

Article

A Systems-Based Model of Platform-Enabled Freight Orchestration for Cross-Border E-Commerce Fulfillment

Shucheng Fan *  and Shaochuan Fu

School of Economics and Management, Beijing Jiaotong University, Beijing 100044, China

* Correspondence: scfan@bjtu.edu.cn

Abstract

Cross-border e-commerce fulfillment depends on coordinated inland container movements across factories, inland container depots (ICDs), and port gateways, yet many container trucking operations still follow synchronous one-truck-one-order execution. This study models the fulfillment network as a platform-enabled socio-technical transportation system in which the ICD acts as a digital–physical coordination node for spatiotemporal decoupling. A drop–buffer–pick task architecture is developed to represent direct execution, relay execution, and delayed dispatch, and a mixed-integer linear programming (MILP) model optimizes task assignment and tractor sequencing under loading-time, port cutoff, inventory, and working-time constraints. In the certified-optimal 10-order instance, gross positive cost decreases from CNY 27,540 to CNY 19,915 (−27.7%); after applying the same post hoc coordination-credit accounting rule, net total fulfillment cost decreases to CNY 18,734 (−32.0%). The 10 orders are served with five tractors under the tested platform configuration, compared with 10 tractors under the restricted benchmark. To address sustainability explicitly, the analysis also reports distance-based emissions and energy-use proxies; the proposed schedule lowers cost and fleet deployment but increases total mileage, showing that economic efficiency and emissions performance do not automatically move together. The evidence is a deterministic baseline for later stochastic, mixed import/export, and collaborative-platform extensions.

Keywords: socio-technical systems; transportation systems; logistics systems; systems engineering; digital freight platform; inland container depot; spatiotemporal decoupling; mixed-integer linear programming



Academic Editors: Tatjana Stanivuk, Jelena Žanić Mikuličić, Ladislav Stazić and Branko Lalić

Received: 2 April 2026

Revised: 30 April 2026

Accepted: 13 May 2026

Published: 17 May 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 Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Cross-border e-commerce has shifted competition from online market access toward fulfillment capability. Orders must be consolidated, moved inland, connected to port gateways, and delivered within promised service windows at a sustainable cost. The upstream container trucking layer, however, remains fragmented: factory loading times are uncertain, port cutoff times are rigid, tractor capacity is unevenly distributed, and empty equipment availability varies across locations. Traditional synchronous one-truck-one-order execution is easy to manage, but it often ties tractors to non-productive waiting and limits revenue productivity.

In maritime-linked e-commerce logistics, this efficiency problem is also a sustainability problem. Drayage mileage, empty repositioning, and port-facing detours translate directly into fuel use, energy consumption, and emissions exposure. For this reason, the paper treats environmental indicators as complementary system outcomes: the model remains

cost-first, but the evaluation reports distance-based emissions and energy-use proxies so that operational efficiency is not confused with automatic environmental improvement.

Digital freight platforms create a different operating logic. By combining order visibility, vehicle availability, and rule-based dispatching, a platform can convert isolated order fulfillment into pooled task orchestration. In this context, an ICD is not only a storage yard. When digitally coordinated, it becomes a digital–physical node through which a tractor can drop a trailer, leave for another task, and let the trailer be picked up later when port timing, loading completion, or capacity conditions are favorable.

Existing research on digital commerce emphasizes consumers, platforms, business models, and last-mile service. Freight platform research emphasizes matching, visibility, and routing, while drayage and dry-port studies provide scheduling tools. Yet the profitability implications of ICD-enabled relay execution in containerized cross-border fulfillment remain insufficiently examined. This paper therefore asks when a platform should choose ICD relay instead of direct execution, how much task pooling can compress fleet needs, and what trade-offs emerge between profitability, schedule tightness, and operational slack.

1.1. Literature Review

Research on cross-border e-commerce has expanded rapidly, but logistics-related work remains uneven. Recent studies link platform service capability, enterprise digital transformation, logistics service innovation, and return logistics to firm performance and operating strategy [1–6]. They show that fulfillment capability is strategically important, but they rarely explain how upstream containerized freight tasks should be orchestrated.

A second stream examines fulfillment service provision and logistics sharing on e-commerce platforms. Platform logistics integration, shared fulfillment, and co-opetition can reshape costs, incentives, and profitability [7–12]. This work mainly studies pricing, service sharing, or collaboration at the strategic level rather than tractor-trailer execution, ICD buffers, and port-facing deadlines.

A third stream studies digital freight platforms and algorithmic coordination. Digital platforms provide visibility, matching, optimization, and analytics, and recent work integrates matching and routing to reduce empty trips [13–18]. However, this literature focuses mainly on spot matching, generic exchanges, or hub-based freight rather than containerized fulfillment in which an ICD functions as a relay node.

A fourth stream comes from operations research on drayage, transfer-enabled routing, dry ports, and inland container systems [19–24]. These studies provide foundations for precedence modeling, transfer nodes, time windows, and inland container coordination, but they usually frame the problem as transport-cost minimization or scheduling feasibility. They seldom theorize the digital platform as an orchestrator that improves profitability by reducing waiting cost, recombining tasks, and increasing tractor productivity.

From a socio-technical systems perspective, operational performance is a joint outcome of technical routines, organizational control, information flows, and physical work-system design. This lens is appropriate here because dispatch algorithms, platform authority, ICD buffering capacity, and tractor behavior jointly shape system performance.

1.2. Research Gap and Study Positioning

This paper positions itself at the intersection of these literature streams by modeling the ICD as a digital-physical coordination node through which a platform can pool tasks, reallocate tractor time, and improve fulfillment profitability via spatiotemporal decoupling. Although recent studies have shown that logistics mode selection and information sharing can materially reshape outcomes in cross-border e-commerce supply chains [25], the present

study is explicitly anchored in a socio-technical systems perspective, under which value depends on the joint design of digital coordination rules, physical transfer nodes, and mobile resources rather than on isolated route decisions alone [26,27]. Accordingly, there is still limited understanding of how a digital freight platform can use an ICD to reorganize container flows, choose between relay and direct execution, and improve the economic return of cross-border fulfillment through spatiotemporal decoupling.

1.3. Research Questions and Contributions

This study is guided by three research questions:

RQ1. *Under what conditions should a digital freight platform prefer relay execution via an ICD over direct execution?*

RQ2. *To what extent can task pooling and spatiotemporal decoupling improve profitability by reducing waiting-related opportunity cost, compressing fleet capacity, and absorbing repositioning mileage?*

RQ3. *What trade-offs emerge between tractor productivity, schedule tightness, and operational slack when execution shifts from isolated orders to coordinated task orchestration?*

The paper makes three contributions. First, it conceptualizes the fulfillment network as a platform-enabled socio-technical transportation system. Second, it introduces an ICD-based spatiotemporal decoupling framework representing direct execution, relay execution, ICD buffering, local empty-container reuse, and backhaul coordination. The main certified benchmark remains export-side, while Appendix A reports small MILP activation tests for mixed import/export street-turn reuse and triangular backhaul. Third, it provides a profitability- and sustainability-aware evaluation logic for inland container scheduling, emphasizing opportunity cost, orders per tractor, fleet compression, economic absorption of repositioning mileage, and distance-based emissions/energy-use proxies rather than route length alone.

1.4. Structure of the Paper

Section 2 defines the platform setting, task architecture, decision boundaries, MILP formulation, benchmark policy, and experiment design. Section 3 reports the computational results. Section 4 discusses implications and boundary conditions, and Section 5 concludes.

2. Materials and Methods

2.1. System Boundary and Network Architecture

The fulfillment network contains ICDs, factory customers, and port gateways. The digital freight platform coordinates tractors across this port–ICD–factory system to improve economic return rather than merely minimize route length.

In systems terms, the setting is a socio-technical transportation system composed of physical nodes, mobile tractors, digital coordination rules, and time-critical constraints [26,27]. The system boundary includes the port, ICD, factory customers, and platform-controlled tractors. External conditions are represented by task-release times, loading readiness, and port cutoffs. Performance depends on how the platform coordinates tasks, vehicles, data visibility, and nodes across time.

Let the physical node set be

$$\mathcal{N} = \mathcal{I} \cup \mathcal{C} \cup \mathcal{P}.$$

Here, $\mathcal{I}$ denotes ICDs, $\mathcal{C}$ factory customers, and $\mathcal{P}$ port nodes. Travel time between nodes i and j is τ_{ij} , and distance is d_{ij} .

The platform coordinates tractors $\mathcal{K}$ and orders $\mathcal{O}$ with order visibility, tractor-status visibility, and task-level dispatching authority. The model assumes centralized dispatch control; bargaining, outsourcing, and decentralized carrier acceptance are outside the boundary.

The ICD provides a buffer for time shifting and task handoff. It lets the platform decouple container-handling urgency from tractor capacity, especially when export orders face loading and port cutoff constraints or import flows create local empty-container reuse opportunities. Figure 1 summarizes the network and task-flow logic.

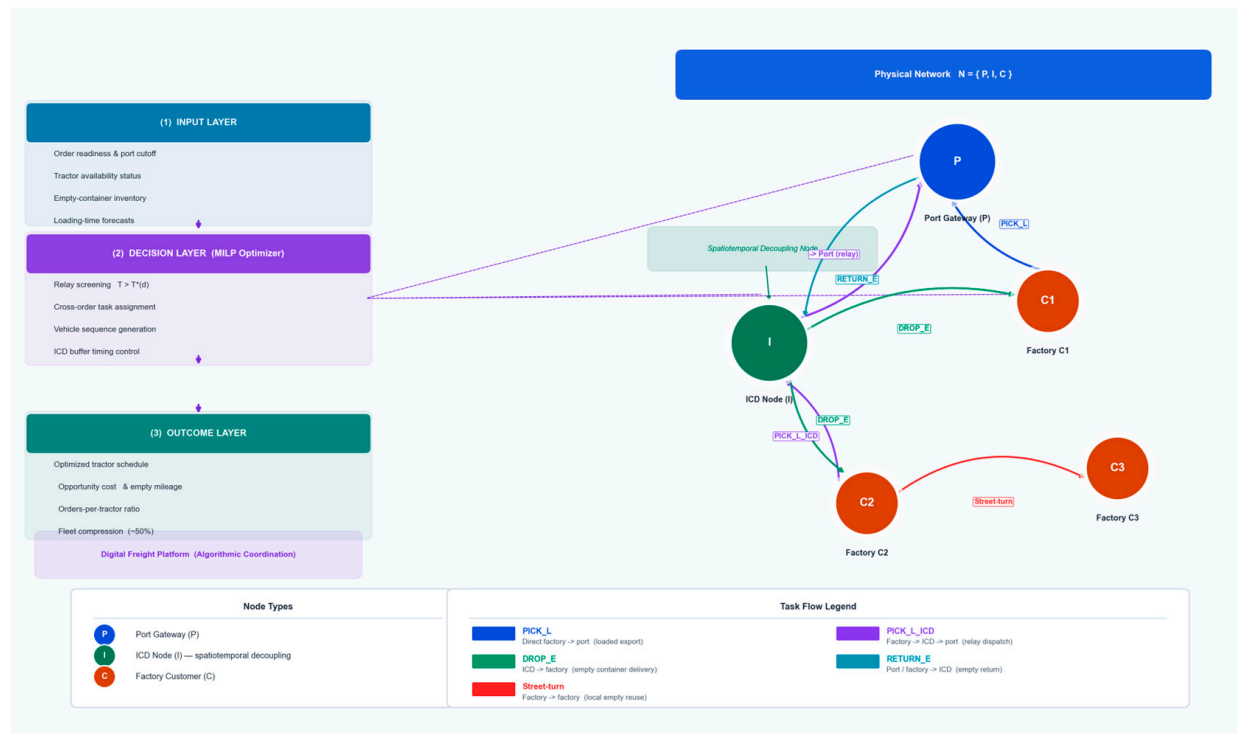


Figure 1. Platform-enabled freight network structure and task flow.

The platform overlays the physical network with an input layer for order readiness, tractor availability, empty inventory, and port cutoffs; a decision layer for relay screening, task assignment, and sequencing; and an outcome layer for optimized schedules, lower opportunity cost, higher orders-per-tractor ratios, and fleet compression.

2.2. Task Structure and Execution Modes

2.2.1. Task Atoms and Unified Operational Semantics

Orders are decomposed into executable task atoms. Each task has an origin, destination, time attributes, execution duration, and inventory effect. A tractor may execute multiple tasks, and different stages of the same order may be assigned to different tractors when precedence, timing, and inventory constraints are satisfied.

The main tasks are DROP_E for empty delivery, PICK_L for direct loaded pickup, BUFFER_L for factory-to-ICD buffering, PICK_L_ICD for ICD-to-port relay pickup, GATEIN for port completion, and import/reuse tasks such as PICK_I, DROP_I, RETURN_E, PICK_E_PORT, and TRANSFER_E_ST. Table 1 summarizes the task dictionary.

This representation turns operations into task recombination problems rather than isolated trips, supporting relay execution, drop-and-hook, street-turn reuse, and back-haul coordination.

Table 1. Task dictionary and inventory effects.

Task Type	Notation	Origin	Destination	ICD Inventory Effect	Description
Empty delivery	DROP_E	ICD i	Factory c	−1	Deliver empty container to factory
Loaded pickup (direct)	PICK_L	Factory c	Port p	0	Pick up loaded export container and drive directly to port
Loaded pickup (to ICD)	BUFFER_L	Factory c	ICD i	+1 (loaded buffer)	Pick up loaded container and return to ICD for buffering
ICD relay pickup	PICK_L_ICD	ICD i	Port p	−1 (loaded buffer)	Pick up buffered loaded container at ICD and deliver to port
Port gate-in	GATEIN	Port p	Port p	0	Complete port entry and cutoff verification
Import pickup	PICK_I	Port p	Import customer c	0	Move import loaded container from port to consignee
Import delivery/unloading	DROP_I	Import customer c	Import customer c	0	Unload the import container at the consignee site
Empty return	RETURN_E	Import customer c	ICD i	+1	Return empty container to ICD after import unloading
Port empty pickup	PICK_E_PORT	Port p	Factory c or ICD i	context-dependent	Pick up empty container at port after export/import operation
Street-turn transfer	TRANSFER_E_ST	Import customer c_1	Export customer c_2	0	Reuse empty container locally without ICD return

2.2.2. Asynchronous Task Chains

For each order $o \in \mathcal{O}$, the platform constructs a task chain $\mathcal{L}(o) \subseteq \mathcal{L}$, together with a precedence set $\mathcal{E} \subseteq \mathcal{L} \times \mathcal{L}$ that defines which tasks must be completed before subsequent tasks can start.

For a typical export order, the task chain can be represented as

$$\text{DROP}_E \rightarrow \text{Loading} \rightarrow \begin{cases} \text{PICK}_L \rightarrow \text{GATEIN}, & \text{direct port delivery,} \\ \text{BUFFER}_L \rightarrow \text{PICK}_{L,ICD} \rightarrow \text{GATEIN}, & \text{ICD buffering and delayed dispatch.} \end{cases}$$

For a typical import order, the task chain is

$$\text{PICK}_I \rightarrow \text{DROP}_I \rightarrow \text{RETURN}_E.$$

This structure enables the platform to treat the ICD as a break point at which time-critical and distance-critical components of a shipment can be separated. In particular, the factory-side “loaded container must be removed within two hours” rule and the port-side “gate-in before cutoff” rule do not need to be satisfied by the same tractor in a continuous trip.

2.2.3. Operational Modes as Task-Recombination Patterns

Operational modes are task chains. Direct export execution follows $\text{DROP}_E \rightarrow \text{Loading} \rightarrow \text{PICK}_L \rightarrow \text{GATEIN}$, while relay execution follows $\text{DROP}_E \rightarrow \text{Loading} \rightarrow \text{BUFFER}_L \rightarrow \text{PICK}_{L,ICD} \rightarrow \text{GATEIN}$. Street-turn reuse connects an import empty directly to a nearby export customer, and triangular backhaul links port-side completion with a later port empty pickup.

The main computational experiment activates direct and ICD-relay execution. Street-turn reuse and triangular backhaul are structurally supported but not part of the main certified export-side benchmark; Appendix A therefore reports small MILP activation tests for these patterns rather than using them as broad performance evidence.

2.3. System-Level Decision Boundaries

The platform's objective is not merely to reduce route length but to improve profitability through better allocation of tractor time, lower opportunity cost, lower port storage exposure, and more effective reuse of local assets. To make the operational logic interpretable, three analytical decision-boundary functions are derived. These functions are not hard constraints of the MILP; rather, they provide managerial screening rules and economic intuition for why certain orchestration patterns emerge.

2.3.1. Factory-Side Boundary: Drop-and-Hook Versus Live Load

Consider an export order for which an empty container must be delivered from the ICD to a factory and later picked up after loading. Let d denote the one-way distance between the ICD and the factory, T the factory loading time, c_{km} the unit distance cost, C_{op} the fixed dispatch-related operating cost, and γ the opportunity cost of tractor time.

If the tractor waits on site, the main economic loss is the opportunity cost of non-productive time:

$$C_{\text{wait}} = \gamma T.$$

If the tractor leaves immediately after dropping the empty unit and returns later for pickup, the platform incurs two additional empty travel legs and the associated fixed operating cost:

$$C_{\text{extra}} = 2dc_{km} + 2C_{op}.$$

The net profitability advantage of drop-and-hook over live load can therefore be expressed as

$$\Pi^{DH}(d, T) = \gamma T - (2dc_{km} + 2C_{op}).$$

The break-even loading time is

$$T^*(d) = \frac{2dc_{km} + 2C_{op}}{\gamma}.$$

If $T > T^*(d)$, releasing the tractor is economically preferable; if $T < T^*(d)$, direct waiting is preferable. This decision boundary is especially useful for platform dispatching because it converts operational experience into a computable switching rule.

2.3.2. Port-Side Boundary: ICD Buffering Versus Early Gate-In

For containers ready well before port cutoff, the platform decides whether to gate-in early or buffer at the ICD. Let t_o^{ready} be loading completion, t_o^{cut} port cutoff, Δ^{free} the free storage window, and ρ the storage rate. If the container arrives at the port at t_A^{sin} , storage exposure is

$$C_{\text{port}}(t_A^{\text{sin}}) = \rho \cdot \max\left\{0, (t_o^{\text{cut}} - t_A^{\text{sin}}) - \Delta^{\text{free}}\right\}.$$

If the unit is buffered at the ICD and dispatched at $t_B^{\text{dep},*}$, the profitability advantage is

$$\Delta\Pi_o^{\text{BUF}} = C_{\text{port}}(t_A^{\text{sin}}) + \gamma(t_B^{\text{dep},*} - t_o^{\text{ready}}) + \Delta R_o - C_{\text{back}}.$$

ICD buffering is preferred when $\Delta\Pi_0^{BUF} > 0$ and port feasibility is preserved. In the experiments, $\Delta^{free} = 72$ h and the horizon is 24 h, so port storage costs are not triggered; this boundary is analytical rather than separately stress-tested.

2.3.3. Empty Reuse Boundary: Street-Turn Versus Return-to-ICD

After an import container is unloaded, the platform may return the empty unit to the ICD or reuse it locally for a nearby export order. Let $d_{im \rightarrow b}$ be the import customer-to-ICD distance, $d_{b \rightarrow ex}$ the ICD-to-export customer distance, and $d_{im \rightarrow ex}$ the direct import-to-export distance. Let C_{icd}^{op} denote ICD handling cost, C_{st}^{op} street-turn coordination cost, C^{clr} any avoided cleaning/inspection cost, and R_{uv}^{cond} and R_{uv}^{tw} mismatch penalties. The gain is

$$\Pi_{uv}^{ST} = c_{km}[(d_{im \rightarrow b} + d_{b \rightarrow ex}) - d_{im \rightarrow ex}] + (2C_{icd}^{op} - C_{st}^{op}) + C^{clr} - (R_{uv}^{cond} + R_{uv}^{tw}).$$

Street-turn reuse is preferred when $\Pi_{uv}^{ST} > 0$ and the pair is operationally compatible. Table 2 summarizes the decision boundaries.

Table 2. Profitability-oriented decision boundaries.

Boundary	Decision	Trigger Condition	Key Variables
Factory-side: drop-and-hook vs. live load	Release tractor during factory loading or wait continuously	Prefer drop-and-hook when $\Pi^{DH}(d, T) > 0$	$T, d, \gamma, c_{km}, C_{op}$
Port-side: ICD buffering vs. early gate-in	Buffer at ICD and dispatch later or gate-in early	Prefer ICD buffering when $\Delta\Pi_0^{BUF} > 0$	$t_0^{ready}, t_0^{cut}, \rho, \Delta^{free}, C_{back}$
Empty reuse: street-turn vs. return-to-ICD	Reuse import empty locally or return to ICD	Prefer street-turn when $\Pi_{uv}^{ST} > 0$	$d_{im \rightarrow b}, d_{b \rightarrow ex}, d_{im \rightarrow ex}, R_{uv}^{cond}, R_{uv}^{tw}$

These boundaries are screening rules, not hard MILP constraints. The factory-side boundary is empirically activated in the main benchmark. Port buffering remains analytically represented under the current 72 h free-storage setting, while street-turn and triangular-backhaul patterns are tested only through the small activation examples in Appendix A.

2.4. System Optimization Model

2.4.1. Sets, Parameters, and Decision Variables

The platform optimization model is built on the following core sets:

- $\mathcal{K}$: set of tractors;
- $\mathcal{O}$: set of orders;
- $\mathcal{L}$: set of task atoms;
- $\mathcal{L}^{gate} \subseteq \mathcal{L}$: set of port-related tasks;
- $\mathcal{L}^{load} \subseteq \mathcal{L}$: set of loaded-pickup tasks triggered by factory loading completion;
- $\mathcal{E} \subseteq \mathcal{L} \times \mathcal{L}$: precedence relations between tasks;
- $\mathcal{T}$: discrete inventory-tracking time periods, used only for the empty-container inventory balance;
- $\mathcal{R}_o$: feasible execution-mode set for order o , used when an order has optional direct, relay, buffering, street-turn, or backhaul patterns.

The main parameters include travel distance d_{ij} , travel time τ_{ij} , task execution duration τ_l , task opening time t_l^{open} , task latest feasible time t_l^{close} , loading completion time t_l^{ready} , unit transport cost c_{km} , tractor time opportunity cost γ , port storage rate ρ , free port storage window Δ^{free} , delay penalty P_l , inventory bounds E_i^{init} and E_i^{max} , and maximum tractor working time T_k^{max} . Here, τ_l includes both the in-task travel time from orig(l) to dest(l) and node-level service or handling time.

The main decision variables are:

- $x_{k,l} \in \{0, 1\}$: equals 1 if tractor k executes task l ;
- $w_{l,m}^k \in \{0, 1\}$: equals 1 if tractor k executes task m immediately after task l ;
- $z_{o,r} \in \{0, 1\}$: equals 1 if execution mode r is selected for order o ;
- $b_{k,l}, e_{k,l} \in \{0, 1\}$: equal 1 if task l is respectively the first or last task on tractor k 's route;
- $s_l \geq 0$: start time of task l ;
- $c_l \geq 0$: completion time of task l ;
- $Idle_k \geq 0$: non-productive time of tractor k within the scheduling horizon;
- $S_l \geq 0$: early-arrival storage exposure for port-related task l ;
- $L_l \geq 0$: lateness of port-related task l ;
- $E_{i,t} \geq 0$: available empty-container inventory at ICD i in period t .

For reference, all principal notation is consolidated in Table 3.

Table 3. Summary of principal sets, parameters, and decision variables.

Symbol	Type	Description
K	Set	Set of tractors (vehicles)
O	Set	Set of customer orders
L	Set	Set of task atoms
Lgate	Set	Subset of port gate-in tasks
Lload	Set	Subset of loaded-pickup tasks
E	Set	Precedence pairs (l, m), where m must follow l
T	Set	Discrete time periods for inventory tracking
Ro	Set	Feasible execution modes for order o
dij	Param.	Travel distance from node i to node j (km)
τ_{ij}	Param.	Travel time from node i to node j (h)
τ_l	Param.	Execution duration of task l, including in-task travel and service/handling time (h)
tlopen	Param.	Earliest start time of task l
tlclose	Param.	Latest feasible completion time (deadline) of task l
tready	Param.	Factory loading completion time for l in Lload
ckm	Param.	Unit transport cost (CNY/km); calibrated at 3.5 CNY/km
ϕ_l	Param.	Fixed handling cost of task l (CNY); 100 CNY/task
γ	Param.	Tractor time opportunity cost (CNY/h); calibrated at 200 CNY/h
ϱ	Param.	Port storage charge rate (CNY/h); calibrated at 2.08 CNY/h (equivalent to 50 CNY/day)
Δ_{free}	Param.	Free port storage window (h); set to 72 h
Δ_{max}	Param.	Maximum allowed factory residence window (h); set to 2 h
Pl	Param.	Unit delay penalty for port task l (CNY/h)
ω_l	Param.	Priority weight for port-deadline task l
Eiinit	Param.	Initial empty-container inventory at ICD i
Eimax	Param.	Maximum empty-container capacity at ICD i
Tkmax	Param.	Maximum working time of tractor k (h); set to 24 h
Tkavail	Param.	Available duty horizon of tractor k (h)
M	Param.	Big-M constant for disjunctive linearization
$x_{k,l}$	Var. (binary)	1 if tractor k executes task l; 0 otherwise
$w_{l,m}^k$	Var. (binary)	1 if tractor k performs task m immediately after l
$z_{o,r}$	Var. (binary)	1 if order o uses execution mode r
$b_{k,l}, e_{k,l}$	Var. (binary)	Route-start and route-end indicators for tractor k
s_l	Var. (cont.)	Start time of task l
c_l	Var. (cont.)	Completion time of task l
Idlek	Var. (cont.)	Non-productive time of tractor k (h)
S_l	Var. (cont.)	Early-arrival port storage exposure of task l (h)
L_l	Var. (cont.)	Lateness of port task l (h)
$E_{i,t}$	Var. (cont.)	Empty-container inventory at ICD i in period t

In the present formulation, the core assignment remains task-atom based. Tasks that are common to all feasible execution patterns of an order are covered directly, while mode-specific tasks are activated only when the corresponding execution mode is selected. After mode activation, direct execution and relay execution are determined by tractor identity,

timing, and route continuity. When two sequential tasks of the same order are assigned to the same tractor without meaningful buffer delay, the order behaves as a direct or live-load execution. When they are separated in time or tractor identity through the ICD, the order behaves as relay-based execution.

2.4.2. Objective Function

The MILP minimizes total scheduling cost:

$$\min Z = Z_{mov} + Z_{op} + Z_{opp} + Z_{stor} + Z_{vio}.$$

Movement cost includes in-task mileage and inter-task repositioning:

$$Z_{mov} = \sum_{k \in \mathcal{K}} \sum_{l \in \mathcal{L}} c_{km} d_l x_{k,l} + \sum_{k \in \mathcal{K}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{L}} c_{km} d_{\text{dest}(l), \text{orig}(m)} w_{l,m}^k.$$

Operating, opportunity, storage, and lateness costs are

$$\begin{aligned} Z_{op} &= \sum_{l \in \mathcal{L}} \phi_l \sum_{k \in \mathcal{K}} x_{k,l}, \\ Z_{opp} &= \sum_{k \in \mathcal{K}} \gamma \text{Idle}_k, \\ Z_{stor} &= \sum_{l \in \mathcal{L}^{gate}} \rho S_l, \\ Z_{vio} &= \sum_{l \in \mathcal{L}^{gate}} \omega_l P_l L_l. \end{aligned}$$

The results distinguish gross positive cost from net total fulfillment cost, which subtracts a post hoc near-pickup coordination credit when the realized schedule avoids a redundant dispatch leg. This credit is not a separate optimization objective and is applied consistently to benchmark and proposed schedules.

2.4.3. Core Constraints

Task Coverage

Each mandatory task must be executed by exactly one tractor:

$$\sum_{k \in \mathcal{K}} x_{k,l} = 1, \forall l \in \mathcal{L}.$$

For order-specific mode alternatives, coverage is conditional on the selected execution pattern. Let $\mathcal{R}_o$ denote the feasible execution-mode set of order o , and let $z_{o,r}$ equal 1 if mode r is selected. The model selects one execution mode for each order:

$$\sum_{r \in \mathcal{R}_o} z_{o,r} = 1, \forall o \in \mathcal{O}.$$

For mode-specific task atoms $l \in \mathcal{L}_{o,r}^{mode}$, task coverage is activated only under the selected mode:

$$\sum_{k \in \mathcal{K}} x_{k,l} = z_{o,r}, \forall l \in \mathcal{L}_{o,r}^{mode}, \forall o \in \mathcal{O}, \forall r \in \mathcal{R}_o.$$

Common task atoms that appear in all feasible modes retain the direct coverage constraint above. This convention prevents direct, relay, buffering, street-turn, and backhaul alternatives from being forced simultaneously.

Vehicle Continuity

A tractor can only continue from a task that it has executed, and each task can have at most one immediate successor and one immediate predecessor on the same tractor route:

$$\begin{aligned} w_{l,m}^k &\leq x_{k,l}, w_{l,m}^k \leq x_{k,m}, \forall k \in \mathcal{K}, \forall l, m \in \mathcal{L}. \\ \sum_{m \in \mathcal{L}} w_{l,m}^k &\leq 1, \sum_{m \in \mathcal{L}} w_{m,l}^k \leq 1, \forall k \in \mathcal{K}, \forall l \in \mathcal{L}. \end{aligned}$$

To ensure that tasks assigned to the same tractor form one connected route rather than disconnected fragments, route-start and route-end indicators are used:

$$\begin{aligned} \sum_{m \in \mathcal{L}} w_{l,m}^k + e_{k,l} &= x_{k,l}, \forall k \in \mathcal{K}, \forall l \in \mathcal{L}, \\ \sum_{m \in \mathcal{L}} w_{m,l}^k + b_{k,l} &= x_{k,l}, \forall k \in \mathcal{K}, \forall l \in \mathcal{L}, \\ \sum_{l \in \mathcal{L}} b_{k,l} &\leq 1, \sum_{l \in \mathcal{L}} e_{k,l} \leq 1, \forall k \in \mathcal{K}. \end{aligned}$$

Temporal Feasibility Between Consecutive Tasks

If tractor k performs task m after task l , the start time of m must allow for completion of l and travel to the origin of m :

$$s_m \geq c_l + \tau_{\text{dest}(l), \text{orig}(m)} - M(1 - w_{l,m}^k), \forall k, l, m.$$

Task completion time is defined as

$$c_l = s_l + \tau_l, \forall l \in \mathcal{L}.$$

Order-Level Precedence

For any precedence pair $(l, m) \in \mathcal{E}$, task m cannot start before task l is completed:

$$s_m \geq c_l, \forall (l, m) \in \mathcal{E}.$$

This constraint ensures physical consistency for multi-stage orders, including export chains, import chains, and ICD relay patterns.

Factory Loading Completion and Two-Hour Removal Rule

For loaded export pickup tasks, the tractor may only begin pickup after loading has been completed and must remove the loaded container within the allowed residence window:

$$t_l^{\text{ready}} \leq s_l \leq t_l^{\text{ready}} + \Delta^{\text{max}}, \forall l \in \mathcal{L}^{\text{load}}.$$

Here, Δ^{max} is typically set to 2 h.

Port Cutoff Feasibility

All port-related tasks must satisfy the latest gate-in deadline:

$$c_l \leq t_l^{\text{close}}, \forall l \in \mathcal{L}^{\text{gate}}.$$

In the present experiments, port cutoff is treated as a hard feasibility requirement; therefore $L_l = 0$ for all feasible schedules. The lateness variable and penalty term are retained to allow soft-deadline extensions in future applications.

Early-Arrival Storage Exposure

To capture the cost of arriving at the port too early, storage exposure is defined as

$$S_l \geq (t_l^{close} - \Delta^{free}) - c_l, S_l \geq 0, \forall l \in \mathcal{L}^{gate}.$$

Lateness at the Port

Late arrival is captured by

$$L_l \geq c_l - t_l^{close}, L_l \geq 0, \forall l \in \mathcal{L}^{gate}.$$

ICD Empty-Container Inventory Balance

The ICD must satisfy inventory flow balance over time:

$$E_{i,t} = E_{i,t-1} + \Delta E_{i,t}^+ - \Delta E_{i,t}^-, \forall i \in \mathcal{I}, \forall t \in \mathcal{T}.$$

subject to

$$0 \leq E_{i,t} \leq E_i^{max}, \forall i \in \mathcal{I}, \forall t \in \mathcal{T}.$$

Here, $\Delta E_{i,t}^+$ and $\Delta E_{i,t}^-$ denote task-induced empty-container inflow and outflow at ICD i during period t .

Tractor Working-Time Limit

To preserve operational feasibility and labor compliance, each tractor must satisfy a maximum working-time constraint:

$$\sum_{l \in \mathcal{L}} \tau_l x_{k,l} + \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{L}} \tau_{\text{dest}(l), \text{orig}(m)} w_{l,m}^k \leq T_k^{max}, \forall k \in \mathcal{K}.$$

Idle Time Definition

The opportunity-cost variable $Idle_k$ is defined as the unused portion of the tractor's available duty horizon:

$$Idle_k \geq T_k^{avail} - \sum_{l \in \mathcal{L}} \tau_l x_{k,l} - \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{L}} \tau_{\text{dest}(l), \text{orig}(m)} w_{l,m}^k, \forall k \in \mathcal{K},$$

$$Idle_k \geq 0, \forall k \in \mathcal{K}.$$

This formulation allows the model to penalize non-productive time directly and thereby to represent profitability improvement through better tractor utilization.

2.5. Benchmark Policy and Performance Indicators

To evaluate the managerial value of platform orchestration, the optimized solution is compared with a conventional benchmark in which each order is executed in a largely isolated manner. This benchmark reflects a practical one-truck-per-order logic under which direct task continuity is preferred and cross-order relay opportunities are not systematically exploited.

Operationally, the benchmark is defined as a continuous tractor-to-order coupling policy. Each tractor is pre-assigned to a single order and may execute only tasks belonging to that order. No cross-order sequencing is allowed. Relay through the ICD is disabled in the benchmark, so an order cannot be split into factory-side and ICD-side legs served by different tractors. Drop-and-hook splitting across different tractors is likewise not permitted: if a tractor performs the empty-container delivery for an order, that same tractor remains responsible for the subsequent loaded movement until the order is completed. For

import tasks, the empty unit is returned to the ICD under standard continuity rather than reused locally.

The benchmark therefore serves as a meaningful managerial reference for quantifying the value of task pooling and relay execution. It is not intended to be a weak heuristic baseline, but a restricted MILP that enforces isolated order execution explicitly.

The paper uses the following performance indicators.

2.5.1. Total Fulfillment Cost

The overall economic performance is measured by the objective value Z , together with its internal composition into transport cost, operating cost, opportunity cost, port storage cost, and violation cost.

2.5.2. Orders-per-Tractor Ratio

To evaluate tractor productivity, the study reports the orders-per-tractor ratio:

$$OVR = \frac{|\mathcal{O}|}{|\mathcal{K}^{used}|}.$$

Here, $|\mathcal{K}^{used}|$ is the number of tractors actually deployed.

2.5.3. Capacity Compression Rate

To quantify how much fleet requirement is reduced relative to one-truck-per-order logic, the study uses

$$CCR = 1 - \frac{|\mathcal{K}^{used}|}{|\mathcal{O}|}.$$

A higher CCR indicates stronger fleet compression through task pooling.

2.5.4. Average Daily Mileage per Tractor

The mileage productivity of tractors is measured by

$$\bar{D} = \frac{D^{tot}}{|\mathcal{K}^{used}|}.$$

Here, D^{tot} is total fleet mileage.

2.5.5. Empty-Mileage Ratio

To measure non-revenue travel, the empty-mileage ratio is defined as

$$\eta^{empty} = \frac{D^{rep} + D^{return}}{D^{tot}}.$$

Here, D^{rep} denotes repositioning mileage and D^{return} denotes return-to-base mileage.

2.5.6. Distance-Based Environmental Proxy

To make the environmental dimension explicit, the study also reports a transparent distance-based proxy.

Under identical vehicle and fuel assumptions, distance-based emissions and energy-use proxies are measured as:

$$E^{dist} = D^{tot} e_k m, Q^{dist} = D^{tot} q_k m.$$

Here, $e_k m$ and $q_k m$ denote generic per-kilometer emissions and energy-use factors. These indicators are not calibrated carbon inventories, but they make visible whether a cost-efficient schedule also reduces or increases mileage-related environmental exposure.

2.5.7. Tractor Utilization and Schedule-Preserving Slack

The busy-time utilization of tractor k is measured as

$$U_k = \frac{T_k^{busy}}{T_k^{avail}}.$$

Here, T_k^{busy} includes task execution duration and inter-task travel time. The complement of this measure reflects available slack and is used to interpret the trade-off between high utilization and schedule resilience.

Taken together, these indicators allow the study to assess profitability improvement not only as direct cost reduction, but also as higher tractor productivity, stronger fleet compression, better balance between workload intensity and operational flexibility, and the environmental implications of additional or reduced mileage.

2.6. Computational Experiment Design and Data Generation

Computational experiments use synthetic but operationally calibrated instances generated from port, ICD, and factory coordinates, loading-ready times, port cutoffs, and tractor availability. The instances were generated with fixed seed = 42 on a network with one port, one ICD, and five factory nodes; the ICD–port distance is approximately 80 km, ICD–factory distances range from 25 to 60 km, factory–port distances range from 35 to 90 km, and loading times are drawn from [1.5, 4.5] h. The main benchmark remains export-side to isolate ICD relay effects; Appendix A separately tests street-turn and triangular-backhaul activation.

The MILP is solved with Gurobi 10.0.1 using time limits of 120 s (5 orders), 300 s (10 orders), 600 s (15 orders), and 900 s (20 orders). The 5-order and 10-order cases are solved to certified global optimality. The 15-order and 20-order cases do not reach certified optimality within the run budget; no incumbent solution from these runs is used as quantitative performance evidence.

Key parameters are: $c_{km} = 3.5$ CNY/km, $\phi_l = 100$ CNY, $\gamma = 200$ CNY/h, $\rho = 2.08$ CNY/h, $\Delta^{free} = 72$ h, $\Delta^{max} = 2$ h, $T_k^{max} = 24$ h, and $P_l = 500$ CNY/h for GATEIN tasks. The one-truck-per-order benchmark is implemented as a restricted MILP enforcing the continuity rules in Section 2.5.

3. Results

This section reports the computational results of the proposed platform-enabled freight orchestration model under deterministic perfect-information conditions. Four layers of evidence are presented. First, a detailed 10-order instance is used to interpret the internal structure of the optimized solution, including cost composition, task sequencing, mileage allocation, and workload balance. Second, a scale-up analysis over 5, 10, 15, and 20 orders is used to evaluate computational growth and the limits of exact tractability under increasing task density. Third, Section 3.1 reports comparative-static sensitivity of the analytical relay boundary so that the switching logic is not tied to a single ad hoc parameter setting. Fourth, a small post-optimal perturbation screen is used to identify where the deterministic schedule is timing-sensitive; this screen is not a stochastic MILP, but it clarifies the robustness limits of the certified schedule.

3.1. Relay-Versus-Direct Execution Boundary

The first result concerns the economic boundary between direct execution and ICD-enabled relay execution. Figure 2 visualizes the break-even relationship between the distance from the ICD to the factory and the loading time required for the relay strategy to become economically preferable. The threshold increases linearly with distance, indicating that relay execution becomes harder to justify as the detour associated with the ICD grows. Under the calibrated analytical expression in Section 2.3.1, the break-even loading time is approximately 2.05 h at 30 km, 2.75 h at 50 km, and 3.80 h at 80 km.




Figure 2. Relay-versus-direct execution boundary under ICD-based platform dispatching.

This result has an immediate platform interpretation. When a factory is relatively close to the ICD, even a moderate loading delay can justify drop-and-hook execution because the tractor can be released and reassigned before the additional travel cost becomes dominant. By contrast, when the factory is farther away, direct execution remains preferable unless the loading dwell time is sufficiently long to offset the extra empty legs. The boundary therefore converts what is often treated as tacit dispatching experience into a transparent switching rule. Rather than using relay execution indiscriminately, the platform can apply ICD-based handoff selectively under conditions in which tractor time is more valuable than the incremental distance cost.

3.1.1. Comparative-Static Sensitivity of the Relay Boundary

To clarify that the relay boundary is not tied to a single arbitrary calibration, it is useful to read $T^*(d)$ as a comparative-static expression. At a representative distance of 50 km, lowering the tractor-time opportunity cost from 200 CNY/h to 150 CNY/h raises the break-even loading time from 2.75 h to 3.67 h. Conversely, raising γ to 250 CNY/h lowers the threshold to 2.20 h. Holding $\gamma = 200$ CNY/h constant, increasing the unit distance cost from 3.0 to 4.0 CNY/km shifts $T^*(50)$ from 2.50 h to 3.00 h. Table 4 reports the break-even loading times used in this comparative-static screen.

Table 4. Comparative-static sensitivity of the relay boundary at 50 km.

Scenario	Tractor-Time Value γ (CNY/h)	Unit Distance Cost c_{km} (CNY/km)	Break-Even Loading Time at 50 km
Lower time value	150	3.5	3.67 h
Lower distance cost	200	3.0	2.50 h
Baseline calibration	200	3.5	2.75 h
Higher distance cost	200	4.0	3.00 h
Higher time value	250	3.5	2.20 h

These comparative-static results show a transparent directional pattern. Relay becomes easier to justify when released tractor time is more valuable, and harder to justify when ICD-related detour cost rises. This sensitivity check does not constitute a full stochastic model, but it does show that the analytical switching logic is structurally stable and managerially interpretable under plausible parameter perturbations.

3.1.2. Post-Optimal Perturbation and Slack Screen

To respond more directly to timing uncertainty while keeping the paper within its deterministic baseline scope, the certified 10-order schedule was also screened under several simple perturbations. The screen is deliberately conservative: it does not re-optimize the MILP under each scenario and therefore should not be interpreted as a stochastic or robust optimization result. Instead, it asks whether the already-certified schedule contains enough timing margin to absorb small shocks without resequencing.

This screen supports two narrower conclusions. First, the deterministic schedule is not being presented as disruption-proof; several port-facing tasks have limited cutoff margin. Second, the observed idle time in the workload chart should not be read as a calibrated universal slack threshold. It is an instance-specific feature of the certified deterministic schedule, and a general resilience threshold would require explicit stochastic or robust modeling. Table 5 summarizes the corresponding perturbation scenarios and feasibility screen.

Table 5. Post-optimal perturbation and slack screen for the certified 10-order schedule.

Scenario	Perturbation	Cost or Feasibility Screen	Result for the Certified 10-Order Schedule	Interpretation
Baseline	Original loading, travel, and cutoff values	Net cost = CNY 18,733.56; tractors = 5; relay orders = 6; cutoff violations = 0	Feasible certified schedule	Baseline deterministic evidence
Loading +20%	Increase factory pickup duration proxy by 20%	Compare extra pickup duration with pickup-window slack	1 of 10 pickup tasks has insufficient local pickup-window slack	Some resequencing may be needed under loading shocks
Travel +10%	Increase distance-based travel-cost proxy by 10%	Transport-cost proxy rises from CNY 8201.48 to CNY 9021.63	6 of 10 port gate-in tasks have no more than 0.5 h cutoff margin	Travel inflation mainly affects tightly timed port-facing tasks
Cutoff tightened by 1 h	Reduce all gate-in latest times by 1 h in a fixed-schedule screen	Count gate-in completions that would exceed tightened cutoff	6 of 10 gate-in tasks become at risk	The certified schedule is feasible but not universally robust to cutoff compression
Combined stress	Loading +20%, travel +10%, cutoff −1 h	Qualitative fixed-schedule screen	Feasibility cannot be guaranteed without re-optimization	Rolling-horizon or robust extensions are needed for operational deployment

3.2. Cost Composition of the Optimized Scheduling Plan

For consistency across the detailed-instance analysis, Figures 3, 5 and 6 all refer to the same benchmark case with 10 orders, 36 executable tasks, and 5 tractors. Figure 3

reports the cost structure of this solution in gross positive-cost terms. To avoid mixing units, the revised figure reports percentage shares on the slices, while the corresponding absolute CNY amounts are provided in the legend and discussed in the text. The positive cost components sum to CNY 19,914.92, comprising opportunity cost of CNY 9913.44 (49.8%), transport cost of CNY 8201.48 (41.2%), and operating cost of CNY 1800.00 (9.0%). In addition, the optimized schedule realizes a near-pickup coordination credit of CNY 1181.36, which reduces the net total fulfillment cost to CNY 18,733.56. Port storage and violation costs are both zero, confirming that the schedule satisfies all port deadlines without early-arrival penalties.




Figure 3. Cost composition of the MILP-optimal solution in the 10-order instance. Slice labels indicate percentage shares; absolute component values are reported in the legend and text.

This composition is highly informative from a managerial perspective. The dominant role of opportunity cost indicates that the main source of profitability leakage is not the physical act of moving containers itself, but the time during which tractors remain underutilized, tied to loading delays, waiting windows, or poorly synchronized execution. In other words, the economic problem is not only one of route shortening, but one of time monetization. The relatively small proportion of operating cost further suggests that ICD-related handling actions do not materially undermine the value of platform coordination. As long as task pooling meaningfully releases tractor time, the additional operational actions required for relay execution can be economically justified.

This result supports the central argument of the paper: profitability improvement in containerized cross-border fulfillment depends heavily on the platform's ability to reduce non-productive tractor time, not merely on reducing line-haul distance.

Benchmark Comparison

To quantify the economic value of platform-enabled orchestration, Table 6 compares the certified-optimal platform solution with the exact one-truck-per-order benchmark generated

by the restricted MILP. The 10-order case is emphasized because it is the largest instance solved to certified global optimality in both the proposed and benchmark formulations.

Table 6. Performance comparison: platform orchestration vs. one-truck-per-order benchmark (10-order instance).

Metric	Benchmark (One-Truck-per-Order)	Proposed (Platform Orchestration)	Improvement
Gross positive cost before coordination credit (CNY)	27,540	19,915	−27.7%
Net total fulfillment cost after coordination credit (CNY)	27,540	18,734	−32.0%
Net cost per order (CNY/order)	2754	1873	−32.0%
Tractors deployed in the tested solution	10	5	−50.0%
Orders per tractor (OVR)	1.00	2.00	+100.0%
Capacity compression rate (CCR)	0.0%	50.0%	+50.0 pp
Total mileage (km)	1872	2343	+25.2%
Empty-mileage ratio (nempty)	43.0%	43.4%	+0.4 pp
Opportunity-cost share of net total cost	72.0%	52.9%	−19.1 pp
Port cutoff violations	0	0	—

Note: pp denotes percentage points. The higher total mileage under the proposed scheme reflects increased inter-task repositioning that enables cross-order task chaining; as discussed in Section 3.5, this additional mileage is economically absorbed into denser execution networks. The empty-mileage ratio remains close across the two schemes because the benchmark also incurs return-to-depot legs; the key difference lies in the cost-per-order and fleet-size dimensions rather than empty-mileage proportion alone.

Because the proposed schedule increases total mileage while reducing cost and fleet size, a distance-based environmental proxy is also reported to avoid implying that the cost-minimizing schedule is automatically emission-minimizing. Let e_{km} denote a generic emissions factor per vehicle-kilometer and q_{km} denote a generic energy-use factor per vehicle-kilometer. Under identical vehicle and fuel assumptions, the total emissions proxy is proportional to total mileage. In the 10-order comparison, the benchmark records 1872 km, whereas the proposed schedule records 2343 km. The corresponding proxy is therefore $1872 e_{km}$ for the benchmark and $2343 e_{km}$ for the proposed schedule, implying a 25.2% higher distance-based emissions proxy for the cost-minimizing platform schedule. The same proportional relationship applies to the energy-use proxy, $1872 q_{km}$ versus $2343 q_{km}$.

This proxy does not represent a calibrated carbon inventory, because vehicle type, load factor, fuel technology, congestion, and driving-cycle effects are not modeled. It nevertheless makes the economic–environmental trade-off explicit: the platform schedule reduces net fulfillment cost and fleet deployment, but it does so partly by using additional repositioning mileage. A full green-logistics extension should therefore treat emissions or energy use as an explicit objective or constraint rather than assuming that economic and environmental efficiency always move together.

Notes: Benchmark values are exact outputs of the restricted MILP and are rounded to whole CNY and kilometers for readability. The gross positive cost reports the direct cost components before the near-pickup coordination credit; the net total fulfillment cost applies the same accounting convention to both schemes. Since the one-truck-per-order benchmark has no near-pickup coordination credit, its gross and net values are the same. For the 5-order certified-optimal case, the proposed solution achieves a total cost of CNY 7899 with three tractors and OVR = 1.67 (CCR = 40%), versus an exact benchmark cost of CNY 11,048 with five tractors (OVR = 1.00), implying a cost reduction of approximately 28.5%.

3.3. Computational Scaling of the Optimization Model

The second result concerns computational scalability under the tested fleet settings. Figure 4 reports the runtime response of the exact MILP for four benchmark configurations: 5 orders/3 tractors, 10 orders/5 tractors, 15 orders/7 tractors, and 20 orders/10 tractors. Across these configurations, the task volume increases from 18 to 74 tasks. The 5-order and 10-order instances are solved to certified global optimality in 117.56 s and 228.82 s, respectively. For the 15-order and 20-order cases, the solver exhausts the respective run budgets (600 s and 900 s) without reaching a certified optimum. The experiment log records total wall-clock elapsed times of 1170.75 s and 2794.38 s, reflecting full model construction and solution time, but no solver outcome from these larger runs is used in the manuscript as quantitative evidence of platform performance. All later interpretation for the 15-order and 20-order settings is therefore limited to pre-specified stress-test fleet configurations rather than solver-verified optimization results.




Figure 4. MILP runtime growth with increasing problem size.

The pattern in Figure 4 indicates a marked loss of tractability as the problem scale expands. This result should be interpreted as evidence of the growing combinatorial burden associated with task assignment, temporal synchronization, and relay sequencing, rather than as a limitation of any single benchmark instance. From a managerial and methodological perspective, the implication is clear: exact optimization remains valuable as a benchmark and policy-evaluation device, but routine platform operation at larger scales will likely require decomposition strategies, heuristics, or rolling-horizon control to maintain computational responsiveness.

3.4. Vehicle-Level Task Sequencing Under Platform Orchestration

Figure 5 provides a Gantt-chart representation of the optimized tractor schedule for the same 10-order benchmark instance analyzed in Figures 3, 6 and 7. The figure shows how five tractors jointly execute 36 tasks over the planning horizon through a combination

of empty delivery, loaded pickup, ICD relay pickup, and port gate-in actions. Rather than following one-truck-one-order continuity, the tractors perform cross-order task chaining, which is the core operational expression of platform-enabled orchestration.

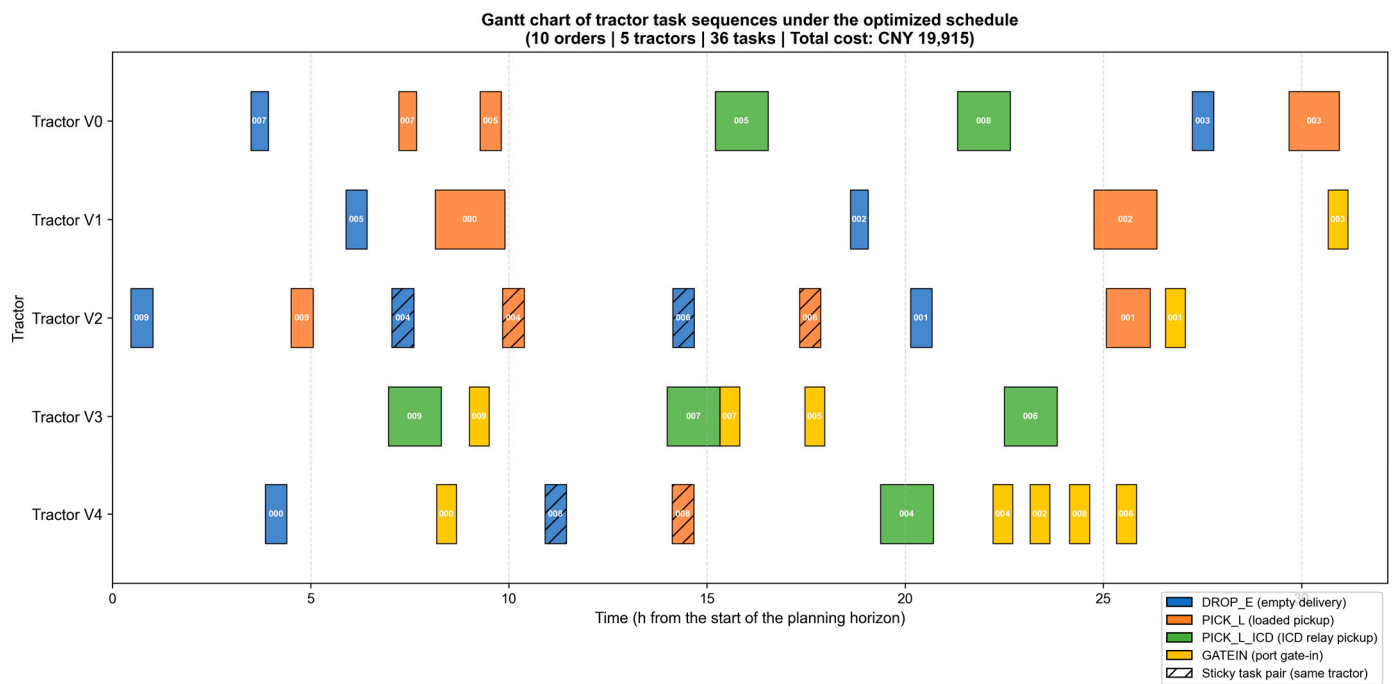


Figure 5. Gantt chart of tractor task sequences for the same 10-order benchmark instance under the optimized platform schedule.




Figure 6. Mileage composition of the optimized scheduling plan.




Figure 7. Fleet compression, orders-per-tractor ratio, and average daily mileage under the tested scaling configurations. Note: The 5-order and 10-order values are certified-optimal solver outcomes. The 15-order and 20-order values are stress-test fleet targets rather than solver-verified outcomes.

Several features of the schedule are noteworthy. First, tractors execute heterogeneous task portfolios rather than fixed repetitive routes, indicating that the platform reallocates vehicles according to task readiness and network timing rather than order ownership. Second, some tasks are scheduled as continuous paired operations on the same tractor, while others are separated through relay or ICD buffering. This mixed pattern suggests that the model does not force a single operational mode; instead, it selects continuity or handoff according to economic and temporal conditions. Third, the planning horizon shows that tasks are spread across the day in a way that preserves feasible sequencing while maintaining a relatively dense workload on each tractor.

From a platform perspective, the Gantt chart makes visible the transformation from isolated trip execution to network-level orchestration. The economic value of the ICD does not lie only in storage, but in enabling this kind of cross-order chaining. Without such a coordination node, much of the observed sequencing flexibility would collapse into waiting or deadheading.

3.5. Mileage Structure and Detour Absorption

Figure 6 decomposes total mileage in the 10-order optimized instance into three components: inter-task repositioning mileage, within-task execution mileage, and return-to-base mileage. The total mileage is 2343.3 km, of which 1326.0 km (56.6%) corresponds to within-task execution, 617.3 km (26.3%) corresponds to inter-task repositioning, and 400.0 km (17.1%) corresponds to final return-to-base movement.

Two insights emerge from this structure. First, productive mileage remains the dominant component, which means that the platform successfully converts most tractor movement into economically meaningful execution rather than pure repositioning. Second, however, empty or semi-empty repositioning remains substantial. This indicates that

profitability improvement cannot be interpreted as the complete elimination of detours; instead, it comes from the platform's ability to absorb detours into a denser execution network. A repositioning leg is economically less damaging when it connects to the next revenue-generating task quickly and when the same tractor can continue serving multiple orders within one planning horizon.

Accordingly, the relevant question is not whether the ICD introduces additional mileage in an absolute sense, but whether the additional mileage enables enough task chaining to improve economic return overall. The present result suggests that the answer is yes: even though repositioning remains visible, it does not dominate the mileage structure, and the pooled schedule still yields strong system-level productivity gains.

3.6. Capacity Compression and Tractor Productivity Under Increasing Order Volume

One of the most significant results of the study is reported in Figure 7, which summarizes the tested scaling configurations through the orders-per-tractor ratio (OVR), the implied capacity compression rate, and the average daily mileage per tractor. Across the four configurations, the fleet setting changes from 3 tractors for 5 orders to 10 tractors for 20 orders. For the two certified-optimal instances, OVR rises from 1.67 (5 orders, 3 tractors) to 2.00 (10 orders, 5 tractors), reflecting capacity compression of 40.0% and 50.0%, respectively. For the 15-order and 20-order stress-test configurations, the pre-specified fleet settings imply OVR values of 2.14 and 2.00, with associated compression rates of 53.3% and 50.0%; however, these larger-case values are scenario targets rather than solver-verified outcomes.

Accordingly, the evidence should be interpreted in two layers. For the two certified cases, the monotonic increase in OVR confirms that platform-level task recombination allows each tractor to serve more orders without a proportional fleet increase. For the 15-order and 20-order settings, the reported values should be read only as stress-test indications that similar compression logic may remain feasible under denser task pools; they do not prove larger-scale optimal performance. The non-monotonic pattern across all four configurations therefore reflects the distinction between certified optimization results and scenario-based fleet targets, rather than a verified decline in returns.

The mileage indicator in Figure 7 should likewise be read as a complementary productivity signal rather than as proof of indefinite mileage expansion. In managerial terms, the platform gains first from better utilization and then from better composition of the available task pool. This interpretation is also consistent with Figure 4: as scale grows, both economic coordination opportunities and optimization difficulty rise simultaneously.

3.7. Workload Balance and Schedule-Preserving Slack

The final result concerns vehicle-level workload distribution. Figure 8 compares the number of tasks assigned to each tractor in the 10-order optimized instance with the corresponding idle time. The number of tasks per tractor ranges from 5 to 9, while idle time ranges from approximately 14.9 h to 18.6 h. The implied utilization rates shown in the figure range from roughly 40.4% to 52.2%.

At first glance, these utilization levels may appear modest relative to highly compressed line-haul operations. In the present deterministic instance, however, the remaining idle time should be interpreted more narrowly as schedule-preserving slack rather than as proof of a universal optimal slack level. The platform is operating under multiple timing constraints, including loading completion, port deadlines, relay feasibility, and task continuity. Under these conditions, forcing every tractor to operate at near-saturation could require excessive repositioning or introduce schedule fragility. Instead, the optimized solution distributes tasks in a relatively balanced manner and preserves non-negligible idle time across all tractors.

This pattern reveals an important managerial trade-off, but it should not be read as identifying a quantitative slack threshold. Profitability improvement does not require maximizing the utilization of every individual tractor at every moment. Rather, in this certified 10-order schedule, the platform gains value by maintaining enough task density to avoid severe underuse while retaining enough slack to preserve feasible chaining. The perturbation screen in Section 3.1.2 further shows why this interpretation must remain conditional: several port-facing tasks are timing-sensitive even though the schedule contains substantial aggregate idle time.




Figure 8. Vehicle workload, utilization, and idle time under the optimized schedule.

3.8. Summary of Results

Taken together, the certified results and the comparative-static boundary checks indicate that ICD-enabled platform orchestration improves economic performance through four related mechanisms. First, it creates a transparent boundary for choosing between relay and direct execution. Second, it reduces profit leakage associated with tractor waiting and underutilization, as reflected in the central role of opportunity cost. Third, it reorganizes task sequences in a way that absorbs detours into denser execution chains rather than treating each order as a standalone trip. Fourth, the certified cases show substantial capacity compression, while the 15-order and 20-order stress-test settings are used only to indicate tractability pressure and possible fleet targets rather than solver-verified large-scale performance.

These findings establish that the value of platform-enabled freight orchestration lies not only in route optimization, but in the platform's ability to transform time, space, and task fragmentation into economically coordinated execution.

4. Discussion

The results indicate that the value of platform-enabled freight orchestration cannot be reduced to conventional route minimization. In the present setting, economic improvement emerged from a broader coordination logic in which the platform reorganized task

continuity, redistributed waiting time, pooled geographically separable jobs, and used the ICD as a digital–physical node for controlled task handoff. From this perspective, the core managerial question is not simply how to move containers at the lowest direct transport cost, but how to allocate tractor time and network slack so that a given order volume can be fulfilled with higher economic return.

4.1. The ICD as a Digital–Physical Coordination Node

A central implication of this study is that the ICD should be interpreted as a coordination mechanism rather than merely an intermediate storage facility. In traditional operational logic, an inland depot is often treated as a passive location where containers wait until the next movement stage becomes available. The present results suggest a different interpretation. When embedded in a platform-controlled scheduling system, the ICD functions as a digital–physical coordination node that allows the platform to separate task timing from tractor continuity.

This distinction is important because the economic bottleneck in cross-border containerized fulfillment is often not the physical impossibility of moving a container, but the temporal mismatch between where tractors are available and when specific tasks become executable. Factory loading may finish late, port cutoff times may be rigid, and empty-container reuse opportunities may emerge asynchronously across locations. The ICD absorbs these mismatches by enabling a platform to release one tractor before the next execution stage is ready. In this sense, the ICD does not merely hold containers; it holds temporal pressure on behalf of the platform.

The results reinforce this interpretation in two ways. First, the relay-versus-direct boundary in Figure 2 shows that relay execution becomes economically attractive precisely when tractor waiting time is sufficiently costly relative to ICD detour distance. Second, the task-sequencing pattern in Figure 5 demonstrates that once a relay node is available, tractors can be reconnected across otherwise unrelated orders. This shifts the operational unit of analysis from the individual shipment to the task network. The managerial implication is that the strategic value of an ICD depends less on its storage role per se than on whether the platform can exploit it to recompose execution flow.

4.2. Relay Execution Should Be Selective, Not Universal

The findings also suggest that ICD-based relay execution should be treated as a conditional platform option, not as a universally superior operating mode. Figure 2 makes this clear. Relay execution is economically favored only when the expected loading time is long enough, relative to the distance between the ICD and the factory, to justify decoupling the tractor from the order. When the ICD is close to the factory, the threshold loading time is relatively low; when the ICD is farther away, the threshold rises materially.

This has two managerial consequences. First, platform dispatch rules should be state-dependent rather than rule-of-thumb based. It is not sufficient to assume that all export jobs should be handled through relay because an ICD exists. Nor is it efficient to insist on direct continuity merely because a one-truck-one-order rule is easier to manage operationally. Instead, the platform should use tractable decision boundaries to decide when relay is worth the additional repositioning. Second, the role of algorithmic orchestration becomes more valuable when loading-time uncertainty or delay heterogeneity is high. In such cases, the platform benefits from being able to redeploy tractors quickly rather than tying them to factory-side dwell time.

More broadly, this result refines the meaning of “flexibility” in digital freight platforms. Flexibility does not imply unlimited mode switching or complete operational fluidity. Rather, it means that the platform can choose among execution structures according to economically

interpretable conditions. In the present context, the relay option creates value not because it is always cheaper, but because it expands the decision set available to the platform.

4.3. Profitability Improvement Is Driven by Time Reallocation More than by Handling Cost Reduction

The cost composition in Figure 3 offers a particularly important insight for platform operations. Opportunity cost is the largest single cost component in the optimized solution, exceeding both direct transport cost and operating cost. This indicates that the principal source of economic inefficiency in the studied fulfillment network is not excessive handling expenditure or even excessive line-haul mileage alone, but the underuse of tractor time.

This is a crucial finding because many freight optimization approaches implicitly assume that profitability improvement should be pursued primarily through shorter routes or fewer handling actions. The present results suggest a different interpretation. In platform-mediated container fulfillment, profit leakage is often generated by the fact that tractors remain tied to processes that do not create immediate value, such as factory loading delays, coordination gaps between successive jobs, or temporally mismatched pickup windows. Accordingly, the value of orchestration lies in time reallocation: a tractor released from one incomplete process can be inserted into another executable task without waiting for the original order to progress.

This also explains why the operating-cost share remains comparatively small. ICD-enabled relay execution may introduce additional handling actions, but these actions are economically acceptable if they unlock enough tractor time to reduce opportunity loss elsewhere in the network. In other words, the platform should not evaluate relay execution solely by counting handling steps. It should evaluate it by comparing the cost of one more handoff with the value of one more reusable tractor hour.

For logistics managers, this implies that profitability-oriented scheduling should monitor waiting time and task readiness with the same intensity traditionally devoted to route distance. For platform operators, it suggests that the key data assets are not only map-based routes, but also reliable forecasts of loading completion, gate-in timing, and task readiness states.

4.4. Task Pooling Compresses Fleet Requirements by Changing the Unit of Coordination

The most visible economic effect of the proposed orchestration logic is the fleet-compression pattern reported in Figure 7. In the certified solutions, the orders-per-tractor ratio rises from 1.67 to 2.00; the larger stress-test settings extend the displayed range to 2.14 in the 15-order case before easing to 2.00 at 20 orders, but those latter values are not solver-verified. Accordingly, the pattern should not be interpreted simply as a mechanical consequence of higher volume. Rather, it suggests how the platform changes execution economics once tractors are assigned at the task level rather than at the order level.

Under isolated order execution, the platform effectively assigns capacity at the order level. A tractor remains associated with one order for long stretches, even if some parts of that time are economically unproductive. Under the proposed task-pooling logic, the platform instead coordinates at the task level. This enlarges the combinatorial set of feasible continuations after each completed action. Once the platform is allowed to connect tractors to whichever task is economically most attractive at the next decision point, the certified cases already show clear scale benefits relative to one-truck-one-order continuity, while the larger stress-test settings indicate that similar coordination logic may remain structurally plausible under denser pools, although certified optimality at those scales was not achieved within the run budget.

The practical consequence is significant. A digital freight platform can increase throughput without increasing fleet size proportionally, provided that enough compatible

tasks exist within the same scheduling horizon. This is especially relevant for cross-border e-commerce, where order flows are volatile and peak periods can create pressure to expand capacity quickly. The current findings suggest that part of this pressure can be absorbed not by acquiring more tractors, but by restructuring execution logic. Capacity compression, therefore, is not simply a function of volume density; it is a function of whether the platform can translate volume density into cross-order task compatibility.

From a business-model perspective, this result also matters because it links platform value creation to asset-light expansion. The platform does not need to own proportionally more tractors to grow fulfillment volume. Instead, it needs the coordination ability to extract more economically useful output from the existing fleet. This is an important bridge between freight optimization and platform economics.

4.5. Additional Mileage Is Not Necessarily Wasteful if It Is Absorbed into a Denser Execution Network

The mileage analysis in Figure 6 adds nuance to the interpretation of efficiency. The optimized plan does not eliminate repositioning mileage. Inter-task repositioning remains a significant component of total distance. At first glance, this might seem to weaken the claim that platform orchestration improves profitability. However, the relevant question is not whether some extra repositioning exists, but whether the additional mileage is economically absorbed by denser task chaining and better time allocation.

In the present setting, more than half of total mileage remains within-task execution mileage, which indicates that the system preserves a strong productive movement core. The remainder reflects the cost of maintaining connectivity between tasks. Yet this cost becomes acceptable when repositioning is immediately followed by another revenue-generating or deadline-critical task. In that case, repositioning is not merely deadhead movement; it is an enabling movement that connects otherwise incompatible tasks.

This insight is especially important for evaluating ICD-enabled relay strategies. A narrow route-minimization perspective may conclude that any detour to an ICD is undesirable. The present results show that this conclusion can be misleading. An ICD detour may increase travel in an isolated sense, but it may still improve economic return if it allows the platform to redeploy a tractor sooner, avoid excessive waiting, or create a higher-density downstream task chain. In such cases, the platform effectively trades a small amount of additional distance for a much larger gain in time productivity and fleet compression.

The certified cases support this interpretation directly. Average mileage per tractor remains economically meaningful, but the benefit of orchestration is not simply that each tractor drives more kilometers. The larger 15-order and 20-order values in Figure 7 are descriptive values from pre-specified stress-test configurations and are not interpreted as optimized or certified performance outcomes. They are therefore used only to illustrate capacity-planning scenarios under denser task pools. The platform improves profitability not by making trucks drive indefinitely more, but by making daily mileage more economically productive through better task composition.

4.6. Schedule-Preserving Slack Can Be Economically Rational

Figure 8 reveals another important managerial pattern: tractor utilization in the optimized schedules is moderate rather than extreme, with visible idle time remaining across all vehicles. In conventional fleet-management thinking, such idle time may be interpreted as inefficiency. The present results suggest a more careful interpretation. In an environment governed by loading delays, gate-in deadlines, precedence constraints, and relay feasibility, some idle time can preserve schedule feasibility rather than represent operational waste.

This point matters because the platform's goal is not to maximize the utilization of each individual tractor in isolation. Its goal is to minimize total system cost while preserving fulfillment feasibility. If the model were to push every tractor toward near-saturation, it would likely have to accept tighter task chaining, more brittle route continuity, and greater vulnerability to even small timing deviations. Under such conditions, a schedule might appear efficient *ex ante* but become fragile in practice. By contrast, a schedule with some remaining slack can maintain feasible handoff structure and reduce the need for emergency repositioning.

The current findings therefore imply that profitability improvement in digitally coordinated freight systems is consistent with a balanced workload profile, not necessarily with maximal instantaneous utilization. This insight is important for both theory and practice. Theoretically, it suggests that utilization should be interpreted jointly with timing feasibility and synchronization. Practically, it suggests that platform performance dashboards should avoid treating idle time as uniformly negative. Some idle time may be economically justified if it supports lower overall opportunity cost, better task timing, and more reliable port compliance.

That said, the present model operates under deterministic perfect-information conditions. The perturbation screen in Section 3.1.2 shows that the certified schedule remains sensitive to loading, travel, and cutoff shocks in specific locations. Therefore, the current result should not be interpreted as a universal target utilization rate or an optimal slack threshold. Rather, it should be understood as an instance-specific observation that economically efficient orchestration may preserve operational breathing room while still requiring re-optimization or robust control under stronger uncertainty.

4.7. Implications for Systems Research and Digital Commerce

Beyond the freight-planning context, the study has broader implications for systems research and digital commerce. The present findings position the fulfillment network as a socio-technical transportation system whose performance depends on coordinated interactions among physical nodes, mobile tractor resources, digital rules, and time-critical constraints [26,27]. From a systems perspective, this study shows that value emerges not from optimizing individual trips in isolation, but from managing interdependence across tasks, vehicles, and nodes within a defined system boundary. The ICD functions as a boundary component within this system: it buffers temporal mismatches between the physical and digital layers and enables the platform to maintain coherent execution flow despite non-synchronous order conditions.

This study contributes to that perspective in two ways. First, it shows that the platform's economic role extends beyond bilateral matching between demand and supply. In the freight setting, the platform is an orchestrator of sequential interdependence: it determines how one task creates or constrains the next. Second, it demonstrates that digital value creation can occur through the management of time and compatibility in physical networks. The platform creates economic value not only by reducing search cost, but also by increasing the number of economically feasible task combinations that can be executed with a limited asset base. In other words, the technological contribution of the platform lies not only in routing computation, but in the integration of visibility, synchronization, and rule-based control over a temporally coupled work system.

In this sense, the paper extends digital-commerce thinking from marketplace design toward operational execution architecture. The relevant performance indicators are no longer only conversion, retention, or transaction volume, but also orders handled per tractor, capacity compression, opportunity-cost reduction, and fulfillment feasibility under temporal constraints. This is particularly relevant in cross-border settings, where the

physical execution layer is often the hidden determinant of whether digital demand can be profitably fulfilled.

4.8. Managerial Implications

Several practical implications follow from the results.

First, platform operators should implement rule-based relay screening rather than using fixed dispatch habits. The relay-versus-direct boundary provides a basis for deciding when ICD handoff is likely to be profitable and when it is not.

Second, logistics managers should evaluate fleet performance using task-density and productivity indicators, such as orders per tractor and capacity compression, rather than focusing only on trip-level route efficiency. A platform can create value by improving how much economically useful work is extracted from each tractor.

Third, dispatch systems should explicitly track task readiness information, especially factory loading completion and port timing windows. Because opportunity cost is central to overall performance, poor visibility into readiness states can quickly undermine profitability.

Fourth, operators should interpret the ICD as a coordination asset rather than a storage asset alone. Its main value lies in enabling task handoff, deferred dispatch, and local compatibility exploitation, not simply in holding containers.

Finally, managerial performance targets should balance utilization with schedule feasibility. Pursuing extremely high tractor utilization may appear attractive in static reporting, but it can reduce the platform's ability to maintain robust sequencing and absorb operational variation.

4.9. Boundary Conditions of Interpretation

The discussion should be interpreted within the scope of the present study. The results are generated under a deterministic perfect-information baseline and do not yet include stochastic travel times, real-time order arrivals, behavioral responses by independent carriers, or endogenous pricing decisions. The perturbation screen in Section 3.1.2 is a fixed-schedule diagnostic rather than a stochastic or robust optimization model. The model also assumes centralized platform authority over dispatching and sequencing, so the conclusions are strongest for internally coordinated or tightly governed platform settings rather than for loosely coupled subcontracting networks. In addition, the main certified benchmark remains export-side; mixed import/export reuse patterns are tested only through small activation examples in Appendix A rather than used as broad performance evidence. Accordingly, the current findings should be read as a structured demonstration of the economic logic of platform orchestration rather than as a universal prescription for all freight contexts.

Even with these limitations, the main message remains robust: the profitability of cross-border e-commerce fulfillment can be materially improved when a digital freight platform reorganizes inland container movements as a task network rather than as isolated trips, and when the ICD is used as a controlled node for spatiotemporal decoupling rather than as passive infrastructure.

5. Conclusions

This study examined how a digital freight platform can improve the profitability of cross-border e-commerce fulfillment by orchestrating inland container truck movements through an ICD-based spatiotemporal decoupling structure. Rather than treating the ICD as a passive storage facility, the paper conceptualized it as a digital–physical coordination node through which the platform can pool tasks, release tractors from non-productive waiting, and recombine execution stages across orders. On this basis, the study developed

a drop–buffer–pick task framework and a profitability-oriented MILP model that jointly captured direct execution, relay execution, ICD buffering, local empty-container reuse, and tractor sequencing under loading-time, gate-in, inventory, and working-time constraints.

The results show that the economic value of platform-enabled orchestration lies in more than route shortening. First, the relay-versus-direct execution boundary demonstrates that ICD-based handoff becomes economically attractive when the value of released tractor time exceeds the additional mileage associated with the ICD detour. The comparative-static sensitivity analysis confirms that this switching logic moves transparently with time-value and distance-cost parameters rather than depending on a single calibration. Second, the cost analysis shows that opportunity cost is the largest single cost component in the optimized fulfillment plan, which indicates that profitability improvement depends fundamentally on the platform's ability to reduce tractor waiting and underutilization. Third, the mileage and scheduling results reveal that the platform does not eliminate all repositioning movement, but it absorbs otherwise unproductive detours into denser task chains that generate higher system-level output. Fourth, the certified 5-order and 10-order cases show substantial capacity-compression effects, whereas the 15-order and 20-order settings are used only to expose the tractability limits of the exact MILP rather than to claim verified large-scale superiority. Taken together, these findings indicate that profitability improvement is driven by time reallocation, task recombination, and asset-productivity gains rather than by transport-cost reduction alone.

From a systems perspective, the study contributes a socio-technical coordination framework for analyzing platform-enabled freight execution as an interdependent transportation system rather than a set of isolated trips. It shows that system-level performance depends on managing interdependence among tasks, vehicles, and nodes within a defined system boundary, and that the ICD functions as a critical boundary component that buffers temporal mismatches and enables the platform to maintain coherent execution flow across the fulfillment system. At the application level, the framework also extends digital-commerce research by showing how upstream fulfillment performance depends on platform-mediated orchestration of sequential physical tasks, demonstrating that algorithmic task recombination is itself a source of platform value beyond market intermediation. The paper further contributes to the logistics and operations literature by reframing the ICD as a profit-enabling coordination asset and by providing interpretable decision boundaries for relay execution, buffering, and task pooling.

From a managerial perspective, the findings suggest that platform operators should adopt state-dependent dispatch rules rather than relying on fixed operating habits. ICD relay should be used selectively under conditions in which factory-side loading delay or timing mismatch makes tractor release economically valuable. Fleet performance should be evaluated not only through trip-level distance efficiency but also through orders-per-tractor ratios, capacity compression, opportunity-cost exposure, emissions/energy-use proxies, and slack-aware workload balance. More broadly, the results indicate that digital freight platforms can improve economic return without proportionally expanding fleet size, provided that enough task compatibility exists within the planning horizon and that the platform has sufficient visibility into task readiness and vehicle availability.

At the same time, the study has several limitations. The current model is deterministic and assumes perfect information regarding task readiness, travel time, and port timing. It does not incorporate stochastic loading uncertainty, real-time order arrivals, endogenous carrier responses, pricing decisions, or contractual interactions among independent fleet participants. The perturbation screen in Section 3.1.2 is only a fixed-schedule diagnostic and should not be read as a full robust-optimization result. In addition, the profitability interpretation is based on a fixed-order setting in which all orders must be served; revenue heterogeneity, selective order acceptance, dynamic pricing, and outsourcing decisions are not yet modeled explicitly. The current evaluation is also cost-first. Although Section 3.2

reports a distance-based emissions and energy-use proxy, the model does not include calibrated fuel consumption, vehicle-specific emissions factors, load-dependent energy use, or carbon-cost constraints; economic and environmental performance are therefore not jointly optimized in the present baseline. The present formulation also assumes centralized dispatch control, so it does not represent how benefits would be divided across external contractors or collaborative carrier networks. Furthermore, the main computational experiments are conducted on export-side instances; while Appendix A provides small MILP activation tests for street-turn reuse and triangular backhaul, these operational modes are not validated through the main certified benchmark comparison. These simplifications were intentional in order to isolate the core economic logic of ICD-enabled orchestration, but they limit the generalizability of the results to more dynamic platform environments.

Future research can extend the present framework in several directions. A first avenue is to integrate real-time and stochastic decision-making, including uncertain loading completion, rolling task arrivals, and dynamic port conditions. A second avenue is to incorporate explicit revenue-side mechanisms such as order selection, differential service pricing, outsourcing, and penalty-sensitive service-level management, which would allow a direct revenue-maximization or profit-maximization formulation. A third avenue is to examine behavioral and governance issues in platform-based freight orchestration, including incentive alignment among shippers, carriers, and platform operators. A fourth avenue is to extend the present cost-first baseline toward multi-objective formulations that evaluate emissions, energy use, and economic performance jointly. Finally, empirical validation using platform transaction data and field deployment cases would strengthen the external validity of the proposed coordination logic.

In conclusion, this paper argues that the profitability of cross-border e-commerce logistics can be materially improved when a digital freight platform reorganizes inland container execution as a pooled task system and uses the ICD as a controlled relay node for spatiotemporal decoupling. The sustainability implication is equally important: efficiency-oriented orchestration should be evaluated together with mileage-related emissions and energy-use indicators, because lower operating cost and smaller fleet deployment may still involve additional repositioning distance. In digitally coordinated freight networks, economic improvement therefore depends not only on moving containers more cheaply, but also on how effectively the platform can reorganize time, task continuity, asset availability, and environmental exposure across the fulfillment system.

Author Contributions: Conceptualization, S.F. (Shucheng Fan); methodology, S.F. (Shucheng Fan); software, S.F. (Shucheng Fan); validation, S.F. (Shaochuan Fu); formal analysis, S.F. (Shucheng Fan); data curation, S.F. (Shucheng Fan); Writing—original draft preparation, S.F. (Shucheng Fan); Writing—review and editing, S.F. (Shucheng Fan) and S.F. (Shaochuan Fu); supervision, S.F. (Shaochuan Fu). All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data and computational artifacts supporting the reported results are available from the corresponding author upon reasonable request.

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

Appendix A. Small MILP Activation Tests for Extended Task Patterns

This appendix reports two small MILP activation tests for task patterns that are structurally supported by the unified framework but not included in the main export-side certified benchmark comparison. Table A1 summarizes the first test, and Table A2 summarizes the second test. The tests are intentionally minimal and should not be interpreted as broad empirical validation. Their purpose is to demonstrate that the same binary se-

lection logic can activate mixed import/export street-turn reuse and triangular backhaul transitions under the same calibrated cost structure.

The first test contains one import order and one export order. After the import container is unloaded, the empty unit can either return to the ICD and then be dispatched to the export factory, or it can be reused locally through a street-turn handoff. The reduced binary MILP selects exactly one of these two patterns and minimizes the corresponding transition cost:

$$y^{STD}, y^{ST} \in \{0, 1\}, y^{STD} + y^{ST} = 1,$$

$$\min C^{STD} y^{STD} + C^{ST} y^{ST}.$$

Table A1. Small MILP activation test for the local street-turn pattern.

Test Item	Standard Return-to-ICD Pattern	Local Street-Turn Pattern
Solver/status	PuLP/CBC/Optimal	PuLP/CBC/Optimal
Transition cost	CNY 410.0	CNY 217.5
Selected pattern	–	Street-turn
Cost advantage	–	CNY 192.5

The selected task sequence is PICK_I → DROP_I → TRANSFER_E_ST → PICK_L → GATEIN.

The street-turn gain is computed from the same distance and cost calibration used in Section 2.3.3: $d_{im_to_b} = 30$ km, $d_{b_to_ex} = 30$ km, $d_{im_to_ex} = 25$ km, $c_km = 3.5$ CNY/km, and the net operating-cost advantage of local reuse is 70 CNY. The resulting gain is $3.5 \times [(30 + 30) - 25] + 70 = 192.5$ CNY, so the MILP selects the street-turn pattern.

The second test checks a triangular-backhaul transition at the port. A tractor that has completed a port gate-in task may either return through the ICD before a later port-side empty pickup, or it may directly connect the gate-in completion to a port empty-pickup task. The reduced binary MILP again selects exactly one transition pattern:

$$y^{STD}, y^{TRI} \in \{0, 1\}, y^{STD} + y^{TRI} = 1,$$

$$\min C^{STD} y^{STD} + C^{TRI} y^{TRI}.$$

Table A2. Small MILP activation test for the triangular-backhaul transition.

Test Item	Standard Port-Return Transition	Triangular Backhaul Transition
Solver/status	PuLP/CBC/Optimal	PuLP/CBC/Optimal
Transition cost	CNY 560.0	CNY 0.0
Selected pattern	–	Triangular backhaul
Cost advantage	–	CNY 560.0

The selected task sequence is GATEIN → PICK_E_PORT → DROP_E.

This second activation test uses the port-ICD distance of 80 km and $c_km = 3.5$ CNY/km. The standard transition requires a port-to-ICD movement followed by an ICD-to-port repositioning movement, while the triangular transition links the two port-side tasks directly. The reported costs are transition costs only; task-execution costs common to both alternatives are excluded from this activation comparison. The zero value for the triangular transition refers only to the absence of additional repositioning between two consecutive port-side tasks, not to the total cost of the subsequent empty-container movement. The simplified test therefore selects the triangular-backhaul transition.

Appendix B. Calculation Transparency for the Certified 10-Order Benchmark

This appendix summarizes the main arithmetic used in the certified 10-order comparison. Tables A3–A5 provide the corresponding cost accounting, mileage decomposition, and solver-status summary. The purpose is reproducibility and auditability rather than additional performance evidence. All values correspond to the same 10-order instance discussed in Section 3.2 unless otherwise stated.

Table A3. Cost accounting for the proposed 10-order schedule.

Component	Value
Transport cost	CNY 8201.48
Fixed operating cost	CNY 1800.00
Tractor opportunity cost	CNY 9913.44
Port storage cost	CNY 0.00
Violation/risk penalty	CNY 0.00
Gross positive cost	CNY 19,914.92
Near-pickup coordination credit	-CNY 1181.36
Net total fulfillment cost	CNY 18,733.56

Table A4. Mileage decomposition and environmental proxy.

Indicator	Benchmark	Proposed Platform Schedule
Total mileage	1872 km	2343 km
Within-task execution mileage	Not separately reported for benchmark	1326.0 km
Inter-task repositioning mileage	Not separately reported for benchmark	617.3 km
Return-to-base mileage	Not separately reported for benchmark	400.0 km
Distance-based emissions proxy	1872 e_{km}	2343 e_{km}
Distance-based energy-use proxy	1872 q_{km}	2343 q_{km}

The proxy comparison assumes the same vehicle type and a constant per-kilometer emissions or energy-use factor across both schedules. Under that simplifying assumption, the proposed schedule has a 25.2% higher distance-based environmental proxy than the benchmark, despite its lower net fulfillment cost and smaller tractor requirement. This result reinforces that the present model is a cost-first baseline and that future green-logistics extensions should optimize economic and environmental objectives jointly.

Table A5. Solver-status summary used for interpreting scalability.

Orders	Tractor Setting	Task Count	Reported Time	Solver Status	Interpretation
5	3	18	117.56 s	Certified optimal	Used as certified performance evidence
10	5	36	228.82 s	Certified optimal	Used as certified performance evidence
15	7	56	1170.75 s elapsed	Not certified optimal	Used only as tractability-limit evidence
20	10	74	2794.38 s elapsed	Not certified optimal	Used only as tractability-limit evidence

References

1. Liu, A.; Osewe, M.; Shi, Y.; Zhen, X.; Wu, Y. Cross-Border E-Commerce Development and Challenges in China: A Systematic Literature Review. *J. Theor. Appl. Electron. Commer. Res.* **2022**, *17*, 69–88. [\[CrossRef\]](#)
2. Zhu, Q.; Ruan, Y.; Liu, S.; Yang, S.-B.; Wang, L.; Che, J. Cross-Border Electronic Commerce's New Path: From Literature Review to AI Text Generation. *Data Sci. Manag.* **2023**, *6*, 21–33. [\[CrossRef\]](#)
3. Yang, Y.; Chen, N.; Chen, H. The Digital Platform, Enterprise Digital Transformation, and Enterprise Performance of Cross-Border E-Commerce—From the Perspective of Digital Transformation and Data Elements. *J. Theor. Appl. Electron. Commer. Res.* **2023**, *18*, 777–794. [\[CrossRef\]](#)
4. Fei, L.-G.; Liu, X.; Jin, Y.-C.; Su, M. Reconstruction of Logistics Services in Cross-Border E-Commerce and Consumer Continuance Intention on Platforms: The Mediating Role of Digital Logistics Services. *J. Theor. Appl. Electron. Commer. Res.* **2025**, *20*, 251. [\[CrossRef\]](#)
5. Jiang, W.; Lu, H.; Wang, Z.; Jing, Y. How Does Port Logistics Service Innovation Enhance Cross-Border e-Commerce Enterprise Performance An Empirical Study in Ningbo-Zhoushan Port, China. *J. Theor. Appl. Electron. Commer. Res.* **2025**, *20*, 188. [\[CrossRef\]](#)
6. Li, Y.; Li, Z. Research on B2C Cross-Border Electronic Commerce Return Logistics Model Selection Based on Estimated Return Rate. *J. Theor. Appl. Electron. Commer. Res.* **2024**, *19*, 1034–1059. [\[CrossRef\]](#)
7. Li, L.; Li, G. Integrating Logistics Service or Not The Role of Platform Entry Strategy in an Online Marketplace. *Transp. Res. Part E Logist. Transp. Rev.* **2023**, *170*, 102991. [\[CrossRef\]](#)
8. Wang, M.; Zhao, L.; Herty, M. Logistics Service Sharing of Sellers on an e-Commerce Platform Considering the Behaviour of Carbon Emission Pooling. *Electron. Commer. Res. Appl.* **2023**, *59*, 101267. [\[CrossRef\]](#)
9. Qin, X.; Liu, Z.; Tian, L. The Strategic Analysis of Logistics Service Sharing in an E-Commerce Platform. *Omega* **2020**, *92*, 102153. [\[CrossRef\]](#)
10. Li, J.; Shen, W.; Liao, Y.; Cai, G.; Chen, X. The Fulfillment Service in Online Marketplaces. *Eur. J. Oper. Res.* **2024**, *315*, 1139–1152. [\[CrossRef\]](#)
11. Khooban, F.; Mutlu, N.R.; de Kok, A.G. Logistics Service Sharing in Cross-Border E-Commerce. *Int. J. Prod. Econ.* **2025**, *279*, 109460. [\[CrossRef\]](#)
12. Hua, L.; Wu, Y. Strategic Analysis of Vertical Integration in Cross-Border e-Commerce Logistics Service Supply Chains. *Transp. Res. Part E Logist. Transp. Rev.* **2024**, *188*, 103626. [\[CrossRef\]](#)
13. Bossong, P.; Reinhardt, A.; Elbert, R. Adoption Drivers and Barriers of Digital Freight Transport Platforms—An Intermodal Case Study. *Electron. Mark.* **2025**, *35*, 43. [\[CrossRef\]](#)
14. Turienzo, J.; Blanco, A.; Lampón, J.F.; Muñoz-Dueñas, M.d.P. Logistics Business Model Evolution: Digital Platforms and Connected and Autonomous Vehicles as Disruptors. *Rev. Manag. Sci.* **2024**, *18*, 2483–2506. [\[CrossRef\]](#)
15. Heinbach, C.; Beinke, J.; Kammler, F.; Thomas, O. Data-Driven Forwarding: A Typology of Digital Platforms for Road Freight Transport Management. *Electron. Mark.* **2022**, *32*, 807–828. [\[CrossRef\]](#)
16. Park, A.; Chen, R.; Cho, S.; Zhao, Y. The Determinants of Online Matching Platforms for Freight Services. *Transp. Res. Part E Logist. Transp. Rev.* **2023**, *179*, 103284. [\[CrossRef\]](#)
17. Tang, W.; Chen, X.; Lang, M.; Li, S.; Liu, Y.; Li, W. Optimization of Truck–Cargo Online Matching for the Less-Than-Truck-Load Logistics Hub under Real-Time Demand. *Mathematics* **2024**, *12*, 755. [\[CrossRef\]](#)
18. Haji Sami, E.; Shahnejat Bushehri, A.; Amirnia, A.; Yarahmadi, A.; Keivanpour, S. Integrated Sequential Matching and Routing Approach for Efficient and Eco-Friendly Freight Logistics. *Transp. Res. Part C Emerg. Technol.* **2025**, *179*, 105290. [\[CrossRef\]](#)
19. Chen, R.; Jia, S.; Meng, Q. Dynamic Container Drayage Booking and Routing Decision Support Approach for E-Commerce Platforms. *Transp. Res. Part E Logist. Transp. Rev.* **2023**, *177*, 103220. [\[CrossRef\]](#)
20. Wang, D.; Xie, X.; Moon, I. Two-Echelon Container Drayage with Multiple Transportation Modes and Overall-Network Time Interdependence. *Transp. Res. Part E Logist. Transp. Rev.* **2025**, *204*, 104436. [\[CrossRef\]](#)
21. Shiri, S.; Huynh, N. Optimization of Drayage Operations with Time-Window Constraints. *Int. J. Prod. Econ.* **2016**, *176*, 7–20. [\[CrossRef\]](#)
22. Cortés, C.E.; Matamala, M.; Contardo, C. The Pickup and Delivery Problem with Transfers: Formulation and a Branch-and-Cut Solution Method. *Eur. J. Oper. Res.* **2010**, *200*, 711–724. [\[CrossRef\]](#)
23. Kurtuluş, E. Optimizing Inland Container Logistics through Dry Ports: A Two-Stage Stochastic Mixed-Integer Programming Approach Considering Volume Discounts and Consolidation in Rail Transport. *Comput. Ind. Eng.* **2022**, *174*, 108768. [\[CrossRef\]](#)
24. Wei, W.; Dong, Z.; Fan, J. Integrated Location Selection and Scheduling Problems for Inland Container Transportation. *Sustainability* **2023**, *15*, 15992. [\[CrossRef\]](#)
25. Zhang, X.; Zha, X.; Dan, B.; Liu, Y.; Sui, R. Logistics Mode Selection and Information Sharing in a Cross-Border E-Commerce Supply Chain with Competition. *Eur. J. Oper. Res.* **2024**, *314*, 136–151. [\[CrossRef\]](#)

26. Fox, W.M. Sociotechnical System Principles and Guidelines: Past and Present. *J. Appl. Behav. Sci.* **1995**, *31*, 91–105. [[CrossRef](#)]
27. Agote-Garrido, A.; Martín-Gómez, A.M.; Lama-Ruiz, J.R. Manufacturing System Design in Industry 5.0: Incorporating Sociotechnical Systems and Social Metabolism for Human-Centered, Sustainable, and Resilient Production. *Systems* **2023**, *11*, 537. [[CrossRef](#)]

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.