A Branch-and-Price Algorithm for the Two-Echelon Vehicle Routing Problem with Time Windows and Time-Dependent Travel Times

Jiří Spitzer*[1], Gonçalo Homem de Almeida Correia[1], Serge P. Hoogendoorn[1], and Shadi Sharif Azadeh[1]

[1]Transport & Planning Department, Delft University of Technology, Netherlands

Short summary

Urban logistics frequently relies on two-level distribution networks, where large vehicles deliver goods to intermediate facilities from which smaller vehicles complete last-mile deliveries in access-restricted city centers, where travel times vary considerably throughout the day. Fluctuating traffic conditions directly affect route planning and time window feasibility. Yet most studies on vehicle routing simplify or ignore this effect, assuming static travel times throughout the planning horizon. In this paper, we propose a branch-and-price algorithm for the two-echelon vehicle routing problem with time windows and time-dependent travel times. A baseline model is formulated and evaluated through a preliminary computational study, yielding promising results that demonstrate the viability of the approach. We further discuss the limitations of the current model and outline targeted improvements aimed at solving instances of practical size to optimality, as well as establishing solid benchmarks for emerging heuristics.
Keywords: branch-and-price, time-dependent, two-echelon, vehicle routing

1 Introduction

Urban logistics poses operational challenges driven by physical constraints such as restricted road access, limited parking infrastructure, and increasing regulatory measures aimed at reducing the presence of large vehicles in city centers. Such restrictions are motivated by growing concerns over road safety and the environmental impact. An often-used approach to tackle these challenges is to restructure the distribution network into two levels, where large vehicles transport freight from a central depot to intermediate distribution facilities (also called satellites or transshipment points), typically located at the periphery of the urban area, from which smaller vehicles complete the delivery to end customers within the city center. The problem of jointly optimizing the routing of vehicles across both levels is known in the literature as the two-echelon vehicle routing problem (2E-VRP) (Sluijk et al., 2023).

Despite the practical relevance of time-dependent travel times in urban logistics, research on the time-dependent variant of the 2E-VRP remains scarce (Soysal et al., 2015; Sluijk et al., 2023). Oversimplified assumptions that ignore dynamic traffic conditions in urban environments where travel times fluctuate considerably between peak and off-peak hours may lead to insufficient solutions (Ichoua et al., 2003). To the best of our knowledge, the only exact solution method for the 2E-VRP with time-dependent travel times (2E-TDVRP) is an MILP model proposed by Soysal et al. (2015). Their study focuses on environmental and social consequences of logistics operations, such as energy usage, congestion, and traffic-related air pollution. However, they only solved small instances up to two satellites and 16 customers. More recent studies focused on approximate solution methods. Examples can be found in Gutierrez et al. (2024) and Reyes et al. (2024).

To the best of our knowledge, this study is the first to propose an exact branch-and-price algorithm for the two-echelon vehicle routing problem with time windows and time-dependent travel times (2E-TDVRP-TW). Exact methods play a dual role in the operations research literature: they not only produce exact solutions, but also establish rigorous benchmarks against which heuristic and metaheuristic approaches can be evaluated. In this paper, we present a baseline algorithm and evaluate the results of an initial computational study on a set of small instances. We discuss the current limitations and outline concrete enhancements expected to improve the algorithm's performance, enabling it to solve instances of a size comparable to those presented in Mhamedi et al. (2022).

2 Methodology

Our algorithm builds upon the branch-and-price algorithm proposed by Mhamedi et al. (2022), which we extend to accommodate time-dependent travel times in the second-echelon. This section describes the baseline formulation and solution method, covering the problem formulation, column generation procedure, pricing problem, and branching strategy. The baseline model presented here constitutes a functional starting point.

Mathematical formulation

The problem is formulated on a directed graph $G = (N, A)$, where N denotes the nodes and A the set of arcs. The set of nodes is partitioned as $N = N^D \cup N^S \cup N^C$, where N^D, N^S, and N^C are disjoint sets representing depots, satellites, and customers, respectively. Each customer $i \in N^C$ has a demand q_i and must be visited within time window $[\underline{w}_i, \bar{w}_i]$, where $\underline{w}_i$ denotes the earliest service time and $\bar{w}_i$ the latest completion time. A vehicle arriving before $\underline{w}_i$ must wait until service can begin. Each node is associated with a service time θ, with $\theta_d = 0, \forall d \in N^D$, and $\theta_i > 0, \forall i \in N^S \cup N^C$.
The graph G can be split into two parts corresponding to the first and second echelons. The first-echelon includes nodes N^D and N^S, and the corresponding set of arcs denoted $A^{(1)}$, while the second-echelon network contains nodes N^S and N^C, and the corresponding arcs denoted $A^{(2)}$. The set of the first-echelon arcs is defined as $A^{(1)} = \{(d, s) \cup (s, d) : d \in N^D, s \in N^S\} \cup \{(s_1, s_2) : s_1, s_2 \in N^S, s_1 \neq s_2\}$, and the set of second-echelon arcs as $A^{(2)} = \{(s, i) \cup (i, s) : s \in N^S, i \in N^C\} \cup \{(i_1, i_2) : i_1, i_2 \in N^C, i_1 \neq i_2\}$. The complete set of arcs A is defined as $A = A^{(1)} \cup A^{(2)}$.
We consider a homogeneous fleet of first-echelon vehicles with capacity $Q^{(1)}$. Each first-echelon route starts from a depot at time 0, supplies one or more satellites, and return to the same depot. Similarly, a homogeneous fleet of vehicles with capacity $Q^{(2)}$ is considered in the second-echelon, where each route starts from a satellite at the earliest possible time, visits one or more customers, and returns to the same satellite. We assume that $0 < Q^{(2)} < Q^{(1)}$. Furthermore, following the assumption of Dellaert et al. (2019) and Mhamedi et al. (2022), each second-echelon route must be supplied by exactly one first-echelon vehicle. A first- and second-echelon route is feasible if each node is visited at most once and vehicle capacity is not exceeded. Additionally, second-echelon routes must respect customer time windows.
The goal is to find a set of feasible first- and second-echelon routes so that each customer is visited exactly once and the total travel time is minimized.
Let M and L denote the sets of all feasible first- and second-echelon routes, respectively. Let $L_m \subseteq of L$ denote the subset of second-echelon routes supplied by first-echelon route $m \in M$. Let $c_m \forall m \in M$ and $c_l \forall l \in L$ be the total travel time of the corresponding routes, and a_{il} be the binary parameter equal to 1 if customer i is visited on route l, and 0 otherwise.
We define binary decision variables x_m and y_l, which equal 1 if and only if route m and route l is selected, respectively. The problem can be formulated as a set partitioning problem as follows

$$\min \quad \sum_{m \in M} c_m x_m + \sum_{l \in L} c_l y_l \tag{1}$$

$$\text{s.t.} \quad \sum_{l \in L} a_{il} y_l = 1 \qquad \forall i \in N^C \tag{2}$$

$$\sum_{l \in L_m} \sum_{i \in N^C} a_{il} q_i y_l \leq Q^{(1)} x_m \qquad \forall m \in M \tag{3}$$

$$x_m \in \{0, 1\} \qquad \forall m \in M \tag{4}$$

$$y_l \in \{0, 1\} \qquad \forall l \in L \tag{5}$$

The objective function (1) minimizes the total travel time. Constraints (2) ensure that each customer is visited exactly once. Constraints (3) ensure that the total load on second-echelon routes supplied by a given first-echelon vehicle does not exceed its capacity. Constraints (4), and (5) define the domain of the decision variables.

Time-Dependent Travel Times

We assume that first-echelon vehicles travel at a constant free-flow speed throughout the planning horizon. Let $\rho_m^{(1)} = (v_0, v_1, v_2, ..., v_k)$ be a sequence of nodes representing first-echelon route m and

t_i^m the ready time at node i on route m, indicating the earliest possible departure time from i. The total travel time is given by $c_m = t_{v_k}^m$, and the ready time at each node is computed as

$$t_i^m = \begin{cases} 0, & \text{if } i = 0 \\ t_{i-1}^m + \dfrac{d_{(i-1,i)}}{\alpha} + \theta_s, & \text{if } i > 0 \end{cases} \tag{6}$$

where $d_{i-1,i}$ is the distance between $i-1$ and i, and α is the first-echelon travel speed parameter. To account for traffic conditions in urban areas, where travel time can vary significantly during peak and off-peak hours, we adopt the time-dependent travel time function proposed by Ichoua et al. (2003). The planning horizon is divided into multiple time intervals. Travel speed on each arc is defined by a piecewise constant function, from which a piecewise linear travel time function is obtained. This function satisfies the FIFO property. In the resulting time-dependent network, the travel time on arc (i, j) depends on the departure time from i. In practice, arcs with the same speed profiles can by grouped to avoid storing redundant information.
Let $\rho_l^{(2)} = (v_0, v_1, v_2, ..., v_k)$ be a sequence of nodes representing second-echelon route l, e the earliest possible departure time from a satellite supplied by a given route, and $\delta_i^l(e)$ a function that returns ready time at node i on route l given the departure time e. The total travel time is then $c_l = \delta_{v_k}^l(e)$, and the ready time at each node is calculated as

$$\delta_i^l(e) = \begin{cases} e, & \text{if } i = 0 \\ \max\{\underline{w}_i + \theta_i, \delta_{i-1}^l(e) + \tau_{i-1,i}(\delta_{i-1}^l(e)) + \theta_i\}, & \text{if } i > 0 \end{cases} \tag{7}$$

where τ is the travel time function propose by Ichoua et al. (2003) that returns the travel time on $(i-1, i)$ given the departure time at $i-1$.

Branch-and-Price

In practice, the number of depots and satellites is usually limited and all feasible first-echelon routes can be enumerated a priori. To ensure sufficient coverage of all satellites, we adopt the approach proposed by Mhamedi et al. (2022) and create $\left\lceil \frac{\sum_{i \in N^C} q_i}{Q^{(1)}} \right\rceil$ copies of all routes that visit only one satellite. On the contrary, the number of feasible second-echelon routes is typically very large, so these routes are generated iteratively using the branch-and-price framework.
Branch-and-price is an exact solution method that combines branch-and-bound with column generation. The integer program (1) - (5) is referred to as the master problem (MP). We start with a restricted master problem (RMP) containing only an initial subset of second-echelon routes $L' \subseteq L$ and solve its LP relaxation using column generation. To obtain a feasible initial subset of second-echelon routes, in this baseline model, we generate back-and-forth second-echelon trips, routes that visit exactly one customer, supplied by back-and-forth first-echelon trips, so that all customers are covered and minimum number of first-echelon vehicles is used.
Column generation alternates between solving the LP relaxation of the RMP and one or more pricing problems. Dual variables obtained from the linear program are used in the pricing problem to generate new columns with negative reduced cost. If such a column exists, it is added to the RMP and the LP relaxation is re-solved. This process repeats until no improving column exists, proving LP optimality. The solution of the LP relaxation provides a lower bound for the original integer program. If the solution is fractional, the problem is branched. Column generation is applied at each node of the branch-and-bound tree.

Pricing Problem

Let $\sigma_i, i \in N^C$, and $\pi_m, m \in M$ be the dual variables of constraints (2) and (3), respectively. The reduced cost $\bar{c}_l$ of decision variable y_l is calculated as

$$\bar{c}_l = c_l - \sum_{i \in N^C} a_{il}(\sigma_i + q_i \pi_m) \tag{8}$$

We define the pricing problem as a variant of the elementary shortest path problem with resource constraint (ESPPRC) and use a forward labeling algorithm to solve it. For the sake of conciseness, we provide only a high-level overview and refer the reader to Feillet et al. (2004) for a detailed description of the ESPPRC and the labeling algorithm, and to Mhamedi et al. (2022) for the full description of their adaptation to the two-echelon variant.

ESPPRC requires specified start and terminal nodes. In our case, the start node is a given satellite s, and the terminal node is a copy of the satellite, denoted s', which has no outgoing arcs. The algorithm starts at the start node and iteratively extends so called labels - representing partial paths - while tracking resource consumption and enforcing feasibility. Labels are propagated through the network until the terminal node s' is reached. Each label E contains the following information

$m(E)$: the first-echelon route supplying the partial path

$r(E)$: the sequence of nodes visited along the partial path

$v(E)$: the last visited node

$\bar{c}(E)$: the cumulated reduced cost

$\delta(E)$: ready time at $v(E)$

$c(E)$: the cumulated travel time (route cost)

$q(E)$: total quantity delivered along the partial path

$S(E)$: a subset of successor nodes of $v(E)$ denotes $\gamma(v(E))$, indicating nodes that can be reached from $v(E)$ without violating path feasibility

In each column generation iteration, one pricing problem is solved for each satellite $s \in N^S$. For every first-echelon route $m \in M$ that visits satellite s, a label is created with the following initial values: $m(E) = m$, $r(E) = (s)$, $v(E) = s$, $\bar{c}(E) = 0$, $\delta(E) = t_s^m$, $c(E) = 0$, $q(E) = 0$, $S(E) = \psi(E)$, where $\psi(E)$ is a function that returns a subset of successor nodes $\gamma(v(E))$ reachable from $v(E)$ without violating path feasibility. Note that ready time $\delta(E)$ is obtained from the first-echelon route m supplying satellite s.
New label E' can be created by extending label E to node $j \in S(E)$ as follows

$$m(E') = m(E) \tag{9}$$
$$r(E') = r(E) \cup j \tag{10}$$
$$v(E') = j \tag{11}$$
$$\bar{c}(E') = \bar{c}(E) + \max\{\underline{w}_j - \delta(E), \tau_{v(E),j}(\delta(E))\} + \theta_j - \sigma_j - q_j \pi_{m(E)} \tag{12}$$
$$\delta(E') = \max\{\underline{w}_j, \delta(E) + \tau_{v(E),j}(\delta(E))\} + \theta_j \tag{13}$$
$$c(E') = \delta(E') - t_s^m \tag{14}$$
$$q(E') = q(E) + q_j \tag{15}$$
$$S(E') = \psi(E') \tag{16}$$

Extension of a label E to node j is feasible if

$$\{j\} \cap \gamma(v(E)) \setminus \{i : i \in r(E)\} \neq \emptyset \wedge q(E) + q_j \leq Q^{(2)} \wedge \delta(E) + \tau_{v(E),j}(\delta(E)) \ + \theta_j \leq \bar{w}_j \tag{17}$$

Equation (17) states that a label E to node j is feasible only if j is a successor of $v(E)$, j has not yet been visited on the partial path $r(E)$, the time window of customer j is respected, and the vehicle capacity $Q^{(2)}$ is not exceeded.
The algorithm ends when all labels are processed and all possible extensions obtained. To reduce the number of labels, a dominance test is applied, and only non-dominated labels are stored. We state that a label E_2 is dominated by a label E_1 if

$$v(E_1) = v(E_2) \tag{18}$$
$$\bar{c}(E_1) \leq \bar{c}(E_2) + \min\{0, (Q^{(2)} - q(E_2))(\pi_{m(E1)} - \pi_{m(E2)})\} \tag{19}$$
$$c(E_1) \leq c(E_2) \tag{20}$$
$$q(E_1) \leq q(E_2) \tag{21}$$
$$S(E_2) \subseteq S(E_1) \tag{22}$$

Expression (18) ensures that only labels ending at the same node are compared. Conditions (19) - (21) state E_1 can dominate E_2, only if its reduced cost, resource consumption, and travel time are lower than those of E_2. Condition (22) ensures that E_1 can dominate E_2 only if all nodes reachable from E_2 are also reachable from E_1.

Branching

To find the optimal solution to the original integer program (1) - (5), we employ two branching rules compatible with the pricing problem and use the depth-first search. We branch on the first-echelon routes first. To forbid a first-echelon route, the upper bound of its binary decision variable in the RMP is fixed to 0, and the route is excluded from the pricing problem. Conversely, to enforce the use of a first-echelon route, the lower bound of its binary decision variable in the RMP is fixed to 1.
Once the solution contains no fractional values for the decision variables representing the first-echelon routes, we branch on customer arcs. To forbid an arc, it is removed from the pricing problem, and all columns containing the arc are removed from the RMP. To ensure that an arc (i, j) is included in the solution, all arcs $(i, a), a \neq j$ and $(b, j), b \neq i$ are removed from the pricing problem, and all columns containing such arcs are removed from the RMP.

3 Preliminary results and discussion

To validate the baseline algorithm and evaluate its performance, an initial computational study was conducted. We generated a set of test instances for the two-echelon VRP with time windows and time-dependent travel times. Customer locations were generated randomly, with travel distances computed as Euclidean distances. Both distance and time are expressed in arbitrary units. Key characteristics of the instances are presented in Table 1. The study was conducted on a computer equipped with Intel Core i7-1365U and 16GB RAM. The algorithm is implemented in Python 3.13. Gurobi 12.0.2 is used to solve LP relaxation of the RMP.

Baseline model performance

The baseline algorithm solved all test instances to optimality. The preliminary results presented in Table 1 indicate the viability of the proposed solution method for the problem at hand. While the results are promising, the study highlights limitations of the baseline model and indicates directions for improvement.

Table 1: Preliminary computational study results

$\lvert N^D\rvert$	$\lvert N^S\rvert$	$\lvert N^C\rvert$	TW	$Q^{(1)}$	$Q^{(2)}$	H	h	time (s)	opt.
1	2	10	20	120	35	180	6	3.8	✓
1	2	15	20	120	35	180	6	95.7	✓
1	2	20	15	120	30	180	6	145.7	✓
1	2	20	20	120	35	180	10	194.7	✓
1	2	30	15	150	20	180	10	1598.9	✓
1	3	15	20	120	35	160	10	17.6	✓
1	3	20	15	120	25	180	10	42.0	✓
1	3	20	15	120	30	180	10	450.4	✓
1	4	15	20	150	35	220	10	247.1	✓
1	4	20	15	150	25	220	10	296.3	✓

$\lvert N^D\rvert$: number of depots; $\lvert N^S\rvert$: number of satellites; $\lvert N^C\rvert$: number of customers; TW: time window width; $Q^{(1)}$: first-echelon vehicle capacity; $Q^{(2)}$: first-echelon vehicle capacity; H: planning horizon; h: number of time intervals; opt.: optimal solution

As instance size grows, the number of feasible second-echelon routes increase substantially. This growth, however, is influenced by multiple factors. For example, instances with tighter time windows, lower second-echelon vehicle capacity, or sparser customer locations yield fewer feasible routes per satellite, which leads to faster pricing. To improve the efficiency, bi-directional labeling and enhanced dominance rules for the time-dependent ESPPRC can be explored. Moreover, the problem contains first-echelon routes with identical ready times at a given satellite, which leads to the redundant generation of labels representing identical partial paths.
The current branching scheme provides a functional starting point but does not scale well to larger instances. The preliminary results suggest that stronger branching decisions or additional cutting

planes are needed.
The simple construction method used to generate the initial set of feasible second-echelon routes produces a poor starting solution for the restricted master problem. A weak initial solution leads to more column generation iterations at the root node and may result in deeper branch-and-bound trees. Replacing the current heuristic with a more informed construction method is expected to improve the efficiency.

Impact of time-dependent travel times

To illustrate the importance of accounting for time-dependent travel times, we solved one test instance under two settings. In the first setting, time-dependent speed profiles were used, while in the second, speeds were replaced with the average travel speed observed across the network and planning horizon. The results, reported in Table 2, show that although the objective function values are similar, the resulting solutions differ substantially. Solutions optimized using averaged speeds may be suboptimal, or even infeasible, when evaluated under actual time-dependent conditions, highlighting the practical significance of incorporating travel time variability in urban logistics with fluctuating congestion patterns.

Table 2: Impact of time-dependent travel times

TD	time (s)	obj.	solution	
Yes	95.7	364.42	FE	0 - 1 - 0
				0 - 2 - 0
			SE	1 - 9 - 6 - 4 - 10 - 11 - 1
				1 - 7 - 13 - 14 - 3 - 5 - 12 - 1
				2 - 17 - 8 - 15 - 16 - 2
No	202.3	361.56	FE	0 - 1 - 0
			SE	1 - 17 - 14 - 15 - 12 - 1
				1 - 6 - 13 - 7 - 8 - 3 - 11 - 5 - 1
				1 - 9 - 4 - 10 - 16 - 1

TD: time-dependent travel times; obj.: objective value; FE: first-echelon; SE: second-echelon

4 Conclusions

This study highlights the importance of explicitly accounting for traffic conditions in vehicle routing models. A baseline branch-and-price algorithm for the 2E-TDVRP-TW is presented and shown to produce promising results, establishing a solid foundation for further development. Future work will focus on three targeted enhancements: improving the labeling algorithm to accelerate column generation, improving the construction heuristic to generate a higher-quality initial subset of second-echelon routes, and developing a more effective branching strategy. Together, these improvements are expected to substantially enhance the algorithm's scalability.

Acknowledgements

This work has been funded by the European Union under the Horizon Europe grant 101139678 (metaCCAZE). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them.

References

Dellaert, N., Dashty Saridarq, F., Van Woensel, T., & Crainic, T. G. (2019). Branch-and-

price–based algorithms for the two-echelon vehicle routing problem with time windows. *Transportation Science*, *53*(2), 463–479.

Feillet, D., Dejax, P., Gendreau, M., & Gueguen, C. (2004). An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems. *Networks*, *44*(3), 216-229.

Gutierrez, A., Labadie, N., & Prins, C. (2024). A two-echelon vehicle routing problem with time-dependent travel times in the city logistics context. *EURO Journal on Transportation and Logistics*, *13*, 100133.

Ichoua, S., Gendreau, M., & Potvin, J.-Y. (2003). Vehicle dispatching with time-dependent travel times. *European Journal of Operational Research*, *144*(2), 379-396.

Mhamedi, T., Andersson, H., Cherkesly, M., & Desaulniers, G. (2022). A branch-price-and-cut algorithm for the two-echelon vehicle routing problem with time windows. *Transportation Science*, *56*(1), 245–264.

Reyes, F. O. G., Gendreau, M., & Potvin, J.-Y. (2024). A metaheuristic for a time-dependent vehicle routing problem with time windows, two vehicle fleets and synchronization on a road network. *EURO Journal on Transportation and Logistics*, *13*, 100143.

Sluijk, N., Florio, A. M., Kinable, J., Dellaert, N., & Van Woensel, T. (2023). Two-echelon vehicle routing problems: A literature review. *European Journal of Operational Research*, *304*(3), 865-886.

Soysal, M., Bloemhof-Ruwaard, J. M., & Bektaş, T. (2015). The time-dependent two-echelon capacitated vehicle routing problem with environmental considerations. *International Journal of Production Economics*, *164*, 366-378.