Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/AN16.20260117
Published: June 23, 2026

Network Design for Multi-Depot Distribution Planning

Ayça Irmak, Doğa Mert, Doruk Uysal, İrem Arda, Mahmut Aydın and Yaren Turgut
Department of Industrial Engineering
Ozyegin University, Istanbul, Türkiye
ayca.irmak@ozu.edu.tr, doga.mert@ozu.edu.tr, doruk.uysal@ozu.edu.tr, irem.arda@ozu.edu.tr, mahmut.aydin@ozu.edu.tr, yaren.turgut@ozu.edu.tr

Milad Elyasi
Department of Industrial Engineering
Ozyegin University, Istanbul, Türkiye
milad.elyasi@ozyegin.edu.tr

Abstract

This study develops a decision-support framework for a real-world multi-depot distribution planning problem observed at a leading large-scale supermarket chain in Turkey. The operational setting involves two distribution centers, geographically dispersed retail stores, multiple product groups, pallet-based demand, heterogeneous vehicles, depot throughput limits, product-vehicle compatibility rules, time windows, and distance-dependent transportation costs. A mixed integer linear programming (MILP) model is first formulated to represent the integrated depot assignment, product allocation, vehicle activation, and routing decisions. Since the full model becomes computationally challenging as the number of stores increases, an Adaptive Large Neighborhood Search (ALNS)-based matheuristic is designed to explore assignment decisions while using exact optimization-based routing evaluations for candidate solutions. Computational experiments compare the ALNS framework with time-limited MILP runs on representative instances. The results indicate that the proposed method can produce feasible and competitive solutions within shorter computational times than the monolithic MILP on the tested instances. The study therefore positions the ALNS framework as a scalable analytical tool for tactical distribution planning rather than as an empirically validated source of cost savings against current operational practices, for which historical baseline data were not available.

Keywords
Retail logistics; Multi-depot vehicle routing; Adaptive large neighborhood search; Mixed integer linear programming; Heterogeneous fleet optimization

1. Introduction
Efficient distribution planning is a central requirement in large-scale retail logistics, where service reliability, fleet utilization, depot workload balance, and transportation cost must be managed simultaneously. In grocery and supermarket networks, these decisions are especially complex because stores require multiple product groups, vehicles differ in capacity and compatibility, and delivery schedules must respect time-window and route-duration restrictions. Static depot-to-store assignment policies may be simple to implement, but they can be inadequate when product availability, fleet capacity, and regional demand patterns vary across the planning horizon.

This paper studies an anonymized real-world case from a leading large-scale supermarket chain in Turkey. The case network contains two distribution centers serving a large set of stores in the Aegean Region. The planning problem requires deciding which depot should serve each store, which vehicles should be activated, how product-group demands should be allocated, and how feasible delivery routes should be constructed. Because all demand, vehicle capacity, and depot throughput values are represented in pallet units, the problem combines classical vehicle routing logic with practical retail distribution constraints.

The problem is modeled as a rich multi-depot vehicle routing problem with time windows (MDVRPTW). The model extends the standard routing structure by incorporating heterogeneous fleets, product-based store demand, depot-level inventory and dispatch capacities, vehicle-product compatibility, service times, and distance-dependent cost coefficients. A MILP formulation is developed to provide a rigorous optimization representation and to benchmark small instances. However, as the number of stores grows, the number of binary routing variables and timing constraints increases rapidly, making the monolithic MILP computationally difficult to solve within operationally useful time limits.

To address this scalability issue, the study proposes an ALNS-based matheuristic. The ALNS framework operates on high-level depot assignment decisions through destroy and repair operators, while a reduced exact routing subproblem evaluates candidate assignments. This design combines the modeling accuracy of exact optimization with the search flexibility of metaheuristics. The computational analysis focuses on the method's ability to obtain feasible and competitive solutions under limited computational budgets. Importantly, the paper does not claim verified cost reduction relative to the case company's current operations, because historical operational baseline data were not available for a direct before-and-after comparison.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on multi-depot routing, rich vehicle routing, and ALNS-based solution methods. Section 3 presents the problem definition, mathematical formulation, and ALNS design. Section 4 describes the anonymized data preparation and experimental instances. Section 5 discusses the computational results, including the interpretation of Table 1 and Figures 1 and 2. Sections 6 and 7 provide the conclusion, limitations, and future research directions.

2. Literature Review

2.1 Multi-Depot Vehicle Routing with Time Windows

The vehicle routing problem (VRP) and its variants are among the most extensively studied problems in operations research because they capture the operational trade-off between routing cost, fleet utilization, and customer service. The multi-depot vehicle routing problem extends the classical VRP by adding depot assignment decisions to route construction. This extension substantially increases problem complexity, since each customer must be associated with a depot and then sequenced within a feasible route. Reviews of multiple-depot VRP research emphasize that depot assignment, route design, fleet sizing, and operational constraints are often interdependent rather than separable decisions (Montoya-Torres et al. 2015).

The multi-depot vehicle routing problem with time windows introduces additional scheduling requirements by limiting when each customer can be served. Early and representative studies investigated variable neighborhood search, assignment-based heuristics, and exact or decomposition-based approaches for MDVRPTW variants (Polacek et al. 2004; Bae and Moon 2016; Li et al. 2016). More recent work has studied richer variants, including multi-depot multi-trip routing with time windows and release dates (Zhen et al. 2020), dynamic multi-depot routing (Wang et al. 2024), genetic and variable-neighborhood hybrid methods (Su et al. 2024), and multi-depot pickup-and-delivery applications with split deliveries and time windows (Dubey and Tanksale 2023). These studies show that practical MDVRPTW settings often require heuristic or hybrid methods because exact formulations become difficult at realistic scales.

Retail distribution systems add further operational richness. Product compatibility, pallet-based capacity, depot throughput, split demand allocation, and heterogeneous fleets make the problem more realistic but also more computationally demanding. Supply-chain network design and location-routing research has similarly emphasized that integrating distribution, inventory, and resource decisions can produce models that are more faithful to practice but harder to solve exactly (Azizi and Hu 2020; Ganesh et al. 2021; Alikhani et al. 2023).

2.2 Adaptive Large Neighborhood Search and Matheuristic Approaches

Large Neighborhood Search was introduced as a powerful destroy-and-repair framework for vehicle routing problems with side constraints (Shaw 1998). The central idea is to remove part of a solution and reconstruct it in a way that allows the search to move beyond small local changes. Adaptive Large Neighborhood Search later extended this idea by allowing multiple destroy and repair operators to compete, with their selection probabilities updated according to historical performance (Ropke and Pisinger 2006). This adaptive mechanism has made ALNS particularly suitable for rich routing problems because different operators may be effective in different regions of the search space.

ALNS has been widely adopted in routing and distribution settings because it provides a flexible structure for combining diversification and intensification. In the context of multi-depot and time-windowed routing, recent studies use ALNS or hybrid metaheuristics to handle dynamic requests, green routing constraints, heterogeneous fleets, and large-scale customer sets (Wang et al. 2024; Su et al. 2024). The method is also compatible with matheuristic designs, where exact optimization is embedded inside a heuristic framework to evaluate or improve partial solutions. Related inventory and replenishment studies show that matheuristics can be effective when a problem combines realistic resource constraints with combinatorial decision structures (Elyasi et al. 2025).

The present study follows this design tradition. The proposed ALNS is not presented as a set of individually validated components; instead, its parallel multi-start runs, adaptive operator weighting, stagnation handling, elite-pool crossover, and regret-based repair are introduced as theoretically motivated design choices grounded in the ALNS literature. Since no component-level ablation study is conducted, the computational results should be interpreted as evidence for the overall framework rather than as proof that each architectural feature independently improves performance.

2.3 Positioning of the Present Study

This study contributes to the applied optimization literature by formulating and solving an anonymized, industry-inspired, multi-depot retail distribution problem that combines depot assignment, pallet-based product allocation, heterogeneous fleet routing, and time windows. The MILP formulation provides a transparent mathematical representation of the operational setting, while the ALNS-based matheuristic offers a scalable solution strategy for larger instances. The work is positioned as a decision-support methodology and computational feasibility study, not as an empirical validation of realized savings in the case company's live operations.

3. Methods

3.1 Problem Definition

The case problem concerns a two-depot retail distribution system. The depots supply a set of stores with demands for multiple product groups. Each product group has limited stock at each depot, and each depot has a maximum daily dispatching capacity expressed in pallet units. Vehicles are heterogeneous in capacity, type, travel-speed assumptions, fixed operating cost, and product compatibility. Each route starts and ends at the vehicle's assigned depot, and each store must be served within its time window. Service time is proportional to the number of pallets delivered.

Transportation cost has two main components. The first is a pallet-kilometer cost applied to loaded travel, where the cost coefficient differs between urban and intercity arcs according to a distance threshold. The second is an empty-return cost applied when a vehicle returns from its final store to its depot. A penalty is included for unmet demand to preserve model feasibility when available stock or operational capacity is insufficient. The main planning decisions are:

- assignment of stores and product-group quantities to feasible depots and vehicles;
- activation of vehicles from their assigned depots;
- construction of feasible routes satisfying capacity, flow, time-window, and route-duration restrictions;
- minimization of total fixed, loaded-distance, empty-return, and shortage-penalty costs.

3.2 Mathematical Formulation

The MILP formulation below represents the integrated assignment, allocation, and routing problem. The notation is written explicitly to improve transparency and reproducibility.

Sets

D	set of depots $(d \in D)$
S	set of stores ($i,j \in$ S)
G	set of product groups ($g \in$ G)
V	set of vehicles ($v \in$ V)

K	set of vehicle types ($k \in$ K)
N	set of all nodes ($N = D\ S;\ i, j, h\ \in N$)
A	set of directed arcs
V_d	set of vehicles based at depot d

Parameters

$dist_{ij}$	distance between nodes i and j (km)
Q_v	pallet capacity of vehicle v
dem_{ig}	pallet demand of product group g at store i
$stock_{dg}$	stock of product group g at depot d
P_d^{max}	maximum daily pallet throughput of depot d
$d(v)$	starting/ending depot of vehicle v
K	critical distance where cost structure changes (km)
$penalty$	penalty cost for one unsatisfied pallet demand
$\mathrm{c}_{\mathrm{empty}}$	cost per km for an empty return trip
$\epsilon_{\mathrm{pallet}}$	minimum delivery quantity to a store
$\mathrm{c}^{\mathrm{city}}$	cost coefficient for inner city transportation
$\mathrm{c}^{\mathrm{out}}$	cost coefficient for out of city transportation
F_k	fixed cost of vehicle type k
T_k^{max}	maximum route time (hours) per vehicle type k
$\mathrm{speed}_{\mathrm{k}}^{\mathrm{city}}$	inner-city speed of vehicle type k (km/hour)
$\mathrm{speed}_{\mathrm{k}}^{\mathrm{out}}$	out-of-city speed of vehicle type k (km/hour)
α	service time per pallet (hours)
M	a sufficiently large constant
$k(v)$	type of vehicle v

Decision Variables

x_{vij}	$\in \{0,1\}$: 1 if vehicle v travels from node i to node j
w_{vi}	$\in \{0,1\}$: 1 if vehicle v visits store i
z_{vig}	≥ 0 : Quantity of product g delivered by vehicle v to store i
L_{vij}	≥ 0: Load carried by vehicle v on arc (i, j)
p_{vi}	≥ 0: Delivered by vehicle v at store i
s_{ig}	≥ 0: Unsatisfied pallet demand at store i
t_{vi}	≥ 0: Arrival time of vehicle v at node i (hours)
st_{vi}	≥ 0: Service time of vehicle v at store i (hours)

Objective Function

$$\min Z = \sum_{v \in V} F_{k(v)} \left(\sum_{j \in N \setminus \{d(v)\}} x_{v,d(v),j} \right) + \sum_{v \in V} \sum_{(i,j) \in A} \gamma_{ij} \, L_{vij} + \sum_{v \in V} \sum_{i \in S} c_{empty} \, dist_{i,d(v)} \, x_{v,i,d(v)} + penalty \sum_{i \in S} \sum_{g \in G} s_{ig}$$

S.t.

$$\sum_{v \text{ in } V} z_{vig} + s_{ig} = dem_{ig} \qquad \forall i \in S, g \in G \quad (1)$$

$$\sum_{i \in S} \sum_{v \in V_d} z_{vig} \leq stock_{dg} \qquad \forall d \in D, \ g \in G \quad (2)$$

$$\sum_{i \in S} \sum_{g \in G} \sum_{v \in V_d} z_{vig} \leq P_d^{max} \qquad \forall d \in D \quad (3)$$

$$p_{vi} = \sum_{g \in G} z_{vig} \qquad \forall v \in V, \ i \in S \quad (4)$$

$$p_{vi} \leq dem_i \, w_{vi} \qquad \forall v \in V, \ i \in S \quad (5)$$

$$p_{vi} \geq \varepsilon_{\text{pallet}} \, w_{vi} \qquad \forall v \in V, \ i \in S \quad (6)$$

$$\sum_{j \in N \setminus \{d(v)\}} x_{v,d(v),j} \leq 1 \qquad \forall v \in V \quad (7a)$$

$$\sum_{i \in N \setminus \{d(v)\}} x_{v,i,d(v)} \leq 1 \qquad \forall v \in V \quad (7b)$$

$$\sum_{j \in N \setminus \{i\}} x_{vji} = w_{vi} \qquad \forall v \in V, \ i \in S \quad (8a)$$

$$\sum_{j \in N \setminus \{i\}} x_{vij} = w_{vi} \qquad \forall v \in V, \ i \in S \quad (8b)$$

$$\sum_{j \in N \setminus \{d(v)\}} L_{v,d(v),j} - \sum_{i \in N \setminus \{d(v)\}} L_{v,i,d(v)} = \sum_{i \in S} p_{vi} \qquad \forall v \in V \quad (9)$$

$$\sum_{h \in N \setminus \{i\}} L_{vhi} - \sum_{j \in N \setminus \{i\}} L_{vij} = p_{vi} \qquad \forall v \in V, \ i \in S \quad (10)$$

$$L_{vij} \leq Q_v x_{vij} \qquad \forall v \in V, \ (i,j) \in A \quad (11)$$

$$\sum_{i \in S} p_{vi} \leq Q_v \qquad \forall v \in V \quad (12)$$

$$\sum_{(i,j) \in A} dist_{ij} x_{vij} \leq T_{k(v)}^{max} \qquad \forall v \in V \quad (13)$$

$$st_{vi} = \alpha \sum_{g \in G} z_{vig} \qquad \forall v \in V, \ i \in S \quad (14)$$

$$t_{vj} \geq t_{vi} + st_{vi} + T_{ij}^{k(v)} - M\left(1 - x_{vij}\right) \qquad \forall v \in V, \ (i,j) \in A \quad (15)$$

$$t_{vi} \geq E_i - M(1 - w_{vi}) \qquad \forall v \in V, \ i \in S \quad (16)$$

$$t_{vi} + st_{vi} \leq L_i + M(1 - w_{vi}) \qquad \forall v \in V, \ i \in S \quad (17)$$

$$\sum_{(i,j) \in A} \mathrm{T}_{\mathrm{ij}}^{\mathrm{k(v)}} \mathrm{x}_{\mathrm{vij}} + \sum_{i \in S} \mathrm{st}_{\mathrm{vi}} \leq \mathrm{T}_{\mathrm{k(v)}}^{\mathrm{max}} \qquad \forall \mathrm{v} \in \mathrm{V} \quad (18)$$

$$x_{vij}, \ w_{vi} \in \{0,1\} \qquad \forall v \in V, \ (i,j) \in A \quad (19)$$

$$z_{vig}, \ L_{vij}, \ p_{vi}, \ s_{ig}, \ t_{vi}, st_{vi}, \ \geq \ 0 \qquad \forall d \in D, \ g \in G, v \in V, \ (i,j) \in A \quad (20)$$

Equation (1) ensures demand balance by stating that the total amount delivered by all vehicles, together with any shortage, must satisfy the demand of each store for each product group. Equation (2) limits the amount shipped from each depot for each product group by the available stock, while Equation (3) restricts the total outbound quantity from each depot by its overall handling capacity. Equation (4) defines the total quantity delivered by a vehicle to a store as the sum of the delivered amounts over all product groups. Equation (5) links delivery and visit decisions by ensuring that a vehicle can deliver to a store only if it visits that store, and Equation (6) imposes a minimum pallet quantity whenever a visit occurs. Equations (7a) and (7b) ensure that each vehicle departs from and returns to its assigned depot at most once. Equations (8a) and (8b) preserve route continuity by requiring that, for every visited store, exactly one outgoing arc and one incoming arc are selected. Equation (9) establishes load balance at the depot by setting the initial vehicle load equal to the total quantity delivered along its route, while Equation (10) enforces load conservation at each store by reducing the remaining onboard load according to the quantity delivered there. Equation (11) ensures that the load carried on an arc is positive only if that arc is used and does not exceed vehicle capacity, whereas Equation (12) limits the total delivered quantity of each vehicle by its capacity. Equation (13) restricts the total traveled distance of each vehicle by its maximum allowable route length. Equation (14) defines the service time at each store as proportional to the delivered quantity. Equation (15) propagates arrival times along the route and prevents sub-tours through a big-M time relationship. Equations (16) and (17) enforce the lower and upper bounds of the store time windows, respectively. Equation (18) limits the total working time of each vehicle, including both travel and service times. Finally, Equations (19) and (20) define the variable domains by imposing binary restrictions on routing and visit variables, and non-negativity on delivery, load, shortage, arrival time, and service time variables.

3.3 ALNS-Based Matheuristic Design

The ALNS framework is designed to overcome the scalability limitations of the monolithic MILP. It decomposes the decision process into an assignment layer and a routing-evaluation layer. In the assignment layer, each store is temporarily assigned to a depot or marked as unserved when feasibility cannot be maintained. In the routing layer, the fixed depot-level customer sets are evaluated by reduced optimization models that enforce vehicle capacity, time windows, service times, and routing feasibility. The total objective value of a candidate solution is obtained by aggregating depot-level routing costs and unmet-demand penalties.

3.3.1 Solution Representation and Initial Solutions

A solution is represented as a vector of store-to-depot assignments. This compact representation allows the search to modify strategic allocation decisions while avoiding direct manipulation of all route arcs at every iteration. Initial solutions are generated using several construction rules, including nearest-depot assignment, balanced sequential assignment, and randomized assignment. The purpose of the multi-start procedure is to diversify the initial search state. Since no separate comparative run is conducted for initialization strategies, the multi-start mechanism is interpreted as a robustness-oriented design element rather than as an independently validated improvement.

3.3.2 Destroy Phase

At each iteration, a subset of store assignments is removed from the current solution. The destroy operators include random removal, related removal based on geographic proximity, similarity-based removal using demand and time-window attributes, worst-cost removal, depot-focused removal, and load-balancing removal. These operators are included to represent different search behaviors. Random and depot-focused removals support diversification, while related, similarity-based, and worst-cost removals support more targeted reconstruction.

3.3.3 Repair Phase

Removed stores are reinserted through greedy and regret-based repair rules. Greedy repair assigns a removed store to the feasible depot that produces the smallest estimated incremental cost. Regret-based repair gives priority to stores for which postponing insertion would create a large difference between the best and second-best feasible alternatives. This mechanism is motivated by the risk that purely greedy insertion can make early decisions that block later feasibility or consolidation opportunities. Restricted candidate lists and controlled stochastic choices are used to preserve exploration.

3.3.4 Local Search, Adaptive Selection, and Acceptance

After repair, depot-level routing solutions may be improved through route-level local search, such as edge-exchange operations, whenever feasibility is preserved. Destroy and repair operators are selected adaptively using weights updated according to their historical contribution to accepted or improved solutions, following the general logic of ALNS. A simulated annealing acceptance rule allows some worsening moves early in the search and gradually becomes more selective as the temperature decreases. These components are standard design choices in ALNS-style algorithms and are not interpreted as separately proven performance drivers in this study.

3.3.5 Stagnation Handling, Elite Pool, and Parallel Runs

If no improvement is observed over a predefined number of iterations, the destruction rate is increased and more disruptive operators are emphasized to help the search escape local optima. An elite pool stores high-quality and structurally diverse assignment vectors. When stagnation occurs, crossover between elite assignments can generate new candidate solutions that combine favorable depot-allocation patterns. Parallel runs with different random seeds are used to explore different regions of the solution space. These mechanisms are included based on established metaheuristic design rationale; their isolated contributions should be examined through future ablation studies.

4. Data Preparation and Experimental Setting

The data used in this study are derived from an anonymized real-world retail distribution setting. To comply with data privacy and confidentiality requirements, the identity of the case company, proprietary operational identifiers, exact store names, and commercially sensitive details are not disclosed. The company is referred to only as the case company or as a leading large-scale supermarket chain in Turkey. All results are reported at an aggregate computational level rather than as store-level operational recommendations.

The full case network consists of two depots and a large set of retail stores. Each store has product-group demand expressed in pallets, while each depot has limited product-group stock and a maximum daily dispatching capacity. Vehicle data include capacity, type, fixed cost, compatibility with product groups, and route-duration restrictions. Distances are computed from anonymized node coordinates and are used to estimate both transportation cost and travel time. A distance threshold distinguishes urban and intercity arcs, which differ in pallet-kilometer cost and assumed travel speed.

For computational testing, representative subsets of the full network are generated at different store sizes. These instances allow the exact MILP and the ALNS matheuristic to be compared on problems of increasing size and structural complexity. The MILP is solved under predefined time limits, and the ALNS is run using the proposed destroy-repair framework. Because the MILP results are time-limited, the reported MILP values should be interpreted as benchmark incumbents rather than certified global optima for every instance.

5. Results and Discussion

5.1 Numerical Results

Table 1 reports the computational comparison between the time-limited MILP benchmark and the proposed ALNS matheuristic. The table includes the instance size, MILP time limit, MILP incumbent cost, ALNS objective value, ALNS CPU time, and the reported gap indicator. The results show that the ALNS framework consistently returns feasible solutions within shorter CPU times than the MILP time limit. However, the cost differences should not be interpreted as demonstrated savings relative to current company operations; they only compare alternative optimization procedures under the tested computational setting.

Table 1. Comparison of Time-Limited MILP and ALNS Results

Instance	Exact Time Limit (s)	MILP Cost	ALNS Cost	ALNS CPU (s)	GAP (%)
20 stores (2M 1K)	3600	12,848	9,028	768	36.10%
25 stores (2M 2K)	3600	13,932	10,674	652	8.47%
35 stores (2M 1K)	1200	20,004	15,983	454	34.41%
35 stores (3M 2K)	1200	20,288	16,050	706	35.19%
40 stores (3M 2K)	3600	22,781	16,320	1081	40.22%
45 stores (3M 2K)	1600	23,107	16,656	719	17.83%
45 stores (5M 5K)	1200	22,449	19,837	221	34.20%
50 stores (3M 2K)	1200	26,146	21,409	402	25.50%
50 stores (5M 5K)	1200	27,125	21,479	402	39.22%

The gap values in Table 1 vary considerably across instances. For example, the reported gap is 8.47% in the 25-store instance, but it increases to 36.10% in the 20-store instance and 40.22% in the 40-store instance. This variation indicates that instance structure is more important than store count alone. A smaller instance may still be difficult if demand is concentrated near one depot, if product-vehicle compatibility restricts feasible assignments, or if time windows leave little flexibility for route consolidation. Conversely, a larger instance can be relatively easier when stores are geographically clustered, depot workloads are balanced, and product compatibility patterns align well with the available fleet.

The 25-store instance appears to have a comparatively favorable structure: its lower gap suggests that the MILP benchmark and ALNS search identify similar assignment and routing patterns. In contrast, the 20-store instance has a much larger reported gap despite its smaller size, which may reflect tighter time-window interactions, uneven depot utilization, or compatibility constraints that restrict the MILP's ability to find strong incumbents within the time limit. The 40-store instance also exhibits a high reported gap, which is consistent with the rapid growth of routing variables and timing constraints as the customer set expands. These results emphasize that MDVRPTW difficulty is driven by the joint interaction of depot balance, store geography, demand composition, product compatibility, and time-window tightness.

The paired 45-store and 50-store cases further support this interpretation. The two 45-store instances have different gap patterns, suggesting that adding vehicle types or compatibility classes does not automatically simplify the problem. If additional fleet categories introduce more assignment alternatives but also more compatibility restrictions, the search landscape may become more fragmented. Similarly, the two 50-store instances show different gap indicators under the same nominal store count. Therefore, computational performance should be evaluated across multiple structural scenarios rather than by relying only on customer count.

5.2 Graphical Results

Figure 1 visualizes the CPU time comparison between the time-limited MILP runs and the ALNS algorithm. The chart shows that the ALNS framework requires less computational time than the MILP time limits across the reported instances. This supports the main computational motivation of using ALNS: it provides a practical way to obtain feasible solutions when the exact formulation becomes time-consuming.

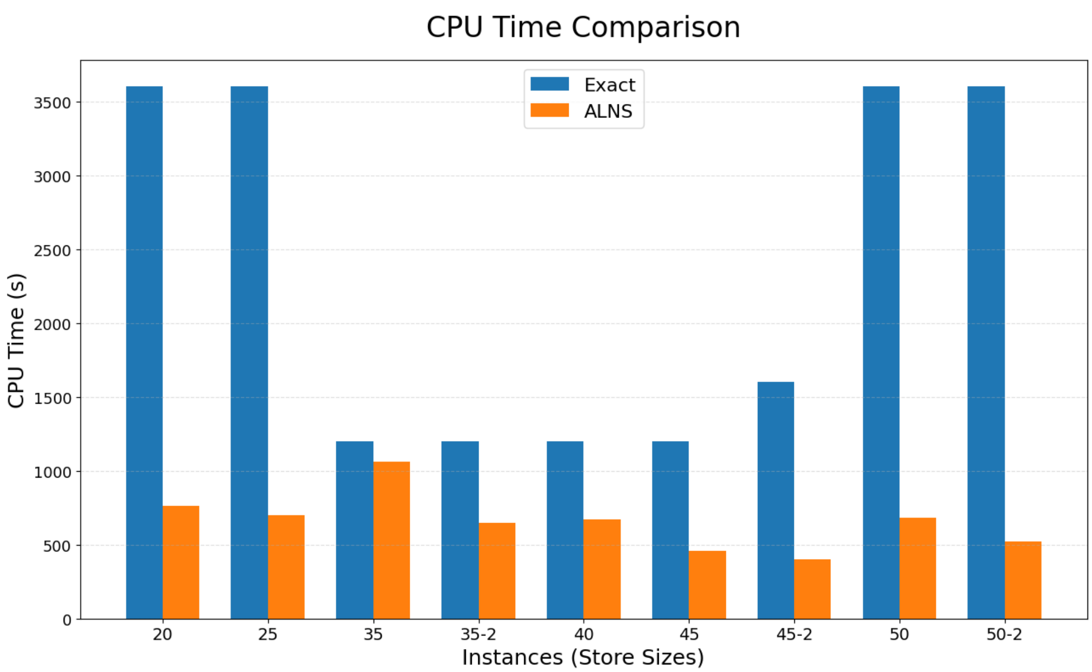

Figure 1. CPU time comparison between the time-limited MILP benchmark and the proposed ALNS matheuristic across tested store-size instances.

Figure 2 presents the corresponding objective-value comparison. The figure should be read together with Table 1. The lower ALNS objective values relative to MILP incumbents indicate that the ALNS search can identify strong feasible solutions under the tested time budgets. However, because the MILP values are time-limited incumbents and not necessarily proven optima, Figure 2 should not be interpreted as an optimality guarantee or as evidence of realized savings in live operations.

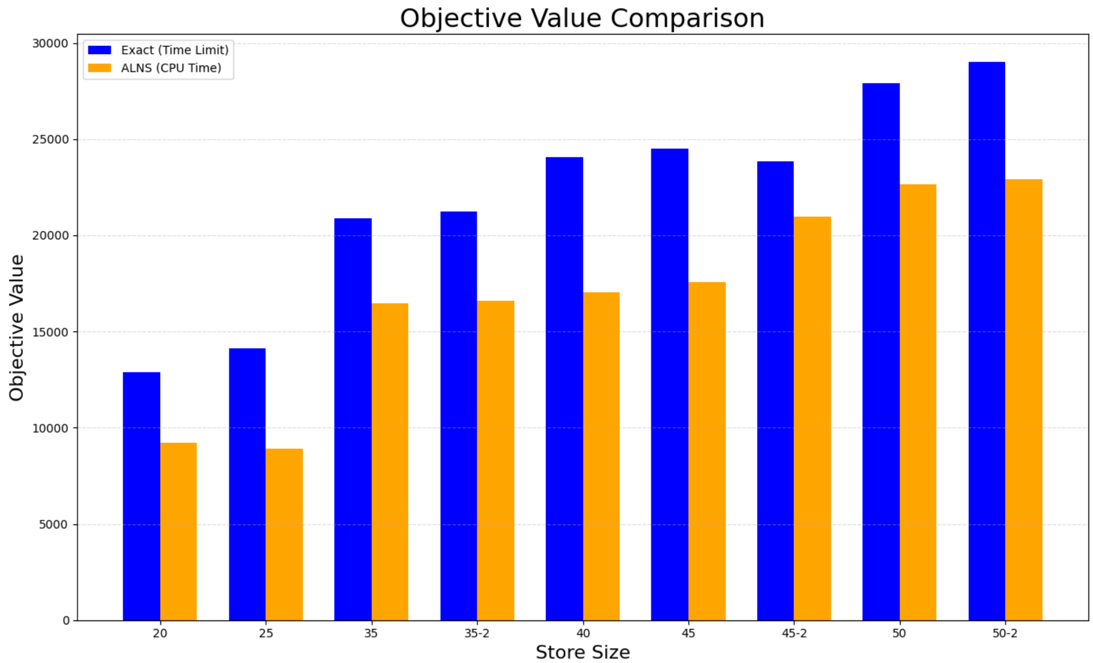

Figure 2. Objective-value comparison between the time-limited MILP benchmark and the proposed ALNS matheuristic.

Overall, the computational results suggest that the ALNS-based matheuristic is suitable as a scalable decision-support approach for rich multi-depot distribution planning. The strongest supported conclusion is computational: the monolithic MILP becomes increasingly difficult under time limits, while the ALNS framework can return feasible and competitive solutions more quickly. The analysis does not validate individual ALNS components through controlled ablation experiments, nor does it prove realized operational savings against the case company's current routing policies.

6. Conclusion

This study formulated and analyzed a rich multi-depot vehicle routing problem with time windows based on an anonymized retail distribution case in Turkey. The problem includes depot assignment, product-group allocation, heterogeneous vehicle capacities, product-vehicle compatibility, pallet-based demand, depot stock and dispatch constraints, service times, route-duration limits, and distance-dependent transportation costs. A MILP model was developed to represent the full optimization problem, and an ALNS-based matheuristic was proposed to address larger instances that are difficult for the monolithic exact model.

The computational results show that the ALNS framework can generate feasible and competitive solutions within shorter CPU times than the time-limited MILP benchmark on the tested instances. The findings mainly support the use of ALNS as a scalable computational strategy for complex distribution planning. Claims about transportation cost reductions relative to existing operational practices are deliberately avoided because historical baseline data from the case company were not available. Similarly, architectural elements of the ALNS design are explained as literature-based and theoretically motivated choices rather than as individually validated performance enhancements.

The study demonstrates that combining exact optimization and adaptive metaheuristic search can provide a practical framework for tactical distribution planning in large-scale retail networks. The proposed approach can help decision makers evaluate alternative depot-assignment and routing scenarios under operational constraints, while maintaining transparency through the underlying mathematical model.

7. Limitations and Future Work

The first limitation of this study is the absence of historical operational baseline data. Because current route plans, realized transportation costs, service performance, and dispatch decisions were not available in a comparable format, the study cannot empirically quantify cost savings against actual company operations. Future research should conduct a controlled real-world validation using historical or pilot-period operational data, comparing proposed routes with executed routes under identical demand and resource conditions.

A second limitation is that the ALNS components are evaluated as a complete framework rather than through component-level ablation experiments. Future work should compare versions of the algorithm with and without adaptive operator weighting, elite-pool crossover, regret-based repair, stagnation handling, and parallel multi-start runs. Such experiments would clarify which components contribute most strongly to performance under different instance structures.

Third, the current study uses representative computational instances derived from an anonymized case network. Future studies could expand the experimental design by incorporating additional demand days, stochastic demand, real travel times, driver regulations, vehicle availability calendars, and rolling-horizon reoptimization. These extensions would make the model more suitable for operational deployment and would allow the decision-support framework to be validated in dynamic retail logistics environments.

Acknowledgments

The authors thank the anonymized case company for providing operational insights and data support for this study. The authors also acknowledge the Department of Industrial Engineering at Ozyegin University for its academic guidance and support throughout the project.

References

Alikhani, R., Eskandarpour, M., and Jahani, H., Collaborative distribution network design with surging demand and facility disruptions, International Journal of Production Economics, vol. 262, article 108912, 2023.

Azizi, V., and Hu, G., Multi-product pickup and delivery supply chain design with location-routing and direct shipment, International Journal of Production Economics, vol. 226, article 107648, 2020.

Bae, H., and Moon, I., Multi-depot vehicle routing problem with time windows considering delivery and installation vehicles, Applied Mathematical Modelling, vol. 40, no. 13-14, pp. 6536-6549, 2016.

Borajee, M., Madani-Saatchi, S. H., and Tavakkoli-Moghaddam, R., A chance-constraint optimization model for a multi-echelon multi-product closed-loop supply chain considering brand diversity: An accelerated Benders decomposition algorithm, Computers & Operations Research, vol. 152, article 106130, 2023.

Crevier, B., Cordeau, J. F., and Laporte, G., The multi-depot vehicle routing problem with inter-depot routes, European Journal of Operational Research, vol. 176, no. 2, pp. 756-773, 2007.

Dubey, N., and Tanksale, A., A multi-depot vehicle routing problem with time windows, split pickup and split delivery for surplus food recovery and redistribution, Expert Systems with Applications, vol. 232, article 120807, 2023.

Elyasi, M., Ekici, A., Altan, B., and Özener, O. Ö., A matheuristic for the joint replenishment problem with and without resource constraints, Annals of Operations Research, vol. 351, pp. 727-749, doi: 10.1007/s10479-024-06230-y, 2025.

Fathollahi-Fard, A. M., Jolai, F., Seydanlou, P., and Tavakkoli-Moghaddam, R., A multi-objective optimization framework for a sustainable closed-loop supply chain network in the olive industry: Hybrid meta-heuristic algorithms, Expert Systems with Applications, vol. 203, article 117566, 2022.

Ganesh, K., Mendoza, A., Nambirajan, R., Narendran, T. T., Olivares-Benitez, E., and Pazhani, S., Multi-period multi-product closed loop supply chain network design: A relaxation approach, Computers & Industrial Engineering, vol. 155, article 107191, 2021.

Li, J., Li, Y., and Pardalos, P. M., Multi-depot vehicle routing problem with time windows under shared depot resources, Journal of Combinatorial Optimization, vol. 31, pp. 515-532, 2016.

Montoya-Torres, J. R., López Franco, J., Nieto Isaza, S., Felizzola Jiménez, H., and Herazo-Padilla, N., A literature review on the vehicle routing problem with multiple depots, Computers & Industrial Engineering, vol. 79, pp. 115-129, 2015.

Polacek, M., Hartl, R. F., Doerner, K., and Reimann, M., A variable neighborhood search for the multi depot vehicle routing problem with time windows, Journal of Heuristics, vol. 10, pp. 613-627, 2004.

Ropke, S., and Pisinger, D., An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows, Transportation Science, vol. 40, no. 4, pp. 455-472, doi: 10.1287/trsc.1050.0135, 2006.

Shaw, P., Using constraint programming and local search methods to solve vehicle routing problems, in Maher, M., and Puget, J. F. (eds.), Principles and Practice of Constraint Programming - CP98, Lecture Notes in Computer Science, vol. 1520, Springer, Berlin, Heidelberg, pp. 417-431, doi: 10.1007/3-540-49481-2_30, 1998.

Su, Y., Zhang, S., and Zhang, C., A lightweight genetic algorithm with variable neighborhood search for multi-depot vehicle routing problem with time windows, Applied Soft Computing, vol. 161, article 111789, 2024.

Wang, S., Sun, W., and Huang, M., An adaptive large neighborhood search for the multi-depot dynamic vehicle routing problem with time windows, Computers & Industrial Engineering, vol. 191, article 110122, doi: 10.1016/j.cie.2024.110122, 2024.
Zhen, L., Ma, C., Wang, K., Xiao, L., and Zhang, W., Multi-depot multi-trip vehicle routing problem with time windows and release dates, Transportation Research Part E: Logistics and Transportation Review, vol. 135, article 101866, 2020.

Biographies

Ayça Irmak is a senior student in the Industrial Engineering Department at Ozyegin University, Istanbul, Türkiye. Her research interests include supply chain management, operations research, production planning, and logistics optimization. She has completed internship programs at Tristone Flowtech Çerkezköy and EAE Electric, gaining experience in production planning, logistics operations and manufacturing processes..

Doğa Mert is a senior student in the Industrial Engineering Department at Ozyegin University, Istanbul, Türkiye. Her research interests include supply chain management, operations research, and logistics optimization. She has completed internship programs at Wilo and ADOTEK**,** gaining experience in production planning and data analysis.

Doruk Uysal is a senior student in the Industrial Engineering Department at Ozyegin University, Istanbul, Türkiye. His research interests include supply chain management, operations research, logistics optimization, and financial analysis. He worked in material management engineering at Vestel, in the production department at Bakgravür, and currently works as an intern in the finance department at Fiba Renewable Energy, gaining experience in operational processes, and financial analysis.

İrem Arda is a senior student in the Industrial Engineering Department at Ozyegin University, Istanbul, Türkiye. Her research interests include supply chain management, operations research, logistics optimization, and data analytics. She worked in data analysis at Migros Ticaret A.Ş. and in revenue management at Corendon Airlines, gaining experience in business intelligence, demand analysis, and operational decision-making.

Mahmut Aydın is a senior student in the Industrial Engineering Department at Ozyegin University, Istanbul, Türkiye. He is actively involved in the plastics and recycling industry as the founder of UCD Geri Dönüşüm, a company specializing in PET recycling and PET sheet production. His work focuses on building integrated production systems that combine recycling, material processing, and packaging solutions.His professional interests include supply chain management, operations research, production planning, and logistics optimization. He has hands-on experience in manufacturing operations, injection molding, and large-scale material flow management.He has completed internship programs at Sanko Tekstil and Eczacıbaşı VitrA Karo, where he gained practical experience in industrial production systems, process optimization, and operational workflows in large-scale manufacturing environments.

Yaren Turgut is a senior student in the Industrial Engineering Department at Ozyegin University, Istanbul, Türkiye.Her research interests include supply chain management, operations research, production planning, and logistics optimization.She has completed internship programs at Tat Nişasta and Sunar Mısır, gaining experience in production processes and planning operations.Currently, she works as a Marketing Intern at Varsapp, where she focuses on startup business models.

Milad Elyasi is an Assistant Professor in the Industrial Engineering Department at Özyeğin University, Istanbul, Turkey. He received his Ph.D. in Industrial Engineering, and his research interests include operations research, mathematical optimization, transportation and energy systems, and decision analytics. His work focuses on the development of optimization models and algorithms for complex industrial and energy-related problems, including sustainable manufacturing, energy-efficient system design, and large-scale optimization under uncertainty. He has published in international journals and conference proceedings and actively supervises graduate research in optimization and applied operations research.