

Article

A Maturity-Aware Proximal ADMM with NG-Route Relaxation for Dynamic Inventory Reallocation in a Multi-Echelon Mandarin Cold-Chain Network

Baowen Liang ¹, Linjie Ma ², Yiran Zhang ³, Yuxuan Su ¹, Haoyu Wang ² and Yiping Jiang ^{1,*}

¹ College of Information Management, Nanjing Agricultural University, Nanjing 210095, China

² College of Engineering, Nanjing Agricultural University, Nanjing 210095, China

³ College of Smart Agriculture, Nanjing Agricultural University, Nanjing 210095, China

* Correspondence: yipjiang@njau.edu.cn

Abstract

The Vehicle Routing Problem with Time Windows (VRPTW) takes on a structurally distinct form when the goods being routed undergo first-order quality decay during transport. In this setting, distance minimisation alone underestimates the true economic cost. A per-customer minimum-quality acceptance constraint further introduces a non-linear feasibility condition that does not appear in the classical formulation. This paper addresses such a setting in the context of loose-skin citrus fruit (e.g., mandarins) distribution, where stock has already undergone several days of cold storage at the origin warehouse, and remaining shelf life makes retail time windows binding rather than decorative. We formulate a *Maturity-Aware Multi-Echelon Dynamic Reallocation Vehicle Routing Problem with Time Windows* (MA-MEDR-VRPTW) on a three-echelon network (origin warehouse → distribution centres → stores) over a seven-day rolling horizon. A first contribution shows that the minimum-quality acceptance constraint admits an analytic transformation into a time-window tightening, which removes per-extension exponential evaluations from the subproblem solver. The algorithmic contribution is a proximal alternating direction method of multipliers (ADMM) with NG-route relaxation (PADMM-MA) whose quality-loss weight is updated by a residual-balancing rule and is decoupled from the outer reallocation linear program (LP) through approximate dynamic-programming-style marginal costs. On twelve Solomon-derived mandarin instances (72 feasible algorithm–instance combinations), PADMM-MA returns a mean seven-day cost of 12,638 CNY against 11,753 CNY for a subgradient baseline (+7.5%) at statistically indistinguishable arrival quality (paired Wilcoxon $p = 0.077$ for $\bar{q}_{arr}$), while cutting mean wall-clock time from 350 to 23 seconds (about $15\times$). The method, therefore, reads as a fast operational heuristic for daily re-planning. An ablation, an exact-MIP benchmark on tractable subproblems, and a scale extension to $n = 100$ customers round out the validation.



Academic Editor: Frank Werner

Received: 8 May 2026

Revised: 3 June 2026

Accepted: 29 June 2026

Published: 7 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

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

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: vehicle routing problem with time windows; alternating direction method of multipliers; NG-route relaxation; proximal ADMM; rolling-horizon optimisation; multi-echelon inventory; cold-chain logistics; perishable goods

MSC: 90C11, 90B06, 90C90

1. Introduction

Mandarin-type citrus fruits (Wenzhou mandarin, Gannan navel orange, Nanfeng mandarin, etc.) are a staple of the Chinese fresh-fruit retail sector and increasingly an export commodity. Unlike fresh-picked produce, a substantial share of in-market mandarin stock has already spent several days in cold storage at origin warehouses before being dispatched to regional distribution centres (DCs). The implications for distribution planning are twofold. First, the remaining shelf life is materially shorter than for freshly harvested stock, so the time windows imposed by retail partners become binding rather than decorative—this is in sharp contrast to standard academic VRPTW benchmarks such as the Solomon $R2xx/RC2xx$ series, whose wide time windows rarely activate in practice. Second, quality continues to decay during transport, so distance minimisation alone underestimates the true economic cost: late delivery of partly ripened fruit entails discounting or outright rejection.

The problem becomes harder when viewed through a multi-echelon, multi-period lens. An unbalanced overnight arrival pattern at the DCs triggers a need for lateral reallocation, but the reallocation itself consumes quality, and the downstream vehicle routing from each DC to its stores must then account for the transferred stock's aged condition. Existing work in cold-chain VRP typically addresses one of these three dimensions in isolation.

Contributions

This paper makes four contributions, each directly addressing a gap identified above.

1. **Unified decomposable model.** The MA-MEDR-VRPTW is cast as a decomposable multi-period mathematical program: an outer linear programme (LP) handles inventory reallocation, an inner maturity-aware VRPTW handles routing at each DC, and the two layers couple through approximate-dynamic-programming (ADP) style marginal costs rather than a Benders-style nesting (Section 3).
2. **Two constructive analytical results with explicit validity boundaries.** Proposition 1 shows that the per-store minimum-quality acceptance constraint under first-order decay kinetics is equivalent to a time-window tightening, which removes all per-extension exponential evaluations from the inner shortest-path problem with resource constraints (SPPRC). Proposition 2, valid under an explicit arc-wise-constant-temperature assumption (Assumption 2, with the boundary clarified in Remark 2), shows that cumulative quality loss is *redundant* as a fifth Pareto dimension, so the standard four-dimensional label suffices for the NG-route (a controlled relaxation of route elementarity) SPPRC. The explicit validity boundary makes the result auditable for practitioners whose fleets experience within-arc temperature variation.
3. **Residual-balancing adaptive penalty.** The proposed algorithm PADMM-MA is a proximal alternating direction method of multipliers (ADMM) with NG-route relaxation whose quality-loss weight β is updated by a residual-balancing rule rather than by subgradient ascent, which avoids the pricing oscillation that would otherwise arise from the non-linear quality penalty (Section 4). On twelve Solomon-derived test instances, this design yields an approximately $15\times$ mean wall-clock speed-up (23 s vs. 350 s) over a subgradient baseline at near-identical mean arrival quality, at a 7.5% mean-cost premium; this trade-off is discussed further in Section 5.2.
4. **Three-tier empirical validation.** The computational study has three complementary tiers: (i) an *optimality-gap study* against an exact mixed-integer program (MIP) on nine tractable single-DC subproblems (Appendix A, Table A1), with heuristic gaps ranging from 0.0% (two pairs, exact match) to +84.2% on tight small instances—locating the regime in which the heuristic leaves room against an exact solver; (ii) a *statistical-test study* on the twelve main-testbed instances (Section 5.2), reporting paired Wilcoxon

and t -tests and Cohen's d for cost, wall-clock time, and arrival quality; and (iii) a *scale-extension study* (Section 5.2) to $n = 100$ customers, which confirms that PADMM-MA remains feasible and that arrival quality is preserved at twice the main-testbed scale.

2. Literature Review

2.1. VRPTW Exact Methods

The state of the art for VRPTW relies on branch-price-and-cut (BPC) with *elementary* shortest-path pricing subproblems. Refs. Costa et al. [1] and Pessoa et al. [2] give complementary reviews of the BPC machinery and report best-known exact results on the Solomon benchmarks. Because exact elementarity is expensive, modern exact algorithms typically work in a relaxation: Baldacci et al. [3] introduced *NG-route relaxation*, a controlled subset of elementarity that has since become standard, and Roberti and Mingozzi [4] added dynamic enlargement of the NG-neighbourhoods. Recent exact algorithms for time-dependent variants extend this machinery in two complementary directions: Castellucci et al. [5] develops a branch-and-Benders-cut algorithm for the time-dependent VRP, while Luo et al. [6] presents an exact branch-price-and-cut algorithm for the time-dependent cold-chain pickup-and-delivery problem with incompatibility constraints. This paper uses NG-route relaxation within a heuristic pricing routine rather than an exact BPC scheme, and therefore does not attempt to close the optimality gap against these methods; the exact-MIP benchmark of Appendix A makes this positioning explicit.

2.2. ADMM in Combinatorial Optimisation

Boyd et al. [7] documented the classical convergence theory and adaptive penalty parameters for convex ADMM, and Zhang and Luo [8] proved convergence of a non-convex proximal ADMM under mild prox-regularity and descent conditions; the present analysis follows the latter framework (see Remark 3). The application of ADMM to vehicle routing is more recent. Yao et al. [9] introduced an ADMM-based decomposition for the VRPTW by splitting the visit-once and time-window constraints across vehicle subproblems and coordinating them through the classical Boyd-style dual update. Chi et al. [10] extended this line with an improved ADMM-based decomposition that partially retains coupling constraints to tighten the dual bound on transportation-network optimisation problems. Both works use an additive subgradient-like penalty update and target single-echelon, quality-agnostic problems. More recently, Xia et al. [11] applied an extended ADMM to electric VRPTW under travel-time uncertainty, again preserving the additive penalty rule. PADMM-MA (Section 4) departs from this line on three points: (i) we relax the visit-once constraint while keeping the time-window and capacity constraints inside an NG-route pricing subproblem, rather than splitting time windows across vehicles as in Yao et al. [9] and Chi et al. [10]; (ii) we replace the additive quality-penalty update with a residual-balancing, multiplicative rule (16), which is motivated by the non-linear structure of the quality-loss term and avoids the oscillation observed empirically under additive updates on our testbed (Section 5.3); and (iii) the outer problem here is a multi-echelon inventory-reallocation LP rather than a single-echelon routing relaxation, so the dual is coupled to the inner pricing through ADP-style marginal costs (Section 4.3) rather than through a shared set of routing variables.

2.3. Cold-Chain and Quality-Aware Routing

The cold-chain VRP literature has grown rapidly in the past decade. Accorsi et al. [12] review cold-chain logistics optimisation with a focus on climate-driven decision support, and the more recent bibliometric review of Qi and Li [13] charts the evolution of the cold-chain VRP from 2008 to 2024, identifying temperature-dependent decay, time-window

integration, and multi-objective formulations as the three dominant research streams. Within the methodological tradition, Osvald and Stirn [14] added time-dependent freshness decay to the VRP, Lado et al. [15] provides the post-harvest kinetic parameters that underpin the Arrhenius model used here, and Tietel et al. [16] documents the importance of storage temperatures in maintaining flavour and quality of mandarins. Recent work in 2024–2025 has extended the cold-chain VRP along three complementary axes. Lin et al. [17] introduces a vehicle routing model with *dynamic* temperature-dependent quality decay for post-harvest precooling, the work closest in spirit to the present paper at the routing level. Liu et al. [18] couples cold-chain routing with a freshness-keeping cost term and a low-carbon objective and reports that under appropriate calibration, freshness investment lowers total cost rather than raising it—an effect consistent with the cost-hierarchy regime identified in our outer LP. On the carbon and network side, Hou et al. [19] optimises urban cold-chain routes under carbon constraints in a time-varying network, and Hou et al. [20] extends this to a multi-depot semi-open configuration. At the supply-chain level, Beshara et al. [21] provides a comprehensive review of citrus supply chains in the EU, the Middle East and Africa, identifying cold-chain management as the dominant research theme on the African continent. Compared with this stream, the present paper embeds the quality-decay kinetics as a closed-form cost term that admits the time-window-tightening reduction of Proposition 1; this turns a non-linear acceptance constraint into a linear TW update and removes all per-extension exponential evaluations from the pricing subproblem.

2.4. Multi-Echelon Reallocation

The inventory-routing problem with lateral transshipment [22] is the closest structural analogue. Most prior work in this stream is either static (a single-period allocation-plus-routing model) or adopts full two-stage or multi-stage stochastic programming with Benders-style nesting. Recent work has begun to integrate perishability into this picture: Chen et al. [23] formulates a two-echelon location-routing model for perishable products and proposes an adaptive large neighbourhood search heuristic on Solomon-derived instances but treats perishability as a cost penalty rather than deriving it from underlying first-order kinetics. The approach taken here is instead an approximate dynamic-programming one: the outer reallocation LP and the inner maturity-aware VRPTW at each DC are coupled via per-kg transport and quality-loss marginals that are refreshed across the rolling horizon (Section 4.3). This sidesteps the computational burden of double nesting and still transmits value-function information from the inner routing outcome to the next-day outer allocation decision.

2.5. Positioning of PADMM-MA

Three features distinguish the present framework from prior ADMM-based routing work. (i) *Penalty update*. Earlier ADMM-based VRPTW decompositions [9–11] coordinate their subproblems with an additive, subgradient-like penalty update. PADMM-MA instead updates the quality-loss weight by a multiplicative residual-balancing rule (16), which is motivated by the non-linear quality-loss term and avoids the pricing oscillation that additive updates exhibit on our testbed (Section 5.3). (ii) *Decomposition*. Those studies split the visit-once and time-window constraints across single-echelon vehicle subproblems. Here, by contrast, the visit-once constraint is relaxed while the time-window and capacity constraints are kept within an NG-route pricing subproblem, and the outer problem is a multi-echelon inventory-reallocation LP coupled to the inner pricing via ADP-style marginal costs rather than shared routing variables. (iii) *Novel versus adapted components*. The NG-route SPPRC machinery [3] and the non-convex proximal-ADMM convergence framework [8] are adapted from existing methods; what is new here is the analytic time-

window-tightening reduction (Proposition 1), the four-dimensional label-sufficiency result (Proposition 2), the residual-balancing quality-weight update, and the ADP-coupled multi-echelon decomposition. Table 1 summarises these distinctions against the most closely related studies.

Table 1. Positioning of the proposed framework against the most closely related studies. “Inv.–rout.”: integrated inventory–routing decisions; “–”: not modelled.

Study	Perish.	Multi-Ech.	Inv.–Rout.	Quality-Decay Model	Solution Method
Yao et al. [9]	No	No	No	–	ADMM, additive penalty
Chi et al. [10]	No	No	No	–	Improved ADMM, additive penalty
Xia et al. [11]	No	No	No	–	Extended ADMM, additive penalty
Chen et al. [23]	Yes	Yes	Yes	Cost penalty	ALNS metaheuristic
Lin et al. [17]	Yes	No	No	Dynamic temp.-dependent	Metaheuristic
Liu et al. [18]	Yes	No	No	Freshness-keeping cost	Low-carbon heuristic
This work	Yes	Yes	Yes	First-order Arrhenius (closed form)	Proximal ADMM + NG-route, residual-balancing

3. Mathematical Model

Before stating the model, we clarify how its two parts fit together, as this governs what is finally solved. The framework is a two-level decomposition over the rolling horizon. At the upper level, a daily inventory-reallocation linear program (12)–(15) decides how much stock each distribution centre (DC) receives from the origin and transfers laterally to its neighbours; its output is the quality and quantity of stock available at each DC at the start of the day. At the lower level, taking that stock as given, a separate maturity-aware vehicle routing problem (3)–(8) is solved for each DC to route its fixed store cluster. The two levels are not solved jointly as one monolithic program; they are coupled sequentially and across days through marginal costs (per-kilogram transport and quality-loss estimates) that the lower-level routing solutions feed back into the next day’s upper-level program (Section 4.3). The objects actually solved are therefore one small LP and D routing subproblems per day, repeated over the horizon; the routing subproblem (3)–(8) is the computational bottleneck and the focus of the algorithm in Section 4.

3.1. Network and Timeline

The system is a three-echelon network $(\{o\}, \mathcal{D}, \mathcal{S})$ consisting of one origin warehouse o , $D = |\mathcal{D}|$ regional distribution centres $d \in \mathcal{D}$, and $|\mathcal{S}|$ stores partitioned into clusters $\mathcal{S}_d \subset \mathcal{S}$ served exclusively by DC d . The planning horizon is H days indexed by $t = 1, \dots, H$ (in this paper, $H = 7$).

Standing Assumptions on the Network and the Data

Three standing assumptions hold throughout. (i) *Complete graph with symmetric arcs.* Each DC and its assigned stores form a complete directed graph; arc costs and travel times are taken to be symmetric, $c_{ij} = c_{ji}$ and $\tau_{ij} = \tau_{ji}$, consistent with the Solomon benchmark convention. (ii) *Deterministic data.* All store demands d_j^t , service times s_i , and travel times τ_{ij} are deterministic and known at the start of each planning day; a stochastic (chance-constrained) extension is left to future work (Section 7). (iii) *Per-DC clusters fixed.* The store partition $\{\mathcal{S}_d\}_{d \in \mathcal{D}}$ is fixed for the entire horizon; the routing-level decisions concern only the daily route structure within each cluster, not the assignment of stores to DCs.

Figure 1 illustrates the three-echelon structure: a single-origin warehouse replenishes the distribution centres, each DC routes to its own fixed store cluster, and lateral transshipment links adjacent DCs.




Figure 1. Three-echelon cold-chain network $(\{o\}, \mathcal{D}, \mathcal{S})$: one origin warehouse, $D = 3$ distribution centres (DCs) each serving a fixed store cluster, with lateral DC-to-DC transshipment.

3.2. Quality-Decay Kinetics

For each store $j \in \mathcal{S}$ product quality follows the first-order law

$$\frac{dq_j}{dt} = -k(T) q_j, \quad q_j(t_0) = q_{j,0}, \quad (1)$$

with Arrhenius rate constant $k(T) = k_{\text{ref}} \exp\left[-\frac{E_a}{R} \left(\frac{1}{T} - \frac{1}{T_{\text{ref}}}\right)\right]$, giving the closed form $q_j(t) = q_{j,0} \exp(-k(T)(t - t_0))$ under a piecewise-constant transport temperature. Here R is the universal gas constant, E_a the activation energy, T_{ref} a reference temperature, k_{ref} the rate constant at T_{ref} , and T the absolute transport temperature in Kelvin (Celsius values are converted by $T = T_{\text{°C}} + 273.15$ before substitution). The first-order decay law (1) with an Arrhenius temperature dependence is not proposed here; it is the standard model in post-harvest quality kinetics, and we adopt the formulation and parameter ranges used in that literature [14–16]. Cold storage at a DC follows the same law with a separate rate k_{storage} . The numerical calibration of these constants, the reference and transport temperatures, and the cold-storage rate is deferred to the experimental setup (Section 5.1).

Each store imposes a minimum acceptance threshold $q_j^{\min}$. Let π_j be the unit value (CNY/kg) of the fruit at store j and d_j^t its demand on day t . Combining the closed-form decay of (1) with this unit value, we define the *unit* quality loss at arrival time t_j^{arr} as

$$L_j(t_j^{\text{arr}}) = \pi_j d_j^t q_{j,0} (1 - \exp\{-k(t_j^{\text{arr}} - t_0)\}). \quad (2)$$

In contrast to (1), Equation (2) is a modelling construction of the present paper: it monetises the standard decay kinetics into a per-customer cost that enters the routing objective directly and, as shown below, admits the time-window reduction of Proposition 1.

Remark 1 (Domain convention for L_j). Equation (2) is defined on the physically meaningful domain $t_j^{\text{arr}} \geq t_0$, where t_0 is the DC dispatch (hand-over) epoch of the vehicle serving customer j and t_j^{arr} is the vehicle's arrival instant at j on the same trip. On this domain, $1 - \exp\{-k(t_j^{\text{arr}} - t_0)\} \in [0, 1]$ and L_j is non-negative and non-decreasing in t_j^{arr} . The constraint $t_j^{\text{arr}} \geq t_0$ is enforced structurally by the service-time flow constraint (6): any vehicle arrives at j strictly after leaving the DC. As an implementation safeguard against spurious negative values from numerical round-off near $t_j^{\text{arr}} \approx t_0$, we treat $L_j(t) := \max\{0, \pi_j d_j^t q_{j,0} (1 - \exp\{-k(t - t_0)\})\}$ in code; on all twelve test instances, the $\max$ operator is never active, confirming the structural enforcement. This clamp is an implementation detail of the numerical solver rather than a modelling assumption: the analytical model and all theoretical results in this paper concern the unclamped form (2) on its physical domain.

Assumption 1 (Quality inspection at arrival). *Product quality at store j is evaluated at the vehicle arrival instant t_j^{arr} rather than at the service-start instant $\max(t_j^{\text{arr}}, a_j)$, where a_j is the ready time. This reflects industrial practice: refrigerated vehicles maintain controlled temperature during on-site waiting, so quality at handover equals quality at arrival.*

3.3. Inner Per-DC VRPTW (Maturity-Aware)

We assume a *homogeneous fleet*: all vehicles share a common capacity W , travel speed, and refrigeration profile, so that arc travel times τ_{ij} and arc costs c_{ij} do not depend on the vehicle index. For each DC d on day t , let K_d be the number of available vehicles, each of capacity W .

The problem has two groups of decision variables: the binary arc-usage variables $x_{ij}^k \in \{0, 1\}$, equal to 1 if vehicle k traverses arc (i, j) and 0 otherwise, and the continuous service-start times $S_j^k \geq 0$ of vehicle k at node j . The arrival time $t_j^{\text{arr},k}$ is not an independent variable: it is determined by the service-start times through the propagation constraint (6), and under Assumption 1, it coincides with the service start at the served node. The maturity-aware inner VRPTW at DC d on day t reads

$$\min \sum_k \sum_{(i,j)} c_{ij} x_{ij}^k + \sum_k \sum_{j \in \mathcal{S}_d} L_j(t_j^{\text{arr},k}) \quad (3)$$

$$\text{s.t. } \sum_k \sum_i x_{ij}^k = 1 \quad \forall j \in \mathcal{S}_d \quad (4)$$

$$\sum_j d_j \sum_i x_{ij}^k \leq W \quad \forall k \quad (5)$$

$$S_j^k \geq S_i^k + s_i + \tau_{ij} - M(1 - x_{ij}^k) \quad \forall (i, j), k \quad (6)$$

$$a_j \leq S_j^k \leq b_j \quad \forall j, k \quad (7)$$

$$q_j(t_j^{\text{arr},k}) \geq q_j^{\min} \quad \forall j, k \quad (8)$$

The objective (3) has two parts, both functions of the decision variables. The first term is the classical total monetised distance cost of the routes selected by x . The second term is the total maturity cost: through the arrival time $t_j^{\text{arr},k}$, which is itself a function of the service-start variables S_j^k , the per-customer loss $L_j(\cdot)$ of (2) depends on how late each vehicle reaches its customers, so later arrivals are penalised. The constraints read as follows. Constraint (4) requires that every store be visited exactly once (the visit-once constraint); (5) enforces the vehicle capacity W ; (6) propagates service-start times along each route through the big- M flow inequality; (7) restricts each service start to the store time window (TW) $[a_j, b_j]$; and (8) is the per-store minimum-quality acceptance constraint. The routing core (4)–(7) is the classical VRPTW [24]; what is proposed here is the maturity cost term in (3) together with the acceptance constraint (8). Here c_{ij} is the (monetised) distance cost, τ_{ij} the travel time, and s_i the service time. The big- M constant in (6) is set to $M = b_{\max} + s_{\max} + \tau_{\max}$, where $b_{\max} = \max_{j \in \mathcal{S}_d} b_j$, $s_{\max} = \max_i s_i$, and $\tau_{\max} = \max_{(i,j)} \tau_{ij}$; this is the tightest instance-wide value that keeps the inequality inactive whenever $x_{ij}^k = 0$ and strictly enforces the service-time propagation whenever $x_{ij}^k = 1$, avoiding both the spurious infeasibility of an under-sized M and the numerical conditioning issues of an over-sized one. On the Solomon-derived testbed used in this paper, M never exceeds a few multiples of the depot horizon.

Constraint (8) makes the formulation non-linear through the exponential in $q_j(\cdot)$. The following reduction makes it algorithmically tractable.

Proposition 1 (Constraint reduction via TW tightening). *Under Assumption 1 and the decay law (1), constraint (8) is equivalent to the linear tightening*

$$b_j \leftarrow \min \left\{ b_j, t_0 + \frac{1}{k(T)} \ln \frac{q_{j,0}}{q_j^{\min}} \right\}. \quad (9)$$

Proof. Substituting the closed form into (8): $q_{j,0} \exp\{-k(T)(t_j^{\text{arr}} - t_0)\} \geq q_j^{\min}$. Taking logs and rearranging yields $t_j^{\text{arr}} \leq t_0 + \frac{1}{k(T)} \ln(q_{j,0}/q_j^{\min})$. Combining with the existing $t_j^{\text{arr}} \leq b_j$ proves the claim. $\square$ $\square$

Assumption 2 (Arc-wise constant temperature). *Along every transport arc (i, j) the vehicle's internal temperature is piecewise constant at some T_{ij} that is known at planning time; along every storage arc at a DC, the temperature is piecewise constant at T_{storage} . Consequently the Arrhenius rate $k(T)$ and hence the quality-decay law $q(t) = q_0 \exp(-k(T)(t - t_0))$ is strictly monotone decreasing in elapsed time for each arc.*

Assumption 3 (Node-intrinsic decay parameters). *For every node j served from DC d on a given day t , the quality-decay parameters $(q_{j,0}, t_{j,0}, \pi_j, d_j^t)$ —initial quality at the DC hand-over instant, hand-over epoch, unit value, and daily demand—are fixed node-level attributes known at the start of day t . They do not depend on the particular partial path by which the subproblem solver reaches j , nor on which vehicle serves j . In particular, the mass-weighted quality blend $\bar{q}_d^t$ produced by the outer LP (Section 3.5) acts as the common $q_{j,0}$ for every node in S_d on day t , and $t_{j,0}$ equals the DC's dispatch epoch for that day.*

Proposition 2 (4-D dominance sufficiency). *Under Assumptions 2 and 3, and the quality-decay law in (1), the NG-route SPPRC used inside PADMM-MA is sound and complete when labels carry the four-dimensional resource vector $(c^{\text{lag}}, w, \tau, \mathcal{N})$ (Lagrangian cost, load, time, NG-set). No separate cumulative quality-loss resource is required.*

Proof. Throughout this proof we fix a single planning day t ; to lighten notation we write $d_{v_k}^t$ for $d_{v_k}^t$ when no ambiguity arises. Fix a node j on a feasible partial path and consider two labels L_1 and L_2 at j with resource vectors $(c_1, w_1, \tau_1, \mathcal{N}_1)$ and $(c_2, w_2, \tau_2, \mathcal{N}_2)$ respectively. Assume standard dominance

$$c_1 \leq c_2, \quad w_1 \leq w_2, \quad \tau_1 \leq \tau_2, \quad \mathcal{N}_1 \subseteq \mathcal{N}_2,$$

with at least one inequality strict.

Let $\sigma = (j = v_0, v_1, \dots, v_m)$ be any feasible extension from j . For each consecutive pair (v_{k-1}, v_k) with travel time $\tau_{v_{k-1}v_k}$, define the cumulative arrival time starting from label L_p ($p \in \{1, 2\}$) by

$$\tau_p^{(k)} = \max(\tau_p^{(k-1)}, a_{v_{k-1}}) + s_{v_{k-1}} + \tau_{v_{k-1}v_k}, \quad \tau_p^{(0)} = \tau_p,$$

where a_v and s_v are the node ready time and service time. A straightforward induction on k using $\tau_1 \leq \tau_2$ and the monotonicity of the max-plus recursion yields

$$\tau_1^{(k)} \leq \tau_2^{(k)} \quad \text{for all } k = 0, 1, \dots, m. \quad (10)$$

By Assumption 2, the quality loss incurred at node v_k on the extension is

$$\ell_{v_k}(\tau) = \pi_{v_k} d_{v_k} q_{v_k,0} [1 - \exp(-k(T_{v_{k-1}v_k})(\tau - t_{v_k,0}))],$$

which is *non-decreasing* in τ for every $k \in \{1, \dots, m\}$ because $-\exp(-k(T) \cdot)$ is non-decreasing. Combined with (10),

$$\ell_{v_k}(\tau_1^{(k)}) \leq \ell_{v_k}(\tau_2^{(k)}) \quad (k = 1, \dots, m). \quad (11)$$

Summing (11) over k gives the cumulative quality-loss ordering $L(\sigma \mid L_1) \leq L(\sigma \mid L_2)$.

For the Lagrangian objective of the SPPRC subproblem, $\text{cost}(\sigma \mid L_p) = c_p + \Delta c(\sigma) + \beta L(\sigma \mid L_p)$, where $\Delta c(\sigma)$ is the extension's lag-cost increment, independent of p . Hence $\text{cost}(\sigma \mid L_1) \leq \text{cost}(\sigma \mid L_2)$ for every feasible σ from j , so any completion of L_2 is matched or beaten by the corresponding completion of L_1 . Feasibility of extensions is preserved because $\tau_1 \leq \tau_2$, $w_1 \leq w_2$, and $\mathcal{N}_1 \subseteq \mathcal{N}_2$. Therefore L_1 dominates L_2 in the full extended-path objective, and no fifth-quality-loss resource is required to detect the dominance. $\square$

Remark 2 (Validity boundary of Proposition 2). *Assumption 2 is essential. If temperature varies within a single arc, the rate $k(T)$ becomes a function of sub-arc position and the quality loss incurred by an extension is no longer uniquely determined by the arrival time τ at the starting node; a fifth label component tracking integrated thermal exposure $\int k(T(\xi)) d\xi$ would then be required. The present analysis, and the accompanying implementation, assumes the arc-wise constant-temperature setting, which is standard for the reefer-truck operation modelled here.*

Directionality of the NG-set relation. *The dominance statement of Proposition 2 uses $N_1 \subseteq N_2$ and not the reverse: under this inclusion, the feasibility set of extensions from L_1 contains the feasibility set of extensions from L_2 , because any forbidden set encoded in N_2 is already encoded in N_1 . The symbol $\subseteq$ on the NG-set component therefore plays the same “weaker resource implies wider extension set” role as $\leq$ on cost, load, and time, and the proof's feasibility-preservation step is in the direction $L_1 \Rightarrow L_2$, not the reverse.*

3.4. Outer Reallocation LP

For each day t , given on-hand inventories I_d^{t-1} and day- t demands $D_d^t = \sum_{j \in \mathcal{S}_d} d_j^t$, the outer LP decides: x_d^t (origin $\rightarrow$ DC d replenishment), $y_{d_1 d_2}^t$ (lateral DC $\rightarrow$ DC transshipment), and u_d^t (slack for stock-out). Let z_d^t be an auxiliary variable for holding overshoot. The LP reads

$$\min \sum_d c_d^o x_d^t + \sum_{d_1 \neq d_2} c_{d_1 d_2}^\ell y_{d_1 d_2}^t + \sum_d h_d z_d^t + \sum_d M_d^t u_d^t \quad (12)$$

$$\text{s.t. } I_d^t = I_d^{t-1} + x_d^t + \sum_{e \neq d} y_{ed}^t - \sum_{e \neq d} y_{de}^t - D_d^t + u_d^t, \quad \forall d \quad (13)$$

$$0 \leq I_d^t \leq I_d^{\max}, \quad z_d^t \geq I_d^t - I_d^{\text{target}}, \quad \forall d \quad (14)$$

$$x, y, u, z \geq 0. \quad (15)$$

The *penalty tier* $M_d^t = \kappa (\bar{c}_{t-1}^{\text{trans}} + \bar{c}_{t-1}^{\text{qloss}})$ dominates all in-network costs, so stock-outs are a last resort. We adopt a default multiplier $\kappa = 10$, chosen so that M_d^t exceeds the in-network unit costs $c_d^o, c_{d_1 d_2}^\ell, h_d$ by two to three orders of magnitude on every day of the testbed; the sensitivity sweep $\kappa \in [3, 20]$ reported in Section 5.4 confirms that the operating regime is insensitive to this value once κ is large enough to make the stock-out tier dominant. The marginals $\bar{c}_{t-1}^{\text{trans}}$ and $\bar{c}_{t-1}^{\text{qloss}}$ are the day- $(t-1)$ per-kg transport and quality-loss unit costs averaged over the inner solves; this closes the feedback loop between the outer and inner problems as an approximate dynamic programming value-function update [25].

Cost-Hierarchy Convention

The unit costs in (12) satisfy the strict ordering $h_d < c_d^o < c_{d_1 d_2}^\ell \ll M_d^t$ on every day t . The gap $c_{d_1 d_2}^\ell > c_d^o$ is structural rather than numerical: a lateral transfer physically embeds an origin-replenishment movement (stock originates from o and must first reach

the sending DC d_1) plus an additional inter-DC leg and the handover latency modelled in Section 3.5. The holding cost h_d is bounded above by the daily marginal transport cost so that the LP does not suppress legitimate replenishment, and the stockout tier M_d^t is set by the data-driven update of the rolling-horizon driver (Step 1 of Section 4.3) two to three orders of magnitude above $c_{d_1 d_2}^\ell$, so that stock-outs are used only as a last resort. Under this ordering the LP uses origin replenishment as the default supply channel, lateral transfer only when an inventory imbalance between DCs would otherwise force a stock-out or a capacity breach, and u_d^t only when no feasible combination of the first two exists; no lateral movement is ever strictly cheaper than an equivalent origin replenishment, and the asymmetric-demand experiments of Section 5.3 exercise the regime in which the second channel becomes binding. Table 2 summarises the principal notation used throughout the paper.

Table 2. Principal notation used throughout the paper.

Symbol	Definition
<i>Sets and indices</i>	
o	Origin warehouse
$\mathcal{D}, d$	Set of distribution centres (DCs); index
$D = \mathcal{D} $	Number of distribution centres
$\mathcal{S}, \mathcal{S}_d, j$	Set of all stores; stores served by DC d ; store index
H, t	Planning horizon (days); day index
<i>Quality-decay parameters</i>	
$q_{j,0}$	Initial quality of fruit at handover from DC to vehicle
$q_j^{\min}$	Minimum acceptance quality threshold at store j
$k(T)$	Arrhenius decay rate at temperature T (Kelvin); see Equation (1)
R	Universal gas constant ($8.314 \text{ J mol}^{-1} \text{ K}^{-1}$)
E_a	Activation energy in Arrhenius rate (50 kJ/mol)
k_{storage}	Cold-storage decay rate at DC (literature-calibrated)
$\bar{q}_d^t$	Mass-weighted quality blend at DC d on day t
$t_{j,0}$	DC dispatch epoch (hand-over time) for store j on day t
<i>Inner VRPTW (per DC d, day t)—decision variables</i>	
x_{ij}^k	Binary: 1 if vehicle k traverses arc (i, j) , 0 otherwise
S_j^k	Service-start time of vehicle k at node j
$V_d^{(k)}$	Number of vehicles dispatched from DC d at PADMM-MA iteration k
<i>Inner VRPTW—parameters</i>	
K_d	Maximum number of available vehicles at DC d
W	Vehicle capacity (kg); homogeneous fleet
c_{ij}, τ_{ij}, s_i	Monetised arc cost; travel time; service time
$[a_j, b_j]$	Time window at store j (tightened by Proposition 1)
$L_j(t_j^{\text{arr}})$	Quality-loss cost at store j given arrival time t_j^{arr} (Equation (2))
M	Big-M constant in flow constraint (6); $M = b_{\max} + s_{\max} + \tau_{\max}$
<i>Outer reallocation LP (day t)—decision variables</i>	
$x_d^t, y_{d_1 d_2}^t, u_d^t$	Origin replenishment, lateral transfer, stock-out slack
z_d^t, I_d^t	Holding-overshoot auxiliary variable; resulting inventory level at DC d
<i>Outer reallocation LP—parameters</i>	
$I_d^{\max}, I_d^{\text{target}}$	Inventory capacity and target at DC d ($I_d^{\max} = 5.0 \bar{D}_d, I_d^{\text{target}} = 1.5 \bar{D}_d$)
$\bar{D}_d$	Average daily demand at DC $d, \bar{D}_d = H^{-1} \sum_t D_d^t$
$c_d^o, c_{d_1 d_2}^\ell, h_d$	Origin-replenishment, lateral-transfer, holding unit costs
M_d^t	Stock-out penalty tier, $M_d^t = \kappa(\bar{c}_{t-1}^{\text{trans}} + \bar{c}_{t-1}^{\text{loss}})$
κ	Stock-out penalty multiplier
<i>PADMM-MA algorithm</i>	
λ^k	Lagrangian multipliers for visit-once constraints (Equation (4))
ρ^k	Augmented-Lagrangian penalty parameter
β^k	Quality-loss weight; updated by residual-balancing rule (Equation (16))
Υ^k	Proximal reference (visit-incidence vector) extracted from routes $\mathcal{R}^k$
α_{prox}	Proximal-anchor weight
$c_{\text{lag}}^{\text{tag}}$	Lagrangian-reduced cost component of an SPPRC label
$\mathcal{N}$	NG-neighbourhood set carried in an SPPRC label
$r^{(k)}$	Quality-loss share of current primal upper bound

3.5. Quality Blending at DCs

After the outer LP executes day t 's decisions, each DC's on-hand quality $\bar{q}_d^t$ is updated by the mass-weighted blend of (a) yesterday's residual stock (aged one day at $T_{\text{storage}} = 5^\circ\text{C}$), (b) freshly received origin stock ($q = 0.95$), and (c) incoming laterally transferred stock (source DC's quality further decayed by a 2 h handover latency). This blended $\bar{q}_d^t$ then overrides the default $q_{j,0}$ in the inner maturity-aware VRPTW at DC d .

3.5.1. Handover-Latency Breakdown

The 2 h figure reflects a composite of four standard handling steps observed in regional-DC lateral-transshipment practice: Unloading at the receiving dock (~ 30 min), inbound quality inspection and palletising (~ 45 min), sortation and re-labelling for downstream routes (~ 30 min), and reloading onto the outbound vehicle (~ 15 min). The 2 h figure is an operational parameter of the network model rather than a literature-derived measurement. Because the Arrhenius-rate quality decay at $T_{\text{transport}} = 10^\circ\text{C}$ produces a per-hour quality drop of roughly $1 - \exp(-k \cdot 1\text{ h})$ — well below 1% under the calibrated k_{ref} — doubling the latency to 4 h changes post-transfer quality by less than one percentage point on every lateral movement, and halving it to 1 h yields a symmetric bound; the main-testbed conclusions are therefore insensitive to this parameter within its plausible operational range.

The initial DC quality at the start of the planning horizon depends on the pre-horizon storage age τ_{storage} (default: 4 days):

$$\bar{q}_d^0 = 0.95 \exp(-k_{\text{storage}} \cdot \tau_{\text{storage}}).$$

3.5.2. Provenance of the 4-Day Default

The 4-day value is an operational-planning parameter rather than a literature-derived constant: it reflects the typical elapsed time between origin-warehouse receipt and the start of the seven-day distribution horizon for already-cold-stored mandarin stock, as opposed to freshly harvested stock with $\tau_{\text{storage}} = 0$. The 4-day value is a modelling assumption, and its impact is tested directly in Section 5.4. The sweep there varies $\tau_{\text{storage}} \in \{2, 4, 6, 8\}$ days on the four representative instances. Mean arrival quality $\bar{q}_{\text{arr}}$ decreases overall from 0.865 to 0.848 over this range. The decrease is non-monotone, with a local dip near $\tau_{\text{storage}} \approx 6$ days that reflects the interaction between storage-age decay and the Proposition 1-induced TW tightening; total cost likewise stays non-monotone and within a $\sim 10\%$ band. The main-testbed conclusions of Section 5.2 — the $15\times$ speed advantage of PADMM-MA and the 7.5% cost advantage of SUBGRAD-MA — are both order-of-magnitude effects and are therefore robust to any τ_{storage} in the plausible 2–8 day operational range.

4. Solution Methodology

Figure 2 gives an overview of how the components introduced in Section 3 interact over one planning day and across the rolling horizon. Each day, the outer LP reallocates inventory, quality blending updates the on-hand stock quality at each DC, PADMM-MA routes each DC's cluster, and the resulting per-kilogram marginal costs feed back into the next day's LP. The remainder of this section details the inner PADMM-MA routine (Section 4.1), the adaptive β update (Section 4.2), and the rolling-horizon driver (Section 4.3).




Figure 2. Workflow of the proposed framework over one planning day and the rolling-horizon feedback (Section 4.3). The dashed return path carries the day- t per-kilogram transport and quality-loss marginals into the day- $(t + 1)$ stock-out penalty tier of the outer LP, closing the approximate-dynamic-programming (ADP) loop.

4.1. PADMM-MA at Each DC

Lagrangian relaxation of the visit-once constraint (4) with multipliers $\lambda \in \mathbb{R}^{|S_d|}$, combined with an augmented proximal penalty, gives the PADMM scheme whose per-iteration structure is summarised in Algorithm 1. The inner SPPRC pricing solver, referred to as FDP-TW-MATURITY, is a forward dynamic-programming NG-route labelling routine over the four-dimensional resource vector of Proposition 2, with time-window feasibility checked via the Proposition 1 tightening and the quality-loss term $\beta L_j(t_j^{\text{arr}})$ folded additively into the reduced cost.

Algorithm 1 PADMM-MA iteration k at DC d .

- 1: Solve $V_d^{(k)}$ successive SPPRC subproblems using FDP-TW-MATURITY with Lagrangian profit $\lambda^k + \alpha_{\text{prox}} Y^k$ and quality weight β^k , subject to the fleet-capacity bound $V_d^{(k)} \leq K_d$. Each subproblem corresponds to pricing a single vehicle's route, and the effective number $V_d^{(k)}$ is the smallest value such that the union of priced routes covers every visit-once constraint (line 4) at least once. Obtain routes $\mathcal{R}^k$.
 - 2: Compute Lagrangian LB and the visit count $n_j = \sum_r [j \in r]$.
 - 3: Heuristic repair + VND local search on $\mathcal{R}^k$ giving the primal UB (C^k, L^k) .
 - 4: Update dual: $\lambda^{k+1} = \lambda^k + \rho^k(1 - n)$.
 - 5: Update proximal: Y^{k+1} extracted from $\mathcal{R}^k$ as the visit-incidence vector $Y_j^{k+1} = \mathbb{1}[j \in \mathcal{R}^k]$, serving as a primal anchor for the proximal regularisation in Step 1.
 - 6: Update ρ^{k+1} via primal/dual residual balancing.
 - 7: Update β^{k+1} via quality-loss-share residual balancing (Section 4.2).
 - 8: **Termination test:** stop if any of (i) iteration count $k + 1 \geq K_{\text{max}}$, (ii) primal-dual residual $\|\mathbf{1} - n\|_{\infty} \leq \varepsilon_{\text{prim}}$ jointly with relative UB change $|C^k - C^{k-1}| / \max(C^{k-1}, 1) \leq \varepsilon_{\text{rel}}$, or (iii) the primal UB has not improved for K_{stall} consecutive iterations; otherwise increment $k \leftarrow k + 1$ and return to Step 1.
-

The three termination criteria are checked in order at each iteration and are disjunctive: any one of them suffices to stop the loop. The termination parameters (K_{max} , $\varepsilon_{\text{prim}}$, ε_{rel} , K_{stall}) and the proximal-anchor weight α_{prox} are applied uniformly to all twelve main-testbed instances; their numerical values are listed in the experimental setup (Section 5.1). Under those settings, the empirical iteration counts of Figure 3 were obtained (e.g., RC201 terminating via criterion (ii), and R101 exhausting the iteration cap without triggering (ii) or (iii)). The returned incumbent is the primal UB with the lowest value of C^k encountered over the whole trajectory, which guarantees that a late-iteration degradation in the monitored dual score cannot worsen the reported solution.

The heuristic repair (Step 3) builds a feasible solution from possibly overlapping NG-route fragments through a four-phase procedure (initialisation, intra- and inter-route deduplication, and regret-2 insertion), followed by a variable-neighbourhood-descent (VND) local search over relocate, Or-opt, and 2-opt moves evaluated against the monetised

objective (3). The full feasibility contract of this repair and the detailed VND template are given in Appendix B. The key coupling with the outer loop is that any customer whose repair leaves unserved ($n_j = 0$) is recovered on a later iteration through the dual update in Step 4, which raises the profit of a route that serves j in the next SPPRC call.

Lagrangian Bound Convergence: PADMM-MA vs Subgrad-MA (Day 1, largest DC subproblem per instance)




Figure 3. Lagrangian dual-score (LB*, surrogate) and best primal upper-bound (UB*) trajectories for PADMM-MA and SUBGRAD-MA on four representative subproblems (Day 1, largest DC). The curves labelled LB* are Lagrangian dual scores used as an internal convergence monitor; they are not certified lower bounds of the original objective (3) (see footnote on Proposition 3).

4.2. Adaptive β via Residual Balancing

A naive subgradient-ascent update for β is known to induce pricing oscillation when the underlying penalty is highly non-linear. The residual-balancing rule used here, adapted from [7], addresses this directly. Let $r^{(k)} = L^{(k)} / (C^{(k)} + L^{(k)})$ be the quality-loss share of the current upper bound. We update

$$\beta^{(k+1)} = \begin{cases} \min(\tau_{\beta} \beta^{(k)}, \beta_{\max}) & r^{(k)} > r_{\text{hi}}, \\ \max(\beta^{(k)} / \tau_{\beta}, \beta_{\min}) & r^{(k)} < r_{\text{lo}}, \\ \beta^{(k)} & \text{otherwise (dead zone)}. \end{cases} \quad (16)$$

The dead zone prevents oscillation around equilibrium; the multiplicative form yields exponential rather than additive corrections, accelerating convergence. The numerical values of the update parameters τ_{β} , r_{lo} , r_{hi} , $\beta_{\min}$, and $\beta_{\max}$ are reported with the rest of the experimental configuration in Section 5.1.

4.3. Rolling Horizon with ADP-Style Feedback

Algorithm 2 summarises the seven-day driver.

Algorithm 2 Seven-day rolling-horizon simulator.

- 1: Initialise state $(I_d^0, \bar{q}_d^0)$; set the initial marginals $\bar{c}_0^{\text{trans}}$ and $\bar{c}_0^{\text{qloss}}$ to small positive non-informative defaults (values in Section 5.1), chosen positive so that the day-1 stockout penalty tier $M_d^1 = \kappa(\bar{c}_0^{\text{trans}} + \bar{c}_0^{\text{qloss}})$ is strictly positive and the outer LP is well-posed; sensitivity to this initialisation is absorbed within two rolling-horizon iterations as the ADP update in Step 9 overwrites them with empirically grounded values.
- 2: **for** $t = 1, \dots, H$ **do**
- 3: Solve outer LP (12)–(15) with penalty tier $M_d^t = \kappa(\bar{c}_{t-1}^{\text{trans}} + \bar{c}_{t-1}^{\text{qloss}})$.
- 4: Quality-blend to obtain $\bar{q}_d^t$ at each DC.
- 5: **for** each DC d **do**
- 6: Tighten TW per Proposition 1 with $\bar{q}_d^t$.
- 7: Solve inner VRPTW via PADMM-MA.
- 8: **end for**
- 9: Refresh marginals from day- t inner solves as demand-weighted per-kg averages $\bar{c}_t^{\text{trans}} = (\sum_d \sum_{j \in \mathcal{S}_d} c_{d,j}^{\text{trans},t}) / (\sum_d D_d^t)$ and $\bar{c}_t^{\text{qloss}} = (\sum_d \sum_{j \in \mathcal{S}_d} L_j(t_j^{\text{arr}})) / (\sum_d D_d^t)$, where $c_{d,j}^{\text{trans},t}$ is the per-unit-demand transport cost attributed to store j by the inner VRPTW solve at DC d on day t , and the denominator aggregates total day- t demand.
- 10: **end for**

The pair $(\bar{c}_t^{\text{trans}}, \bar{c}_t^{\text{qloss}})$ acts as a post-decision state-value estimate in approximate dynamic programming [25]: by carrying the previous day's realised per-kg transport and quality-loss costs into the current day's stockout penalty tier M_d^t , Algorithm 2 couples outer reallocation decisions to anticipated inner-VRPTW consequences without enumerating future scenarios. This is the “ADP-style feedback” referenced throughout the paper; it is a single-step value-function approximation rather than a certified policy-iteration scheme, and no convergence guarantee across the rolling horizon is claimed.

4.4. Convergence Remarks

Proposition 3 (Monotone dual-score trajectory). *Within any fixed day t and fixed β , the running best Lagrangian dual score (We deliberately use “dual score” rather than “lower bound”: because β adapts across outer PADMM iterations, the scores BestLB^k collected at different β values are evaluated under different Lagrangians and are therefore not directly comparable as a single problem's lower bound. Within a fixed- β phase, they are bona fide Lagrangian duals of the relaxed subproblem; across phases, we treat the trajectory only as a convergence indicator, not a certified lower bound. See Section 5.5 for the empirical use of this trajectory) generated by PADMM-MA is non-decreasing in k .*

Proof. By definition $\text{BestLB}^{k+1} = \max(\text{BestLB}^k, \text{LB}^{k+1})$. $\square$

Proposition 4 (β -monotone quality loss). *For a fixed instance, fixed λ , and fixed preprocessing, the optimal total quality loss of the FDP-TW-MATURITY subproblem is non-increasing in β .*

Proof. Let $f(\beta)$ be the min-Lagrangian-cost value at β and $L^*(\beta)$ the corresponding quality loss. For any $\beta_1 < \beta_2$, optimality at β_1 and β_2 gives $f_1 + \beta_1 L_1^* \leq f_1 + \beta_1 L_2^*$ and the symmetric inequality. Adding: $(\beta_2 - \beta_1)(L_1^* - L_2^*) \geq 0$, hence $L_1^* \geq L_2^*$. $\square$

Remark 3 (Scope of Propositions 3 and 4). *Propositions 3 and 4 are structural results about a single SPPRC subproblem under fixed multipliers, not global convergence statements about PADMM-MA. Proposition 3 asserts only that the running maximum over a sequence of dual scores is non-decreasing, a bookkeeping fact that holds for any such sequence; as the footnote on p. 14 records, these scores are evaluated under different Lagrangians whenever β changes across iterations and are therefore not comparable as bounds on a single problem. Proposition 4 assumes λ and β both fixed,*

which is violated on purpose by the residual-balancing update (16). Neither statement implies that the full PADMM-MA iteration converges to a stationary point of the augmented Lagrangian in (3)–(8); that stronger property is inherited from the general non-convex proximal ADMM analysis of Zhang and Luo [8] under the prox-regularity and descent conditions stated there, which we do not re-derive here. Empirical convergence behaviour on the twelve test instances is reported in Section 5.5 and Figure 3; those observations do not substitute for a formal convergence theorem on the full scheme, which remains outside the scope of this paper.

5. Computational Experiments

The algorithmic machinery of Section 4—the inner PADMM-MA scheme (Algorithm 1), the outer reallocation LP (12)–(15), and their coupling via the ADP-style marginal-cost refresh in Algorithm 2—raises three questions that the remainder of this section addresses in turn. The first concerns head-to-head behaviour: how does PADMM-MA compare to an otherwise identical maturity-aware subgradient baseline on cost, wall-clock time, and arrival quality across a representative testbed (Section 5.2, with paired statistical tests and a scale extension to $n = 100$ customers)? The second is structural: which components of the scheme—maturity awareness, lateral reallocation, the adaptive β rule, and the proximal term—are empirically load-bearing, and under what conditions (Section 5.3)? The third asks how sensitive the conclusions are to the key physical and economic parameters of the cold-chain model (Section 5.4). An exact-MIP benchmark on tractable small subproblems (Appendix A, Table A1) then anchors the heuristic results against provable optima where these can be computed.

All experiments were run in MATLAB R2024b (The MathWorks, Inc., Natick, MA, USA) on a desktop PC. The full benchmark comprises $12 \times 6 = 72$ algorithm–instance combinations; all 72 terminated with feasible solutions (0 failures). Total wall-clock time: 144.3 minutes.

5.1. Instances and Configurations

Twelve semi-real mandarin instances are generated by overlaying the maturity and multi-echelon structure on Solomon [24] benchmarks. The Solomon geometries are truncated to 50 customers and embedded in a three-DC, seven-day rolling-horizon framework with mandarin-specific parameters: storage age $\tau_{\text{storage}} = 4$ days, transport temperature $T_{\text{transport}} = 10^\circ\text{C}$, three retail tiers (premium/standard/discount) with unit values 8/6/4 CNY/kg and acceptance thresholds $q^{\min} = 0.75/0.65/0.55$.

5.1.1. Cold-Chain Kinetics Calibration

The decay constants introduced symbolically in Section 3 are fixed here. The universal gas constant is $R = 8.314 \text{ J mol}^{-1} \text{ K}^{-1}$, the activation energy is $E_a = 50 \text{ kJ/mol}$, and the reference temperature is $T_{\text{ref}} = 5^\circ\text{C}$. The transport rate constant k_{ref} is calibrated at $T_{\text{calib}} = 10^\circ\text{C}$ so that the tightest retail channel (premium, $q_0 = 0.85$, $q^{\min} = 0.75$) reaches its maximum allowable transport duration at 60% of the Solomon depot horizon. This 60% target is a testbed-wide calibration choice rather than a physical constant: it is selected so that the Proposition 1 time-window tightening is active on some, but not all, of the twelve instances. The maturity-aware term therefore bites on a meaningful subset of instances without making the benchmark uniformly infeasible, and the main-testbed conclusions are robust to this choice within its plausible operational range, as documented by the τ_{storage} and $T_{\text{transport}}$ sweeps in Section 5.4. DC cold-storage decay uses a separate literature-calibrated rate $k_{\text{storage}} = \ln(0.95/0.50)/21 \approx 0.0306 \text{ d}^{-1}$, reflecting approximately 21 days from quality 0.95 to 0.50 at 5°C [16].

5.1.2. Truncation Rule

The truncation from 100 to 50 customers is deterministic and non-selective: for every Solomon instance the first 50 customer records are retained in the order supplied by the original benchmark file (customers with index $1, 2, \dots, 50$ in the standard distribution), and the remainder is discarded. No filtering on demand, coordinate, or time-window attributes is applied, so the truncated instance preserves the original geographical class (C/R/RC) and time-window regime (narrow/wide) of the full instance. The three-DC assignment is then obtained by running k -means on the (x, y) coordinates of the retained 50 customers with $k = 3$, so each DC serves a contiguous cluster and the clustering is reproducible from the instance file alone. Specifically, the clustering uses Lloyd's algorithm with random seeding by the Mersenne Twister generator (`rng(seed, 'twister')`); k initial centres are drawn uniformly without replacement from the customer set, the update is iterated until the label assignments stabilise or a cap of 50 iterations is reached. The instance-generation seed is set to $\text{seed} = 42$ across the entire main testbed, which makes the clustering bit-exactly reproducible: the same seed regenerates the same partition. Scaling to $n = 100$ in Section 5.2 uses the full original 100 customer record set under the same k -means procedure.

5.1.3. Outer-LP Cost Parameters

The unit costs entering the outer reallocation LP (12) are calibrated as follows. The origin-to-DC replenishment cost is $c_d^o = 0.02 \cdot \text{dist}(o, d)$ CNY/kg and the lateral DC-to-DC transfer cost is $c_{d_1 d_2}^\ell = 0.03 \cdot \text{dist}(d_1, d_2)$ CNY/kg, both proportional to road distance and chosen so that lateral transfers carry a 50% per-kilometre premium over origin replenishment, reflecting the additional handling and handover overhead of an inter-DC leg. The holding-overshoot cost is a flat $h_d = 0.10$ CNY/kg/day at every DC; this charge applies only to the portion $z_d^t = \max\{0, I_d^t - I_d^{\text{target}}\}$ of inventory above the target, not to inventory itself. Inventory bounds are scaled to each DC's average daily demand $\bar{D}_d = H^{-1} \sum_t D_d^t$: the warehouse capacity is $I_d^{\text{max}} = 5.0 \bar{D}_d$ (five days of average throughput) and the target stock is $I_d^{\text{target}} = 1.5 \bar{D}_d$ (one and a half days), so that holding overshoot is triggered only when inventory exceeds 1.5 days of forward cover. Together with $\kappa = 10$, these settings produce the cost ordering $h_d < c_d^o < c_{d_1 d_2}^\ell \ll M_d^t$ discussed after Equation (12), in which the LP defaults to origin replenishment, escalates to lateral transfer only under inventory asymmetry, and resorts to stock-outs only when the first two channels are infeasible.

The twelve instances span a 3×2 design: three geographical types (C—clustered, R—random, RC—mixed) and two time-window widths (narrow 1xx, wide 2xx), with two instances per cell.

Six configurations are compared:

1. PADMM-MA (Full): proposed method with all components;
2. SUBGRAD-MA: subgradient multiplier update with same maturity-aware subproblem;
3. PADMM-NOMAT: maturity awareness disabled ($q_j^{\min} = 0, \beta = 0$);
4. PADMM-NOREALLOC: lateral DC-to-DC transfers forbidden;
5. PADMM-STATIC β : β fixed (no residual balancing);
6. PADMM-NOPROX: proximal term disabled ($\alpha_{\text{prox}} = 0$).

5.1.4. Algorithm Parameters

The PADMM-MA control parameters introduced symbolically in Section 4 take the following values, applied uniformly across all twelve instances. Termination (Algorithm 1): $K_{\text{max}} = 40$, $\varepsilon_{\text{prim}} = 10^{-3}$, $\varepsilon_{\text{rel}} = 10^{-3}$, $K_{\text{stall}} = 5$, and proximal-anchor weight $\alpha_{\text{prox}} = 0.1$. Residual-balancing β update (16): $(\tau_\beta, r_{\text{lo}}, r_{\text{hi}}) = (1.5, 0.05, 0.20)$ and $(\beta_{\text{min}}, \beta_{\text{max}}) = (0.1, 50)$. Rolling-horizon initialisation (Algorithm 2): the day-0 marginals are set to $\bar{c}_0^{\text{trans}} = \bar{c}_0^{\text{qloss}} = 0.5$ CNY/kg.

All implementation, computational experiments, and statistical analysis reported in this section were performed by the authors. AI-assisted editing was limited to language polishing and formatting checks, as detailed in the Acknowledgments.

5.2. Main Results

Table 3 reports the seven-day totals for the three primary methods on all twelve instances.

Table 3. Comparison of solution algorithms across Solomon VRPTW benchmark instances (7-day rolling horizon, 50 stores, 3 DCs). Cost and distance in CNY; $\bar{q}_{arr}$: mean arrival quality; Time in seconds.

Type	Instance	PADMM-MA					Subgrad-MA					PADMM-NoMat				
		Cost	Dist	QLoss	$\bar{q}_{arr}$	Time	Cost	Dist	QLoss	$\bar{q}_{arr}$	Time	Cost	Dist	QLoss	$\bar{q}_{arr}$	Time
C (Clustered)	C101	9648	1844	3431	0.860	5	9182	1849	2962	0.869	54	6622	2250	0	0.845	19
	C104	9594	2348	2874	0.860	20	8715	2260	2082	0.867	247	6573	2201	0	0.860	271
	C201	13,545	2949	5972	0.837	9	12,664	2900	5139	0.836	71	7370	2745	0	0.809	11
	C204	11,620	3090	3906	0.855	10	10,488	3193	2670	0.859	1072	7189	2565	0	0.830	765
R (Random)	R101	12,127	4192	4218	0.860	9	10,363	3289	3356	0.873	17	9586	5868	0	0.827	7
	R105	12,923	4640	4565	0.864	8	11,746	4166	3862	0.864	17	9276	5558	0	0.830	8
	R201	12,629	4991	3920	0.848	9	12,674	5099	3857	0.847	56	8782	5064	0	0.836	15
	R205	12,179	4602	3859	0.850	12	11,961	4748	3495	0.851	326	8310	4592	0	0.826	50
RC (Mixed)	RC101	9911	855	1683	0.863	26	8286	245	667	0.832	32	9462	2089	0	0.830	37
	RC105	10,594	1266	1956	0.870	66	9557	800	1384	0.873	65	9357	1984	0	0.832	105
	RC201	19,384	3770	8242	0.839	40	18,648	4485	6790	0.843	385	10,521	3148	0	0.832	106
	RC205	17,500	3737	6391	0.849	67	16,752	3875	5504	0.853	1861	10,250	2876	0	0.837	29
Avg. Narrow		10,799	2524	3121	0.863	22	9642	2102	2386	0.863	72	8479	3325	0	0.837	74
Avg. Wide		14,476	3856	5381	0.846	24	13,864	4050	4576	0.848	629	8737	3498	0	0.828	162
Avg. All		12,638	3190	4251	0.855	23	11,753	3076	3481	0.856	350	8608	3412	0	0.833	118

The remainder of this subsection works through the patterns the table makes visible.

5.2.1. Quality-Cost Monetisation Changes Routing Structure

PADMM-NOMAT achieves the lowest *apparent* cost (avg. 8608 CNY) because it does not monetise quality loss (The QLoss column for PADMM-NOMAT in Table 3 is identically zero *by construction*: under this configuration the quality-loss weight β and the minimum-quality threshold $q_j^{\min}$ are both set to zero, so the term $L_j(t_j^{\text{arr}})$ does not enter the objective (3). The zero values therefore record the absence of this cost component in the objective, not an absence of physical quality degradation; the latter is implicit in the higher average route distance reported in the same row). However, its average distance is *higher* than both maturity-aware methods (3412 CNY for PADMM-NOMAT vs. 3190 CNY for PADMM-MA and 3076 CNY for SUBGRAD-MA, all from Table 3), showing that the quality-aware objective reshapes the route structure towards shorter, fresher deliveries.

5.2.2. Cost-Speed Positioning of the Two Heuristics

SUBGRAD-MA attains a slightly lower mean seven-day total cost (11,753 CNY) than PADMM-MA (12,638 CNY), a gap of about 7.5%, at near-identical mean arrival quality ($\bar{q}_{arr} = 0.855$ vs. 0.856). The two methods are therefore *not* ranked by solution quality in favour of PADMM-MA; what distinguishes PADMM-MA is wall-clock time. Its mean wall-clock is 23 seconds versus 350 seconds for SUBGRAD-MA (Table 3), a roughly 15 $\times$ speed-up on average, widening further on wide-TW instances: on RC205, SUBGRAD-MA requires 1861 seconds ($\approx$ 31 minutes) against 67 seconds for PADMM-MA. On this basis PADMM-MA is positioned as a *fast operational heuristic* for daily re-planning under a computational budget, rather than as a cost-dominating solver; SUBGRAD-MA remains the preferable choice when the operational window permits offline optimisation and the additional 7.5% cost saving is warranted.

5.2.3. Instance-Class Effects

RC-type instances exhibit the highest total costs (avg. 14,347 CNY for PADMM-MA; Table 4), driven by the mixed geographical structure that increases both distance and quality loss. Wide-TW instances (2xx) show higher quality loss than narrow-TW (1xx) across all methods, consistent with the longer feasible delivery windows, allowing more quality degradation. A notable within-class outlier is RC101, whose total distance (855 CNY under PADMM-MA and 245 CNY under SUBGRAD-MA) is an order of magnitude below the other three RC instances. This is a geometric artefact of the k -means three-DC assignment on the first 50 RC101 customers (Section 5.1): the three cluster centroids happen to fall within a small coordinate sub-region, so that each DC's internal routing distance is correspondingly compressed. The effect is specific to RC101's first-50 truncation; RC105's truncated customer set, though of the same class (RC, narrow TW), has a more dispersed coordinate distribution and its distance totals (1266 CNY under PADMM-MA; 800 CNY under SUBGRAD-MA) lie within the normal RC-narrow range. The pattern is not a data error and does not propagate to the total cost, where RC101 remains within the RC range, because the quality-loss term in (3) does not scale with route distance in the same geometric fashion. The same compressed RC101 geometry also explains the within-class arrival-quality gap between PADMM-MA and SUBGRAD-MA on this single instance ($\bar{q}_{arr} = 0.863$ vs. 0.832 on Table 3): SUBGRAD-MA exploits the very short feasible distances by concentrating service later in each TW, accumulating more decay, while PADMM-MA's quality penalty steers it towards earlier service, even when slack distance is available. The effect is RC101-specific and does not materially shift the twelve-instance mean (0.855 vs. 0.856). A second within-class anomaly is the C201 row, whose mean arrival quality ($\bar{q}_{arr} = 0.837$ for PADMM-MA; 0.836 for SUBGRAD-MA) is markedly lower than the other three C instances (0.855–0.860). C201 combines wide TW with the clustered geometry, allowing the heuristic to push service times late in the day to consolidate stops; the feasible TW window is wide enough that the Proposition 1 tightening is not active, so quality decay accumulates over the full route. The same effect is visible to a smaller degree on R201 and the RC2xx instances; C201 is the most extreme case because its narrow geographical spread amplifies the late-service consolidation behaviour.

Table 4. Average results by instance class and time-window width. The QL% column reports quality loss as a fraction of total cost (QLoss/Cost).

Class	TW	Algorithm	Cost	Dist	QLoss	$\bar{q}_{arr}$	QL%
C	Narrow	PADMM-MA	9621	2096	3152	0.860	32.8%
		Subgrad-MA	8949	2055	2522	0.868	28.2%
		NoMat	6597	2225	0	0.853	0.0%
	Wide	PADMM-MA	12,583	3019	4939	0.846	39.3%
		Subgrad-MA	11,576	3046	3905	0.847	33.7%
		NoMat	7280	2655	0	0.820	0.0%
R	Narrow	PADMM-MA	12,525	4416	4391	0.862	35.1%
		Subgrad-MA	11,055	3728	3609	0.869	32.6%
		NoMat	9431	5713	0	0.829	0.0%
	Wide	PADMM-MA	12,404	4796	3890	0.849	31.4%
		Subgrad-MA	12,317	4924	3676	0.849	29.8%
		NoMat	8546	4828	0	0.831	0.0%
RC	Narrow	PADMM-MA	10,253	1060	1819	0.867	17.7%
		Subgrad-MA	8921	522	1026	0.853	11.5%
		NoMat	9410	2036	0	0.831	0.0%
	Wide	PADMM-MA	18,442	3753	7316	0.844	39.7%
		Subgrad-MA	17,700	4180	6147	0.848	34.7%
		NoMat	10,385	3012	0	0.835	0.0%

5.2.4. Statistical Significance of the Cost and Time Differences

Table 5 reports paired Wilcoxon signed-rank tests, paired t -tests, and Cohen's d for the three primary metrics across the twelve paired (PADMM-MA, SUBGRAD-MA) observations. The cost and wall-clock-time differences are both highly significant (Wilcoxon $p = 0.001$;

Cohen's $d = 1.69$ for cost and $d = -0.59$ for time). In contrast, the mean arrival-quality difference is *not* significant at the 5% level (Wilcoxon $p = 0.077$, $d = -0.11$). The statistical picture matches the narrative above: the two heuristics differ substantially in cost and speed but not materially in delivered quality. The cost significance survives a TW-stratified analysis, with narrow-TW exhibiting a larger effect size ($d = 2.41$) than wide-TW ($d = 1.39$).

Table 5. Paired statistical tests comparing PADMM-MA (A) and SUBGRAD-MA (B) across the twelve main-testbed instances. Wilcoxon signed-rank is the primary non-parametric test; paired t -test is reported for reference. Cohen's d is the standardised effect size for paired samples (positive: $A > B$). Cost and Wall Time are in CNY and seconds respectively; $\bar{q}_{\text{arr}}$ is dimensionless.

Metric/Subset	n	Mean A	Mean B	Wilcoxon p	t -Test p	Cohen's d
Total Cost (all)	12	12,637.85	11,753.02	0.0010	0.0001	+1.689
Wall Time (all)	12	23.43	350.33	0.0010	0.0645	−0.593
$\bar{q}_{\text{arr}}$ (all)	12	0.8546	0.8556	0.0771	0.7225	−0.105
Cost (narrow TW)	6	10,799.45	9641.59	0.0312	0.0020	+2.405
Cost (wide TW)	6	14,476.25	13,864.46	0.0625	0.0190	+1.393

A = PADMM-MA, B = SUBGRAD-MA. Mean A and Mean B columns share the unit of the row's metric. All test statistics are computed from the full-precision per-instance values stored in the released .mat files; the values displayed in Table 3 are rounded to three decimal places for $\bar{q}_{\text{arr}}$ and to integers for Cost and Wall Time, so a re-computation from the rounded display values will give slightly different p -values for $\bar{q}_{\text{arr}}$ (within the precision permitted by the rounding).

5.2.5. Scale Extension to $n = 100$ Customers

Table 6 reports PADMM-MA on the twelve instances when each geometry is rebuilt with 100 customers rather than 50. All twelve doubled instances remain feasible; the cost roughly doubles (mean 20,982 vs. 12,638 CNY), distance increases by roughly 60%, and quality-loss totals grow broadly in proportion to the demand doubling. Wall-clock time scales super-linearly with customer count (mean 120 s at $n = 100$ vs. 23 s at $n = 50$; geometric-mean ratio $\approx 5\times$), consistent with the NG-route SPPRC labelling step dominating runtime at larger scales. The ratio of quality loss to total cost at $n = 100$ is comparable to that at $n = 50$ across all twelve instances, confirming that the maturity-aware objective remains active—rather than being dominated by distance cost—at twice the baseline customer count.

Several observations bear on larger-scale applications. The dominant component as the instance size grows is the inner NG-route SPPRC labelling pass invoked once per pricing call: the outer LP and the four-phase repair/VND scale mildly, whereas the number and length of partial-path labels grow with the per-DC customer count. Consequently the primary bottleneck is expected to be labelling complexity (the number of non-dominated labels) rather than memory, since the four-dimensional resource vector of Proposition 2 keeps each label compact and the NG-neighbourhood bound caps the label set; memory would become binding only at much larger per-DC clusters than considered here. The empirically observed growth—wall-clock rising about $5\times$ for a $2\times$ customer increase, with the per-DC problem size held in check by the fixed three-DC clustering—is super-linear but far from the worst-case exponential blow-up of elementary pricing, consistent with the NG-relaxation controlling, though not eliminating, the combinatorial growth. For deployments substantially beyond $n = 100$, three modifications would most directly improve scalability: increasing the number of DCs (or sub-clustering large clusters) to keep each SPPRC small; dynamic NG-neighbourhood resizing to bound the label set adaptively; and bidirectional or bucket-based labelling to prune dominated labels earlier. A full-scale law study of the labelling pass is left to future work.

Table 6. Scale comparison for PADMM-MA: $n = 50$ (baseline) versus $n = 100$ customers. Demand, cost, quality loss, mean arrival quality, and wall-clock time scale with customer count; arrival quality is essentially preserved, confirming that the maturity-aware objective remains active at larger scales.

Instance	Cost (CNY)		Dist (CNY)		QLoss (CNY)		Time (s)	
	$n = 50$	$n = 100$	$n = 50$	$n = 100$	$n = 50$	$n = 100$	$n = 50$	$n = 100$
C101	9648	13,602	1844	2148	3431	2776	5.4	49.5
C104	9594	16,562	2348	4503	2873	3382	19.6	290.5
C201	13,545	30,427	2949	6722	5971	14,705	9.0	38.1
C204	11,620	26,527	3090	6683	3906	10,843	10.2	68.9
R101	12,127	12,543	4191	2643	4218	3564	9.3	64.3
R105	12,923	18,607	4640	5224	4565	7045	7.9	73.2
R201	12,629	22,538	4991	6743	3920	9458	9.2	70.1
R205	12,179	21,816	4602	6205	3859	9274	11.8	138.5
RC101	9911	13,939	855	1818	1683	2208	25.8	104.9
RC105	10,594	13,938	1266	1825	1956	2200	66.5	350.6
RC201	19,384	30,046	3769	7646	8241	12,487	39.7	77.6
RC205	17,500	31,242	3737	8063	6391	13,266	66.8	112.9
Mean	12,638	20,982	3190	5019	4251	7601	23.4	119.9

5.3. Ablation Study

Two ablations are reported. The first (Table 7) removes one component of PADMM-MA at a time on the twelve main-testbed instances. The second probes the same components under conditions chosen to exercise them: asymmetric DC demand for lateral reallocation, and a range of fixed β values for the adaptive rule. The combined picture is richer than the single-testbed ablation suggests.

Table 7. Ablation study: effect of removing individual PADMM-MA components. $\Delta\text{Cost}\%$ is relative to the Full configuration.

Type	Inst.	Full				NoRealloc				StaticBeta				NoProximal			
		Cost	Dist	QLoss	Time	Cost	Dist	QLoss	Time	Cost	Dist	QLoss	Time	Cost	Dist	QLoss	Time
C (Clustered)	C101	9648	1844	3431	5	9648	1844	3431	11	9398	1925	3101	34	9661	1846	3444	12
	C104	9594	2348	2874	20	9594	2348	2874	24	9338	2314	2652	335	9584	2356	2856	22
	C201	13,545	2949	5972	9	13,545	2949	5972	8	13,265	2998	5642	42	13,545	2949	5972	9
	C204	11,620	3090	3906	10	11,620	3090	3906	10	10,875	3023	3227	550	11,536	3108	3804	14
R (Random)	R101	12,127	4192	4218	9	12,127	4192	4218	9	11,952	4072	4162	16	12,127	4192	4218	8
	R105	12,923	4640	4565	8	12,923	4640	4565	8	12,585	4552	4315	19	12,923	4640	4565	8
	R201	12,629	4991	3920	9	12,629	4991	3920	10	12,189	5001	3470	33	12,612	4991	3903	10
	R205	12,179	4602	3859	12	12,179	4602	3859	11	11,600	4643	3239	61	12,228	4603	3908	11
RC (Mixed)	RC101	9911	855	1683	26	9911	855	1683	25	9909	865	1671	32	9911	855	1683	26
	RC105	10,594	1266	1956	66	10,594	1266	1956	93	10,342	1171	1798	335	10,594	1266	1956	262
	RC201	19,384	3770	8242	40	19,384	3770	8242	69	18,949	3974	7602	432	19,384	3770	8242	67
	RC205	17,500	3737	6391	67	17,500	3737	6391	8	16,872	3949	5550	116	17,500	3737	6391	10
Avg. $\Delta\text{Cost}\%$		–	–	–	–	+0.00%	–	–	1.0×	–2.86%	–	–	7.1×	–0.03%	–	–	1.6×

5.3.1. NoRealloc: Symmetric-Redundant, Asymmetric-Required

On the twelve baseline instances, disabling lateral DC-to-DC transfers has no measurable effect ($\Delta\text{Cost} = 0.00\%$ on every instance in Table 7), because the three-DC assignment produced by k -means on the Solomon coordinates is structurally balanced: no DC ever needs stock from its neighbours. This apparent null result is an artefact of the testbed, not a property of the mechanism. Table 8 reports the same ablation under scaled demand vectors $[1.5, 1.0, 0.5]$ and $[1.8, 1.0, 0.2]$ with a capped origin-replenishment rate (at $1.2 \times$ mean daily demand per DC) and mandatory demand coverage. In every one of the six skewed-demand instances and six extreme-demand instances, PADMM-NOREALLOC is LP-infeasible: the high-demand DC cannot meet its obligations from local inventory plus capped origin supply without a lateral transfer. The PADMM-MA full configuration, in contrast, remains feasible on all twelve asymmetric instances, moving on average 2218 kg (skewed) and 3312 kg (extreme) laterally per seven-day horizon. Lateral reallocation therefore transitions from a cost-saving mechanism at the margin (the $+0.00\%$ symmetric case) to a feasibility-required mechanism under any meaningful demand imbalance; the apparent null result in Table 7 should not be read as a case against the mechanism.

Table 8. Value of lateral reallocation under demand asymmetry. $\Delta\%$ is relative cost increase when lateral DC transfers are disabled.

Instance	Symmetric			Skewed			Extreme		
	Full	NoR	$\Delta\%$	Full	NoR	$\Delta\%$	Full	NoR	$\Delta\%$
C101	9648	9648	+0.00	9570	infeas.	$\uparrow \infty$	8541	infeas.	$\uparrow \infty$
C201	13,545	13,545	+0.00	17,178	infeas.	$\uparrow \infty$	18,918	infeas.	$\uparrow \infty$
R101	12,127	12,127	+0.00	12,240	infeas.	$\uparrow \infty$	9783	infeas.	$\uparrow \infty$
R201	12,629	12,629	+0.00	15,500	infeas.	$\uparrow \infty$	16,757	infeas.	$\uparrow \infty$
RC101	9911	9911	+0.00	13,001	infeas.	$\uparrow \infty$	15,510	infeas.	$\uparrow \infty$
RC201	19,384	19,384	+0.00	22,849	infeas.	$\uparrow \infty$	25,108	infeas.	$\uparrow \infty$
Avg	12,874	12,874	+0.00	—	all infeas.	$\uparrow \infty$	—	all infeas.	$\uparrow \infty$

5.3.2. StaticBeta: Cost Is Not the Decisive Metric

The Table 7 column STATICBETA uses a single fixed β (tuned per instance in the original experiments) and finds mean cost about 2.9% lower than the adaptive rule. We do not contest this number. Table 9 resolves what it means, by sweeping four fixed values $\beta \in \{0.5, 1.0, 2.0, 5.0\}$ on each instance. The cost-minimising fixed β is *not* universal: it is 0.5 on 4/12 instances, 1.0 on 4/12, 2.0 on 3/12, and 5.0 on 1/12. There is therefore no single fixed β that can realise the -2.9% figure without per-instance tuning—which defeats the purpose of an operational heuristic. The adaptive rule’s value lies elsewhere. On four of the twelve instances (C104, C204, R205, RC205), a fixed $\beta \leq 1.0$ produces severe inner-subproblem label growth, driving wall-clock times to 86–272 seconds compared with 5–12 seconds under the adaptive rule (per- β wall-clock timings are recorded in the released .mat files; Table 9 reports costs only, to keep the within-instance comparison readable); a low β weakens quality-loss penalisation and the NG-SPPRC labelling cannot discriminate among time-sensitive labels. The adaptive rule detects the low quality-loss share $r^{(k)} < r_{10}$ and multiplicatively raises β , preventing the label explosion. The trade-off is a $\sim 1\text{--}4\%$ mean-cost premium in exchange for a bounded worst-case runtime; for a fast operational heuristic, that is the right trade.

Table 9. Per-instance cost (CNY) under four fixed quality-loss weights $\beta \in \{0.5, 1.0, 2.0, 5.0\}$ versus the adaptive residual-balancing rule (PADMM-MA Full). The cost-minimising fixed β varies across instances (4 at $\beta = 0.5$, 4 at $\beta = 1.0$, 3 at $\beta = 2.0$, 1 at $\beta = 5.0$); per-instance optima are marked with *. The mean per-instance cost saving from picking the best fixed β is -2.86% versus adaptive, but this saving requires per-instance tuning. On the four instances marked † (C104, C204, R205, RC205), low β drives inner-SPPRC label growth and pushes wall-clock time to 86–272 s, versus 5–12 s under the adaptive rule (per- β wall-clock timings recorded in the released .mat files; not tabulated here).

Type	Instance	$\beta = 0.5$	$\beta = 1.0$	$\beta = 2.0$	$\beta = 5.0$	Best β	Adaptive
C (Clustered)	C101	9561	9398 *	9561	9776	1	9648
	C104 †	9338 *	9655	9972	10,392	0.5	9594
	C201	13,522	13,265 *	13,522	13,863	1	13,545
	C204 †	10,875 *	11,109	11,342	11,651	0.5	11,620
R (Random)	R101	12,184	11,952 *	12,184	12,491	1	12,127
	R105	12,751	12,585 *	12,751	12,970	1	12,923
	R201	12,713	12,451	12,189 *	12,535	2	12,629
	R205 †	11,600 *	11,897	12,195	12,588	0.5	12,179
RC (Mixed)	RC101	10,170	10,039	9909 *	10,082	2	9911
	RC105	10,786	10,564	10,342 *	10,636	2	10,594
	RC201	19,909	19,620	19,331	18,949 *	5	19,384
	RC205 †	16,872 *	17,445	18,018	18,776	0.5	17,500
Mean		12,523	12,498	12,610	12,892	—	12,638

* per-instance cost-minimising fixed β ; ties broken in favour of the smaller β . † instances on which fixed $\beta \leq 1.0$ triggers severe inner-SPPRC label growth (wall-clock 86–272 s) compared with the adaptive rule (5–12 s on the same instances).

5.3.3. NoProximal: A Convergence Safeguard, Not an Empirical Cost Lever

Removing the proximal term changes the mean cost by only -0.03% (Table 7)—statistical noise on this testbed. The proximal term’s role in PADMM-MA is to guarantee the convergence behaviour of the non-convex proximal ADMM under the assumptions of Zhang and Luo [8]; on our twelve instances, this safeguard is not empirically exercised because the dual-primal residual balance remains well-conditioned throughout. The proximal term is kept for theoretical completeness and to preserve robustness on instances outside the current testbed; no claim is made that it yields a measurable cost improvement on this testbed.

5.4. Sensitivity Analysis

Figure 4 reports a four-parameter sweep on four representative instances (C101, R101, RC101, R201), each parameter varied over four levels while holding others at defaults. The sweep is presented in a four-panel layout for completeness; as discussed below, only three of the four parameters produce a non-trivial response on the symmetric main testbed.




Figure 4. Sensitivity of total cost, quality loss, and mean arrival quality $\bar{q}_{arr}$ to four key parameters (averaged over four representative instances). The bottom-right panel is included for completeness: the horizontal cost and quality-loss traces are flat for all four κ values because no stock-out events are triggered on the symmetric main-testbed demand (see the end of Section 5.4).

The three parameters with active responses are as follows:

- **Unit-value multiplier (α_v , 0.5–3.0):** The strongest driver. Increasing the unit-value multiplier six-fold from 0.5 to 3.0 raises the average cost from 9451 to 17,765 CNY (+88%), with quality loss increasing from 1664 to 9954 CNY. Arrival quality $\bar{q}_{arr}$ remains stable at 0.857–0.860, confirming that the algorithm adapts routes rather than sacrificing freshness.

- **Transport temperature** ($T_{\text{transport}}$, 4–16°C): Raising temperature from 4 °C to 16 °C increases the average cost by 17% and the quality loss by 89%. $\bar{q}_{\text{arr}}$ drops from 0.871 to 0.845.
- **Storage age** (τ_{storage} , 2–8 days): Moderate effect. Arrival quality $\bar{q}_{\text{arr}}$ decreases overall from 0.865 at $\tau_{\text{storage}} = 2$ days to 0.848 at $\tau_{\text{storage}} = 8$ days, but the trajectory is non-monotone, with a local dip near $\tau_{\text{storage}} \approx 6$ days visible in Figure 4. Cost response is likewise non-monotone due to the interaction between initial quality and the TW tightening effect.

5.4.1. Note on the κ Panel

The stockout-penalty sweep ($\kappa \in [3, 20]$) is uninformative on the symmetric main testbed because no instance triggers a stock-out: the slack variable u_d^t in (13) remains at zero at the LP optimum for every $\kappa \geq 3$, so κ drops out of the objective entirely. The relevant regime for differentiating κ is one in which stock-outs are feasibility-forced rather than optional, which would require the extreme-asymmetric demand setting of Table 8 combined with a tighter origin-supply cap than used here. We report this panel to make the null result auditable rather than to draw a managerial conclusion; the differentiating regime is exercised separately by the asymmetric-demand experiments reported in that table.

5.4.2. Practical Calibration of the Parameters

The sensitivity results above also indicate how the model would be calibrated in an industrial deployment, and which inputs deserve the most attention. The parameters fall into two groups by their likely source. *Operational data parameters* would be read directly from a firm's records: the unit values π_j and acceptance thresholds $q_j^{\min}$ from retail contracts; demands d_j^t , service times s_i , and travel times τ_{ij} from historical order and telematics data; the transport and storage temperatures from reefer-unit logs; and the cost coefficients c_d^o , $c_{d_1 d_2}^e$, h_d from the carrier tariff and warehousing accounts. *Managerial-judgement parameters* are those without a direct data source and set by planners: the pre-horizon storage age τ_{storage} , the handover latency, and the stock-out penalty multiplier κ , and the target/maximum inventory multipliers; for each of these, the paper reports a sensitivity sweep showing the conclusions are robust within a plausible range. In terms of influence, the sweeps rank the drivers clearly: the unit-value scale and the transport temperature have the largest effect on cost and on arrival quality respectively, the kinetic constants (E_a , k_{ref}) act through the same channel as temperature and matter most for the tightest retail tier, while the handover latency and κ are second-order on balanced demand. Estimating the temperature-related inputs and the highest-value acceptance thresholds accurately is therefore the priority. Finally, applying the framework to a product other than mandarins would require recalibrating only the decay-kinetics block— E_a , k_{ref} , T_{ref} , and the storage rate k_{storage} , taken from the post-harvest literature for that commodity—and the retail-tier values and thresholds; the routing, reallocation, and algorithmic components are product-independent and would carry over unchanged.

5.5. Convergence Observations

Figure 3 shows the per-iteration Lagrangian dual-score and primal upper-bound trajectories for four representative subproblems.

On the R101 and RC105 subproblems (narrow TW, 35–40 iterations), PADMM-MA's dual score climbs substantially faster than SUBGRAD-MA's, consistent with the proximal term's acceleration effect on the relaxed Lagrangian. On the wide-TW instances (RC201, RC205), both methods' primal upper bounds and dual scores meet within 4–8 iterations, indicating that the wide time windows admit many feasible routes and the repair heuristic quickly closes the internal *convergence gap* ($= 100\% \cdot (\text{UB} - \text{BestLB}) / \text{UB}$); this internal gap

is *not* an optimality gap relative to the exact subproblem optimum, which can be measured only against the MIP benchmark in Appendix A.

6. Managerial Insights

1. **Quality-as-cost framing pays off.** Monetising quality loss and minimising the sum rather than the distance alone reshapes the routing structure. Although PADMM-NOMAT reports lower apparent cost, its routing geometry differs systematically from the maturity-aware methods: on narrow-TW R- and RC-type cells PADMM-NOMAT routes are 29–92% longer than under PADMM-MA (R-Narrow: +29%; RC-Narrow: +92%; Table 4), while on wide-TW C- and RC-type cells, they are shorter; the aggregate effect across all twelve instances is $\sim 7\%$ longer routes under PADMM-NOMAT (3412 vs. 3190 CNY mean distance for PADMM-MA; Table 3). Either way, the unmonetised quality degradation under PADMM-NOMAT would translate to discounting or rejection at the store level.
2. **Speed favours PADMM; absolute cost favours the subgradient baseline.** PADMM-MA runs on average $15\times$ faster than SUBGRAD-MA (23 s vs. 350 s on Table 3) at statistically indistinguishable mean arrival quality ($\bar{q}_{\text{arr}} = 0.855$ vs. 0.856; paired Wilcoxon $p = 0.077$ for arrival quality, Table 5), while SUBGRAD-MA attains a 7.5% lower mean seven-day cost (Wilcoxon $p = 0.001$, Cohen's $d = 1.69$). The practical reading of this trade-off depends on the operational window. PADMM-MA is preferable whenever a plan must be produced under a tight time budget—overnight re-planning across many DCs, intra-day re-optimisation after demand or vehicle updates, or interactive what-if analysis—where the 7.5% premium (of order 900 CNY per DC-week against a 12,000–13,000 CNY base) is small relative to the value of returning a plan in seconds rather than minutes. The subgradient baseline remains preferable when the planning window is generous (e.g., a fixed weekly plan computed offline) and the savings are worth a per-instance budget of up to half an hour. As a rough guide, the faster solver is justified whenever the decision must be refreshed more frequently than the slower solver's runtime allows, or whenever the realised cost premium falls below the operation's tolerance for planning latency; on wide-TW instances, where SUBGRAD-MA can take ≈ 31 min (RC205), this threshold is easily met. The choice is therefore operational, not algorithmic.
3. **Temperature control is the strongest quality lever among those tested.** Among the parameters with an active response on the main testbed, transport temperature has the largest effect on arrival quality ($\bar{q}_{\text{arr}} = 0.871$ at 4°C vs. 0.845 at 16°C ; Figure 4), and quality loss rises by 89% over the same range. Whether investment in stronger refrigeration is economically justified depends on local equipment and energy costs that are outside the scope of this paper; the result establishes the direction and approximate magnitude of the effect, not its return on investment.
4. **Lateral reallocation is a feasibility safeguard, not a routine cost lever.** On the balanced main testbed, forbidding DC-to-DC transfers changes cost by 0.00% on every instance (Table 7): when demand is symmetric across DCs, origin replenishment alone suffices and lateral transfer is redundant. Under asymmetric demand, the picture reverses—disabling lateral transfer renders every skewed- and extreme-demand instance infeasible (Table 8), whereas the full model stays feasible by moving on average 2.2–3.3 t laterally per horizon. The lesson is that lateral reallocation should be provisioned for its option value under demand imbalance (regional promotions, supply shocks, uneven store openings), even though it lies dormant when demand is balanced.

Freshness Reprioritises Routes, and the Framework Supports Daily Use

Two cross-cutting lessons close this section. First, freshness considerations change routing priorities relative to classical distance-based routing: the maturity cost steers vehicles towards earlier service of high-value, tight-acceptance stores, even at the expense of slightly longer travel, so the cheapest route by distance is frequently not the cheapest once quality is monetised (Tables 3 and 4). Second, the rolling-horizon design (Figure 2) is intended for daily decision-making: each day's plan is produced in seconds from the current inventory and quality state, and the ADP-style marginal-cost feedback carries the consequences of today's routing into tomorrow's reallocation without re-solving the whole horizon, matching the cadence of operational cold-chain re-planning.

7. Conclusions

This paper has addressed a maturity-aware variant of the multi-echelon VRPTW under first-order quality decay, motivated by the distribution of already-aged mandarin stock through a three-echelon cold chain over a seven-day rolling horizon. The MA-MEDR-VRPTW was cast as a decomposable multi-period mathematical program in which an outer linear programme handles inventory reallocation and an inner maturity-aware VRPTW handles routing at each DC, with the two layers coupled through approximate-dynamic-programming-style marginal costs rather than a Benders-style nesting. On the modelling side, a first-order-kinetics minimum-quality acceptance constraint was shown to reduce analytically to a time-window tightening (Proposition 1). This reduction removes all per-extension exponential evaluations from the inner SPPRC. Under the arc-wise-constant-temperature assumption (Assumption 2, with the validity boundary clarified in Remark 2), the standard four-dimensional NG-SPPRC label suffices, and a separate cumulative-quality-loss resource is redundant (Proposition 2). On the algorithmic side, PADMM-MA is a proximal ADMM with NG-route relaxation whose quality-loss weight β is updated by a residual-balancing rule in place of subgradient ascent. It is coupled to the outer reallocation LP through approximate-dynamic-programming-style marginal costs.

The evaluation uses twelve Solomon-derived mandarin instances (72 feasible algorithm-instance combinations across six configurations), a scale extension to $n = 100$ customers, and an exact MIP benchmark on nine tractable subproblems (Appendix A). Three findings stand out. First, quality-cost monetisation reshapes the routing structure: PADMM-NOMAT achieves lower apparent cost at the price of roughly 7% longer routes on aggregate (with the effect concentrated on narrow-TW R- and RC-type cells) and unmonetised quality degradation. Second, on the main testbed PADMM-MA is on average $15\times$ faster than SUBGRAD-MA (23 s vs. 350 s) at statistically indistinguishable mean arrival quality (paired Wilcoxon $p = 0.077$ for $\bar{q}_{\text{arr}}$), whereas SUBGRAD-MA attains a 7.5% lower mean seven-day cost; the two heuristics therefore serve different operational niches rather than dominating each other. Third, transport temperature is the strongest quality lever among the tested sensitivity parameters.

The exact MIP benchmark (Appendix A) also locates the regime in which both PADMM-MA and SUBGRAD-MA leave substantial room against exact optimisation: on small, tight-TW, clustered instances, the shared primal machinery (four-phase repair + VND over relocate/Or-opt/2-opt) yields gaps of +42% to +84% to the MIP optimum. The most concrete direction for future work is therefore to close this small-instance gap, either via stronger primal heuristics or via an exact-pricing extension of the SPPRC. Three further directions remain open: a chance-constrained variant handling demand and temperature uncertainty; integration with dynamic driver scheduling; and an extension to multi-product cold-chain networks with heterogeneous decay kinetics.

A limitation of the present study should be stated explicitly. All results reported here are obtained on semi-real, Solomon-derived mandarin instances, together with a controlled exact-MIP benchmark and a scale extension; the framework has not yet been validated against real mandarin-distribution data or live operational scenarios. The deterministic-demand, homogeneous-fleet, and fixed-cluster assumptions, although clearly delimited and consistent with the algorithmic and theoretical scope of this paper, remain to be empirically confirmed. Validating the model on real-world datasets and operational deployments is consequently an important direction that would complement—rather than supersede—the methodological contributions established here.

Author Contributions: Conceptualisation, B.L. and Y.J.; methodology, B.L. and L.M.; software, B.L., Y.Z., and H.W.; validation, L.M., Y.S., and H.W.; formal analysis, B.L. and Y.Z.; investigation, L.M. and Y.S.; data curation, Y.Z. and H.W.; writing—original draft, B.L.; writing—review and editing, L.M., Y.Z., Y.S., H.W., and Y.J.; visualisation, Y.Z. and Y.S.; supervision, Y.J.; project administration, B.L. and Y.J.; funding acquisition, B.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research has been supported by the National Innovation and Entrepreneurship Training Program for College Students (No. 202510307092) and the Fundamental Research Funds for the Central Universities (NAU: SKYD2026035).

Data Availability Statement: The Solomon VRPTW benchmark instances used as the geometric base of all experiments are publicly available at <https://www.sintef.no/projectweb/top/vrptw/solomon-benchmark/> (accessed on 28 June 2026) [24]. The mandarin-specific parameter calibrations, the modified Solomon-derived instances, the per-instance experimental results in MATLAB .mat format, the MATLAB implementation of PADMM-MA, and the baseline algorithms generated for this study are available from the corresponding author upon reasonable request.

Acknowledgments: During the preparation of this manuscript, the authors used Claude (Anthropic) for the following limited purposes: (i) English-language editing and copy-editing of author-drafted text; (ii) verification of bibliographic metadata and DOI consistency for the reference list; (iii) consistency checks on mathematical notation and L^AT_EX formatting; and (iv) drafting of approximately 400 words of transition sentences in the Literature Review (Section 2), the factual content of which was independently verified by the authors against the cited primary sources. All mathematical formulations, theoretical results (Propositions 1–4), algorithm design (PADMM-MA, Algorithms 1 and 2), computational experiments, and substantive scientific conclusions are the original work of the human authors. In particular, no AI tool was used to generate, derive, or verify any of the propositions, proofs, or algorithm pseudocode in this paper. The authors have reviewed and edited all AI-assisted output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A. Exact MIP Benchmark on Small Instances

Appendix A.1. Setup

To place both heuristics in reference to an exact optimum where it can be computed, we isolate the single-DC inner VRPTW of three Solomon-derived instances (C101, R101, RC101) on Day 1 of the rolling horizon, at three reduced scales corresponding to $|S_d| \in \{7-8, 8-12, 9-11\}$ effective customer nodes per DC after geographic clustering. We solve each as a time-expanded mixed-integer programme using MATLAB's `intlinprog` (HiGHS backend, default tolerance). The heuristics (PADMM-MA and SUBGRAD-MA) are run on exactly the same subproblem instances, with identical Lagrangian profits and quality-loss weights. Table A1 reports the comparison.

Table A1. Optimality gap comparison: MIP exact vs. heuristics on small instances (single DC subproblem, Day 1). TL = MIP aborted at solver time limit; heuristic cells are correspondingly blank because the MIP optimum against which the gap would be computed is not available. The two RC101 rows with $|\mathcal{S}_d| = 9$ correspond to two different DC subproblems of the same instance (DC3 and DC1, respectively), distinguished here because the k -means three-DC clustering of RC101's first 50 customers yields more than one DC of size 9 across the experiments listed.

Instance	$ \mathcal{S}_d $	MIP Exact			PADMM-MA			SUBGRAD-MA		
		Cost	Time	Gap%	Cost	Time	$\Delta\%$	Cost	Time	$\Delta\%$
C101	7	45	1.9	0.0	82	0.2	+82.2	82	0.1	+82.2
R101	7	158	0.4	0.0	291	0.1	+84.2	280	0.0	+77.2
RC101	8	150	1.7	0.0	225	0.1	+50.0	213	0.0	+42.0
C101	12	146	2.9	0.0	213	0.1	+45.9	218	0.1	+49.3
R101	8	280	1.5	0.0	280	0.0	0.0	280	0.0	0.0
RC101 (DC3)	9	183	2.3	0.0	183	0.0	0.0	183	0.1	0.0
C101	11	108	2.5	0.0	153	0.1	+41.7	153	0.1	+41.7
R101	9	TL	0.9	TL	—	—	—	—	—	—
RC101 (DC1)	9	163	2.1	0.0	180	0.0	+10.4	180	0.0	+10.4

Appendix A.2. What the Benchmark Shows

The exact MIP solves 8 of the 9 size–instance pairs to proven optimality within 3 s; the remaining pair (R101 at $|\mathcal{S}_d| = 9$, the largest scale tested for this instance) was aborted at the solver's time limit. On the eight solvable pairs, the heuristic gaps to the MIP optimum fall into two distinct regimes. In the *tight-TW, clustered* regime (C101 at all three scales, R101 at the smallest scale), both heuristics leave large gaps relative to the MIP, ranging from +41.7% to +84.2%. In the *wider-TW* regime (R101 at $|\mathcal{S}_d| = 8$, RC101 at $|\mathcal{S}_d| = 8$, and the two RC101 DC subproblems at $|\mathcal{S}_d| = 9$, comprising five size–instance pairs of which one—R101 at $|\mathcal{S}_d| = 9$ —was aborted at the MIP solver's time limit), both heuristics match the exact optimum ($\Delta = 0.0\%$) on two of the four solvable pairs (R101 at $|\mathcal{S}_d| = 8$ and RC101 DC3 at $|\mathcal{S}_d| = 9$), with the remaining two at gaps of roughly +10% (RC101 DC1 at $|\mathcal{S}_d| = 9$) and +42–50% (RC101 at $|\mathcal{S}_d| = 8$).

Appendix A.3. Why This Result Is Reported As-Is

The large gaps on tight small instances arise because the primal upper bound in both heuristics is produced by the same combination of a four-phase repair and a VND local search over relocate, Or-opt, and 2-opt moves (Section 4); on small, clustered, time-window-tight instances, this move set does not systematically recover the exact-MIP structure. This is a property of the primal machinery shared by PADMM-MA and SUBGRAD-MA alike, not a deficiency introduced by the residual-balancing rule, which operates on the dual side. It is also consistent with the main testbed's classification, where RC-type (mixed, clustered) instances show the highest absolute costs for both heuristics (Table 3).

Appendix A.4. Role of This Benchmark in the Paper's Claims

The benchmark is included not to claim that PADMM-MA closes the gap to the exact optimum on small instances—it does not, uniformly—but to make the operating regime of the proposed heuristic explicit. On wider-TW instances at 50 and 100 customers, PADMM-MA delivers near-identical arrival quality to SUBGRAD-MA at a $15\times$ speed-up (Section 5.2); on exactly the small, tight, clustered regime where the route-level VND is structurally weak, the MIP benchmark exposes this as a shared limitation. Closing the small-instance gap would require either an exact-pricing extension of the SPPRC or replacing the VND with a

broader primal heuristic; this is one of the concrete directions for future work flagged in the conclusion.

Appendix B. Implementation Details of the Four-Phase Repair and VND

Appendix B.1. Feasibility Contract of the Four-Phase Repair

The four-phase repair is stated as a contract with two regimes. During the procedure, intermediate route matrices may temporarily violate the capacity bound (Equation (3), constraint (5)), the time-window bound (constraint (7)), and the visit-once requirement (constraint (4)): intra- and inter-route deduplication can leave customers unserved, and inter-route reinsertion can overshoot a load or a time window before a subsequent move corrects it. These transient violations are tolerated because the repair is not required to maintain feasibility step by step, only to deliver a feasible incumbent at termination. On termination, the returned route set satisfies (i) capacity feasibility for every vehicle, (ii) time-window feasibility at every visited customer under the Proposition 1-tightened windows, and (iii) the visit-once requirement over the customer set $\mathcal{S}_d$. The regret-2 insertion phase is the mechanism that enforces (iii): it places every still-unserved customer at the position minimising the incremental monetised objective (3) given (i) and (ii). Suppose regret-2 insertion cannot place some customer j , because no remaining position admits j under capacity and the tightened time window. The repair then returns without j , and the PADMM-MA outer loop absorbs the violation via the dual update in Step 4. Specifically, the visit count $n_j = 0$ drives $\lambda_j^{k+1} = \lambda_j^k + \rho^k$ upward. This increases the profit of a route that serves j in the next SPPRC call, so the visit-once constraint is recovered on a subsequent iteration rather than within the current repair.

Appendix B.2. Four-Phase Construction and VND Template

The heuristic repair (Step 3 of Algorithm 1) constructs a feasible VRPTW solution from possibly overlapping route fragments via a four-phase procedure: initialisation from NG-route fragments, intra-route deduplication, inter-route deduplication (retaining the cheapest-to-remove copy), and regret-2 insertion of unserved customers. All route-matrix manipulations use an extract–modify–write pattern to prevent index corruption from shift operations on wide-TW instances. A subsequent VND local search applies relocate, Or-opt, and 2-opt operators in that fixed order, re-starting at the first neighbourhood after every improvement. Formally, at each pass, the current solution is first improved to a local optimum under the relocate neighbourhood (single-customer removal-and-reinsertion across all route pairs) until no improving move remains; control then passes to Or-opt (segments of length two and three re-inserted in all feasible positions), again iterated to local optimality; and finally to 2-opt (intra-route arc-pair reversals). Whenever any neighbourhood yields a strict cost decrease, the search restarts from the relocated neighbourhood in the next pass, following the standard variable-neighbourhood-descent template. The VND terminates when a full pass through all three neighbourhoods produces no improvement, at which point the incumbent is locally optimal with respect to the union of the three move sets. All moves are evaluated against the monetised inner-VRPTW objective (3), i.e., distance cost plus quality-loss cost, so that an increase in travel distance that is more than offset by reduced quality loss is accepted as an improvement.

References

- Costa, L.; Contardo, C.; Desaulniers, G. Exact branch-price-and-cut algorithms for vehicle routing. *Transp. Sci.* **2019**, *53*, 946–985. <https://doi.org/10.1287/trsc.2018.0878>.
- Pessoa, A.; Sadykov, R.; Uchoa, E.; Vanderbeck, F. A generic exact solver for vehicle routing and related problems. *Math. Program.* **2020**, *183*, 483–523. <https://doi.org/10.1007/s10107-020-01523-z>.

3. Baldacci, R.; Mingozzi, A.; Roberti, R. New route relaxation and pricing strategies for the vehicle routing problem. *Oper. Res.* **2011**, *59*, 1269–1283. <https://doi.org/10.1287/opre.1110.0975>.
4. Roberti, R.; Mingozzi, A. Dynamic ng-path relaxation for the delivery man problem. *Transp. Sci.* **2014**, *48*, 413–424. <https://doi.org/10.1287/trsc.2013.0474>.
5. Castellucci, P.B.; Coelho, L.C.; Darvish, M. A new branch-and-Benders-cut algorithm for the time-dependent vehicle routing problem. *Expert Syst. Appl.* **2025**, *265*, 125996. <https://doi.org/10.1016/j.eswa.2024.125996>.
6. Luo, H.; Ma, T.; Li, Z. An exact branch-price-and-cut algorithm for the time-dependent cold chain pickup and delivery problem with incompatibility constraints. *Comput. Oper. Res.* **2025**, *178*, 107007. <https://doi.org/10.1016/j.cor.2025.107007>.
7. 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. <https://doi.org/10.1561/22000000016>.
8. Zhang, J.; Luo, Z.Q. A proximal alternating direction method of multipliers for linearly constrained nonconvex minimization. *SIAM J. Optim.* **2020**, *30*, 2272–2302. <https://doi.org/10.1137/19M1242276>.
9. Yao, Y.; Zhu, X.; Dong, H.; Wu, S.; Wu, H.; Tong, L.C.; Zhou, X. ADMM-based problem decomposition scheme for vehicle routing problem with time windows. *Transp. Res. Part B Methodol.* **2019**, *129*, 156–174. <https://doi.org/10.1016/j.trb.2019.09.009>.
10. Chi, J.; He, S.; Zhang, Y. Improved ADMM-based approach for optimizing intercity express transportation networks: A novel dual decomposition strategy with partial retention of coupling constraints. *Transp. Res. Part E Logist. Transp. Rev.* **2024**, *192*, 103756. <https://doi.org/10.1016/j.tre.2024.103756>.
11. Xia, J.; He, Z.; Wang, S.; Liu, S.; Zhang, S. Mean-standard-deviation-based electric vehicle routing problem with time windows using Lagrangian relaxation and extended alternating direction method of multipliers-based decomposition algorithm. *Soft Comput.* **2024**, *28*, 7139–7160. <https://doi.org/10.1007/s00500-023-09599-3>.
12. Accorsi, R.; Gallo, A.; Manzini, R. A climate-driven decision-support model for the distribution of perishable products. *J. Clean. Prod.* **2017**, *165*, 917–929. <https://doi.org/10.1016/j.jclepro.2017.07.170>.
13. Qi, B.; Li, G. The evolution of the cold chain logistics vehicle routing problem: A bibliometric and visualization review. *Digit. Transp. Saf.* **2024**, *3*, 92–114. <https://doi.org/10.48130/dts-0024-0010>.
14. Osvald, A.; Stirn, L.Z. A vehicle routing algorithm for the distribution of fresh vegetables and similar perishable food. *J. Food Eng.* **2008**, *85*, 285–295. <https://doi.org/10.1016/j.jfoodeng.2007.07.008>.
15. Lado, J.; Rodrigo, M.J.; Zacarías, L. Maturity indicators and citrus fruit quality. *Stewart Postharvest Rev.* **2014**, *10*, 1–6.
16. Tietel, Z.; Lewinsohn, E.; Fallik, E.; Porat, R. Importance of storage temperatures in maintaining flavor and quality of mandarins. *Postharvest Biol. Technol.* **2012**, *64*, 175–182. <https://doi.org/10.1016/j.postharvbio.2011.07.009>.
17. Lin, N.; Kanellopoulos, A.; Akkerman, R.; Zhang, J.; Ruan, J. Vehicle routing in precooling logistics with dynamic temperature-dependent product quality decay. *Eur. J. Oper. Res.* **2025**, *321*, 407–427. <https://doi.org/10.1016/j.ejor.2024.09.041>.
18. Liu, Y.; Tao, F.; Zhu, R. Vehicle-routing problem for low-carbon cold chain logistics based on the idea of cost-benefit. *Complex Intell. Syst.* **2025**, *11*, 157. <https://doi.org/10.1007/s40747-024-01756-x>.
19. Hou, J.; Liu, Y.; Zhang, Y.; Hu, J.; Zhang, A. Optimization of urban cold chain logistics path considering carbon constraints. *Pol. J. Environ. Stud.* **2025**, *34*, 705–717. <https://doi.org/10.15244/pjoes/186444>.
20. Hou, J.; Liu, Y.; Liu, S.; Cai, C. Optimization of multi-depot semi-open cold-chain logistics paths under time-dependent road networks. *Int. J. Comput. Intell. Syst.* **2025**, *18*, 289. <https://doi.org/10.1007/s44196-025-01007-y>.
21. Beshara, S.; Kassem, A.; Fors, H.; Harraz, N. A comprehensive review and mapping citrus supply chains from a sustainability perspective across the European Union, Middle East, and Africa. *Sustainability* **2024**, *16*, 8582. <https://doi.org/10.3390/su16198582>.
22. Coelho, L.C.; Cordeau, J.F.; Laporte, G. Thirty years of inventory routing. *Transp. Sci.* **2014**, *48*, 1–19. <https://doi.org/10.1287/trsc.2013.0472>.
23. Chen, X.; Jiu, Y.; Hu, D. Two-echelon location-routing problem of perishable products based on the integrated mode of in-store pick-up and delivery. *Transp. Res. Rec.* **2024**, *2678*, 622–636. <https://doi.org/10.1177/03611981231218008>.
24. Solomon, M.M. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Oper. Res.* **1987**, *35*, 254–265. <https://doi.org/10.1287/opre.35.2.254>.
25. Powell, W.B. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed.; Wiley: Hoboken, NJ, USA, 2011.

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