyhedral relaxations for optimal pump scheduling of potable water distribution networks. *INFORMS J. Comput.* **2024**, *36*, 1040–1063. [CrossRef]
17. Stuhlmacher, A.; Mathieu, J.L. Chance-constrained water pumping to manage water and power demand uncertainty in distribution networks. *Proc. IEEE* **2020**, *108*, 1640–1655. [CrossRef]

18. Alhazmi, M.; Dehghanian, P.; Nazemi, M.; Oikonomou, K. Uncertainty-informed operation coordination in a water-energy nexus. *IEEE Trans. Ind. Inform.* **2023**, *19*, 6439–6449. [CrossRef]
19. Alhazmi, M.; Dehghanian, P.; Nazemi, M.; Wang, F.; Alfadda, A. Optimal operation of integrated water–power systems under contingencies. *IEEE Trans. Ind. Appl.* **2022**, *58*, 4350–4358. [CrossRef]
20. Ulusoy, A.-J.; Stoianov, I. Distributed solution of the day-ahead pump and valve scheduling problem for dynamically adaptive water distribution networks with storage. *Eur. J. Oper. Res.* **2025**, *323*, 267–275. [CrossRef]
21. Sun, F.; Ma, J.; Yu, M.; Wei, W. Optimized two-time scale robust dispatching method for the multi-terminal soft open point in unbalanced active distribution networks. *IEEE Trans. Sustain. Energy* **2021**, *12*, 587–598. [CrossRef]
22. Song, M.; Sun, W.; Shahidehpour, M.; Yan, M.; Gao, C. Multi-time scale coordinated control and scheduling of inverter-based TCLs with variable wind generation. *IEEE Trans. Sustain. Energy* **2021**, *12*, 46–57. [CrossRef]
23. Yi, Z.; Xu, Y.; Gu, W.; Wu, W. A multi-time-scale economic scheduling strategy for virtual power plant based on deferrable loads aggregation and disaggregation. *IEEE Trans. Sustain. Energy* **2020**, *11*, 1332–1346. [CrossRef]
24. Li, Q.; Zhou, Y.; Wei, F.; Li, S.; Wang, Z.; Li, J.; Zhou, G.; Liu, J.; Yan, P.; Yu, D. Multi-time scale scheduling for virtual power plants: Integrating the flexibility of power generation and multi-user loads while considering the capacity degradation of energy storage systems. *Appl. Energy* **2024**, *362*, 122980. [CrossRef]
25. Boyd, S.; Parikh, N.; Chu, E.; Peleato, B.; Eckstein, J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **2011**, *3*, 1–122. [CrossRef]
26. Zhu, Z.; Guan, G.; Wang, K. Distributed model predictive control based on the alternating direction method of multipliers for branching open canal irrigation systems. *Agric. Water Manag.* **2023**, *285*, 108372. [CrossRef]
27. Dunham, W.; Hencey, B.; Girard, A.R.; Kolmanovsky, I. Distributed model predictive control for more electric aircraft subsystems operating at multiple time scales. *IEEE Trans. Control Syst. Technol.* **2020**, *28*, 2177–2190. [CrossRef]
28. Jenks, B.; Ulusoy, A.-J.; Pecci, F.; Stoianov, I. Distributed nonconvex optimization for control of water networks with time-coupling constraints. *Water Resour. Manag.* **2025**, *39*, 523–546. [CrossRef] [PubMed]
29. Sun, K.; Sun, X.A. A two-level ADMM algorithm for AC OPF with global convergence guarantees. *IEEE Trans. Power Syst.* **2021**, *36*, 5271–5281. [CrossRef]
30. Mohammadi, A.; Kargarian, A. Learning-aided asynchronous ADMM for optimal power flow. *IEEE Trans. Power Syst.* **2022**, *37*, 1671–1681. [CrossRef]