

# A matheuristic based on large neighborhood search for the vehicle routing problem with cross-docking

Philippe Grangier<sup>a,b,\*</sup>, Michel Gendreau<sup>b</sup>, Fabien Lehuédé<sup>a</sup>, Louis-Martin Rousseau<sup>b</sup>

<sup>a</sup>*L'UNAM, Ecole des Mines de Nantes, IRCCyN UMR CNRS 6597, 4 Rue Alfred Kastler, 44307 Nantes Cedex 3, France*

<sup>b</sup>*Department of Mathematics and Industrial Engineering and CIRRELT, Ecole Polytechnique de Montréal and CIRRELT, C.P 6079, Succursale Centre-ville, Montreal, QC, Canada H3C 3A7*

---

## Abstract

The vehicle routing problem with cross-docking (VRPCD) consists in defining a set of routes that satisfy transportation requests between a set of pickup points and a set of delivery points. The vehicles bring goods from pickup locations to a cross-docking platform, where the items may be consolidated for efficient delivery. In this paper we propose a new solution methodology for this problem. It is based on large neighborhood search and periodically solving a set partitioning and matching problem with third-party solvers. Our method improves the best known solution in 19 of 35 instances from the literature.

*Keywords:* Routing, Cross-docking, Transfers, Synchronization, Matheuristic.

---

## 1. Introduction

Cross-docking is a distribution strategy in which goods are brought from suppliers to an intermediate transshipment point, the so-called cross-dock, where they may be transferred to another vehicle for delivery. The transfers are based on consolidation opportunities. There is little or no storage capacity at the cross-dock, so the inventory holding costs are low, and the consolidation process reduces the distribution costs. Cross-docking has been successfully applied to several sectors. The classic example is Walmart: cross-docking is said to have been the key to the growth of the retailer in the 1980s [38].

Problems related to cross-docking include location, assignment of trucks to doors, inner flow optimization, and routing. In particular, the *vehicle routing problem with cross-docking* (VRPCD) [45] consists in designing routes to pick up and deliver a set of goods at minimal cost using a single cross-dock. The trucks bring the goods from the pickup locations to the cross-dock where they can offload some items and load others, and they then make the delivery trips. The exchange of goods at the cross-dock is a consolidation process that aims to minimize the total delivery cost. The VRPCD can be seen as a *pickup and delivery problem with transfers* with a single compulsory transfer point.

In this paper, we propose a matheuristic that relies on large neighborhood search (LNS) to create a pool of routes. These routes are then used in a set partitioning and matching (SPM) problem. This problem is

---

\*Corresponding author.

*Email addresses:* philippe.grangier@cirrelt.ca (Philippe Grangier), michel.gendreau@cirrelt.ca (Michel Gendreau), fabien.lehuede@mines-nantes.fr (Fabien Lehuédé), louis-martin.rousseau@polymtl.ca (Louis-Martin Rousseau)

solved using branch-and-check [41], a hybrid method that relies on both a mixed integer programming (MIP) solver and a constraint programming (CP) solver. We present numerical results showing that our method improves many best known solutions while competing with existing method within similar runtime.

The remainder of this paper is organized as follows. A literature review is presented in Section 2, and the problem is defined in Section 3. Section 4 discusses the solution methodology, and Section 5 presents the computational results.

## 2. Literature review

Our literature review will focus on routing. We refer the reader to [1, 5, 42] for a general overview of cross-docks and the related problems. The cross-docking literature is fairly recent, and only a few papers focus on the routing aspect. Lee et al. [18] solved a VRPCD variant in which all the vehicles have to arrive at the cross-dock simultaneously. They proposed an exact formulation and developed a tabu-search heuristic to solve instances with 10, 30, and 50 nodes. This heuristic was improved by Liao et al. [19]. Wen et al. [45] extended the problem by adding time windows on the nodes and removing the requirement for simultaneous arrival at the cross-dock, instead imposing precedence constraints based on the consolidation decisions. They presented a MIP formulation and proposed a tabu search embedded in an adaptive memory procedure. They solved real-world instances with up to 200 requests, and they compared their results to a lower bound obtained by solving two VRPTW problems. These results were improved by Tarantilis [40] using a multi-restart tabu search. Morais et al. [23] developed a method based on an iterative local search and a set partitioning problem (SPP); they introduced new instances with up to 500 customers. Very recently, Nikolopoulou et al. [24] presented a tabu search for the VRPCD that improved several of the best known results. Petersen and Ropke [27] worked with a Danish flower-distribution company on a variant of the VRPCD with time windows, optional cross-dock return, and multiple trips per day; they call this the *vehicle routing problem with cross-docking opportunity*. They created a parallel adaptive LNS (ALNS) to solve instances with between 585 and 982 requests. Santos et al. [33, 35] proposed two branch-and-price approaches for a VRPCD variant in which there is a cost to transfer items at the cross-dock but there are no temporal constraints. They extended their approach [34] to a problem where returns to the cross-dock for consolidation are optional. They refer to this problem as the *pickup and delivery problem with cross-docking* and show that it can reduce the routing cost by between 3.1% and 7.7% on their instances, which are restrictions of those of [45]. Dondo and Cerdá [10] considered a variant of the VRPCD in which they modeled each door at the cross-dock individually (e.g., handling speeds and travel times to other doors), and the number of trucks was greater than the number of doors. They proposed an algorithm based on an MILP formulation and a sweep heuristic, and they solved instances with up to 70 requests. Enderer [12] studied the *dock-door assignment and vehicle routing problem*; it involves only the assignment of trucks to doors and the routing of the deliveries. He proposed and compared several exact and heuristic methods. In a recent survey of synchronization in cross-docking networks, Buijs et al. [6] classified the VRPCD as a network scheduling problem with synchronization.

A related problem is the *pickup and delivery problem with transfers* (PDPT). Heuristics for this problem include ALNS [20] and GRASP+ALNS [30]; an exact branch-and-cut method [7] has also been proposed. Another related problem is the *two-echelon vehicle routing problem* introduced by [26], for which collaboration between the two echelons is at the heart of the delivery process, is also related to the VRPCD. In particular, the *two-echelon multiple-trip vehicle routing problem with satellite synchronization* [13] deals with the temporal aspects. For a recent surveys on two-echelon routing problems we refer the reader to [8]. Lastly, cross-docking operations correspond to both operation synchronization and resource synchronization as described in [11].

Matheuristic approaches often use one or several (meta)heuristics to generate a pool of routes and then solve an SPP (either during the search or in postprocessing). Examples of the metaheuristics used include local search for the VRP [31], tabu search for the split delivery VRP [3], ALNS for the technician routing problem [28], GRASP and local search procedures for the truck and trailer routing routing problem [44] and the VRP with stochastic demands [22], and iterated local search for seven VRP variants [39]. This last method was applied to the VRPCD in [23]. For more details see the surveys by Doerner and Schmid [9] and Archetti and Speranza [2]; the metaheuristic+SPP technique is classified as a set-covering/SPP approach in the former survey and as a restricted master heuristic in the latter. The LNS+SPM proposed in this paper can also be classified into these categories because of the SPM component. However, the SPM also involves matching and scheduling decisions. It is solved using a *branch-and-check* approach [41], which is a hybrid technique integrating MIP and CP; see Section 4.2. So far, branch-and-check has been used mainly for solving planning and scheduling or resource allocation problems. To the best of our knowledge, this is the second time branch-and-check is used for solving a vehicle routing problem, with Lam and Van Hentenryck [16] proposing a branch-and-price-and-check for a VRP with resource constraints. To our knowledge, this is the first time a metaheuristic has been hybridized with a branch-and-check component to solve a VRP.

### 3. Problem formulation

In this section we present the VRPCD, focusing on the scheduling constraints at the cross-dock.

#### 3.1. Problem statement

In the VRPCD, we consider a cross-dock  $c$ , a set of items  $R$ , and a homogeneous fleet of vehicles  $V$ , each of capacity  $Q$  and based at  $o$ . Each item  $r \in R$  must be collected at its pickup location  $p_r$  within the time window  $[e_{p_r}, l_{p_r}]$  and delivered to its delivery location  $d_r$  within the time window  $[e_{d_r}, l_{d_r}]$ . In the case of early arrival, a vehicle is allowed to wait, but late arrivals are forbidden. We denote by  $P$  the set of pickup locations and by  $D$  the set of delivery locations.

Each vehicle starts at  $o$ , goes to several pickup locations, and then arrives at the cross-dock where it unloads/reloads some items. It then visits several delivery locations and eventually returns to  $o$ . Note that a vehicle must visit the cross-dock even if it does not unload or reload there. The sequence of operations

at the cross-dock is described in Section 3.2. The *pickup leg* is the sequence of operations performed by a vehicle between its departure from  $o$  to perform pickups and its arrival at the cross-dock. The *delivery leg* is the sequence of operations performed by a vehicle between its departure from the cross-dock to perform deliveries and its return to  $o$  at the end of the day.

The VRPCD is defined on a directed graph  $G = (V, A)$ , with  $V = \{o\} \cup P \cup \{c\} \cup D$  and  $A = \{(o, p) | p \in P\} \cup P \times P \cup \{(p, c) | p \in P\} \cup \{(c, d) | d \in D\} \cup D \times D \cup \{(d, o) | d \in D\} \cup \{(o, c), (c, o)\}$ . With each arc  $(i, j) \in A$  is associated a travel time  $t_{i,j}$  and a travel cost  $c_{i,j}$ . Solving the VRPCD involves finding  $|V|$  routes, and a schedule for each route, such that the capacity and time-related constraints are satisfied at a minimal routing cost. An arc-based mathematical formulation can be found in [45].

### 3.2. Cross-dock operations

Following [45], if a vehicle  $k$  has to unload a set of items  $R_k^-$  and reload a set  $R_k^+$  at the cross-dock, the time spent at the cross-dock can be divided into four periods, as shown in Figure 1:

- Preparation for unloading. The duration  $\delta_u$  of this period is fixed.
- Unloading. The duration of this period depends on the quantity of items to unload. For vehicle  $k$  the duration is  $(\sum_{i \in R_k^-} q_i) / s_u$ , where  $s_u$  is the unloading speed in quantity per time unit. All unloaded items become available for reloading at the end of this period.
- Preparation for reloading. The duration  $\delta_r$  of this period is fixed.
- Reloading. The duration of this period depends on the quantity of items to reload. For vehicle  $k$  the duration is  $(\sum_{i \in R_k^+} q_i) / s_r$ , where  $s_r$  is the reloading speed in quantity per time unit. All the items for loading must have been unloaded before the beginning of the reloading operation (preemption is not allowed).




Figure 1: Time chart for vehicle unloading and reloading at the cross-dock. The vehicle must wait to be reloaded because some items are not available when it is ready for reloading.

Items that are not transferred at the cross-dock remain in the vehicle. Thus, if a vehicle does not unload or reload it need not spend any time at the cross-dock and can leave immediately. We do not limit the number of available docks.

#### 4. Algorithm

In this section we describe our algorithm for the VRPCD. The main component is LNS, which is enhanced by the occasional solution of an SPM. When the LNS finds a new solution, the legs in this solution are added to a pool of legs that acts as a memory. The SPM is based on an SPP where the set to partition is  $P \cup D$ , and the candidate partitions are the legs in the pool. Thus, the SPM assists the LNS by selecting good legs that have been previously discovered. We call this LNS+SPM; algorithm 1 presents an outline. Lines 2–15 describe the LNS; see Section 4.1. The SPM is used during the search (line 18); see Section 4.2. The set of legs used in each SPM is managed during the search (line 20).

**Result:** The best found solution  $s^*$

```

1 Pool of legs  $\mathcal{L} := \emptyset$ 
2 Generate an initial solution  $s$ 
3  $s^* := s$ 
4 while stop criterion not met do
5      $s' := s$ 
6     Destroy quantity: select a number  $\Phi$  of requests to remove from  $s'$ 
7     Operator selection: select a destruction method  $M^-$  and a repair method  $M^+$ 
8     Destruction: Apply  $M^-$  to remove  $\Phi$  requests from  $s'$  and place them in the request bank of  $s'$ 
9     Repair: Apply  $M^+$  to reinsert the requests into the request bank in  $s'$ 
10    if acceptance criterion is met then
11         $s := s'$ 
12    end
13    if cost of  $s'$  is better than cost of  $s^*$  then
14         $s^* := s'$ 
15    end
16    Add legs of  $s'$  to  $\mathcal{L}$ 
17    if set partitioning and matching condition is met then
18        Perform set partitioning and matching with the legs in  $\mathcal{L}$ 
19        Update  $s^*$  and  $s$  if a new best solution has been found
20        Perform pool management
21    end
22 end
23 return  $s^*$ 

```

**Algorithm 1:** LNS+SPM

#### 4.1. Large Neighborhood Search

LNS has been widely used since it was introduced by Shaw [37]. It iteratively destroys (removes requests from) and repairs (reinserts requests into) the current solution using heuristics. Its extension, ALNS [29, 32], selects the destruction and repair methods based on their past success. We select these methods randomly (our method is thus LNS-based) because preliminary experiments suggested that the adaptive layer has an extremely limited impact on the quality of the solutions. Parragh and Schmid [25] likewise choose to use LNS and the SPP for their dial-a-ride problem. We also do not consider noise in the cost function. We now present our destruction and repair methods as well as strategies to reduce the computational time.

##### 4.1.1. Destruction methods

When partially destroying a solution, we select a destruction method  $M^-$  and a number  $\Phi$  of requests to remove. Unless stated otherwise, this method is reused until  $\Phi$  is reached. We introduce the transfer removals, and the other removal techniques below are inspired by [29].

*Random removal:*. We randomly remove a request.

*Worst removal:*. We remove a request with a high removal gain. This is defined as the difference in the cost of the solution with and without the request. We then sort the requests in nonincreasing order of removal gain and place them in a list  $N$ . We select the request to remove in a randomized fashion as in [32]: given a parameter  $p$ , we draw a random number  $y$  between 0 and 1. We then remove the request in position  $y^p \times |N|$ .

*Historical node-pair removal:*. Each arc  $(u, v) \in G$  is associated with the cost of the cheapest solution in which it appears (initially this cost is set to infinity). We then remove the request that is served using the arcs with the highest associated costs. A randomized selection, similar to that for the worst removal, is performed.

*Related removals:*. These aim to remove related requests. Let the relatedness of requests  $i$  and  $j$  be  $R(i, j)$ . We use two relatedness measures: distance and time. The distance measure between two requests is the sum of the distance between their pickup points and the distance between their delivery points. The time measure is the sum of the absolute difference between their start times at their pickup points and the absolute difference between their start times at their delivery points. In both cases a small  $R(i, j)$  indicates a high relatedness. A randomized selection, similar to that for the worst removal (but with a nondecreasing ordering), is performed.

*Transfer removal:*. For each pair of routes  $(v_i, v_j)$ ,  $v_i \neq v_j$ , we compute the number of requests transferred from  $v_i$  to  $v_j$ . We then iteratively apply a roulette-wheel selection on the pairs of routes (the score of a pair being the number of requests transferred), and we remove the requests that are transferred between the routes in the selected pair. If there are fewer transferred requests than the target number  $\Phi$  to remove, we remove them all and switch to random removal for the rest.

#### 4.1.2. Repair methods

In LNS, the unplanned requests are stored in a *request bank*. We now explain how we insert these requests into a partial solution.

*Best insertion.*: From the requests  $r$  in the request bank, we choose the one with the cheapest insertion cost considering all possible insertions of  $p_r$  into pickup legs and  $d_r$  into delivery legs.

*Regret insertion.*: For each request  $r$  in the request bank and for each pair of vehicles (pickup vehicle, delivery vehicle), we compute the cost of the cheapest feasible insertion (if any). Note that the pickup and delivery vehicles may be the same (in the case of insertion without transfer). Then, with these insertion options, we compute the  $k$ -regret value of  $r$ ,  $c_r^k = \sum_{i=1}^k f_i - f_1$ , where  $f_1$  is the cost of the cheapest insertion,  $f_2$  is the cost of the second-cheapest insertion, and so on, and  $k$  is a parameter. We insert the request with the highest regret value.

In algorithm 1 we use 2-regret to generate the initial solution. We use best-insertion, 2-regret, 3-regret, and 4-regret as the repair methods.

#### 4.1.3. Reducing the computational time

In LNS, the repair methods take most of the computational time: more than 96% in [13]. We use two strategies to reduce this time: reducing the size of the neighborhoods and efficient time checking.

*Size of the neighborhood.*: Because each request can be transferred, the number of candidate insertions for a request is  $\Theta(|V|^2)$ , which is much larger than for the traditional VRP, where it is only  $\Theta(|V|)$ . For large instances, it takes considerable time to explore the entire reinsertion neighborhood. This does not necessarily lead to a better solution, because not all the transfer opportunities are worthwhile. Thus, for each request we evaluate all the insertions without transfers, and we consider transfer opportunities for only a subset of  $g$  vehicles. For each request  $r$  in the request bank we sort the vehicles for pickup (resp. delivery) in nondecreasing order of the cheapest insertion of  $r$  in the pickup (resp. delivery) leg. We consider insertion with transfers between the first  $g$  vehicles according to this order. A discussion of the choice of  $g$  is presented in Section 5.2.2. We use this restricted-neighborhood exploration in both repair methods defined in Section 4.1.2.

*Efficient time checking.*: If an insertion appears promising (i.e., worth considering with respect to its cost), we need to check that it is feasible. Checking the capacity constraints can easily be done in constant time, but the time constraints are more complex: a straightforward implementation has a linear complexity. Savelsbergh [36] introduced *forward time slacks* that allow us to check in constant time if an insertion is feasible in the VRP with time windows. Masson et al. [21] extended this check to the PDPT. This method has been used several times [13, 14, 20] and has proven successful in significantly reducing the computational time, for example in [13]; see [21] for a detailed description of its implementation. Since the VRPCD can be seen as a PDPT, we reuse it here.

#### 4.2. Set partitioning and matching procedure

In LNS, a solution is rejected solely based on its cost with respect to the best solution so far. A solution that is rejected because it contains bad legs may also contain good legs, which will also be removed. It may also happen that the matching of legs to form routes could be improved (e.g., giving more time flexibility and thus making further improvements easier). We can address these issues by storing the legs found by LNS and using them in an SPP.

The SPM is defined as follows. Given a set of legs  $L$ , select a subset of legs  $\tilde{L}$  such that (1) each request is picked up (resp. delivered) by exactly one leg in  $\tilde{L}$ , and (2) the legs in  $\tilde{L}$  can be matched to form routes that respect the time constraints. Each leg  $l \in L$  has an associated routing cost  $c_l$ , and the objective is to minimize the sum of the costs of the selected legs. We now present the branch-and-check method that we use to solve the SPM. We discuss the master problem (set partitioning), the subproblem (matching and scheduling), and the management of the pool of legs.

Our approach is similar to that of Morais et al. [23], who used the SPP as the intensification phase of their SP-ILS. However, there are two main differences: (1) their pool of legs is much smaller (consisting of the legs that appear in their ten best solutions), and (2) when they can not find a feasible matching (they do not discuss their matching procedure), they perform a repair phase by moving requests from one route to another.

##### 4.2.1. Branch-and-check

We present branch-and-check [41] using the following optimization problem:

$$M1 : \min c^\top x \tag{1}$$

$$Ax \leq b \tag{2}$$

$$H(x, y) \tag{3}$$

$$x \in \{0, 1\}^n \tag{4}$$

$$y \in \mathbb{R}^m \tag{5}$$

Assume that  $H(x, y)$  represents a set of constraints that have a limited impact on the LP relaxation and/or are difficult to efficiently model in a MIP, but that can be handled relatively easily by a CP solver. (1), (2), and (4) form a relaxation (M2) of (M1) that can be solved using branch-and-bound. The general principle of branch-and-check is the following. To solve (M1), we carry out a branch-and-bound on (M2). Whenever an integral solution of (M2) is found, we call a CP solver to check constraints (3). If they are satisfied, we update the best solution so far for (M1). Otherwise, we reject this solution. In both cases the branch and bound process continues.

In the SPM, we solve a classical SPP where the set to partition is  $P \cup D$ , and the candidate partitions are the legs in the pool  $\mathcal{L}$ . A solution to the SPP is a solution to the VRPCD iff we can match the legs to form



a set of routes that respect the time constraints. We determine whether or not this is possible by solving a dedicated matching and scheduling subproblem. Our SPP is the relaxation (M2), and the subproblem checks the constraints  $H(x, y)$ .

#### 4.2.2. Set partitioning

Consider a set of legs  $L = L_p \cup L_d$ , where  $L_p$  is a set of pickup legs and  $L_d$  is a set of delivery legs. For each request  $r \in R$ , we define a binary constant  $\lambda_{r,l}$  that indicates whether or not this request is served by this leg. For each leg  $l \in L$  we use a binary variable  $x_l$  to determine whether or not it is selected. The SPP is then

$$\min \sum_{l \in L} c_l x_l \tag{6}$$

$$\sum_{l \in L_p} \lambda_{r,l} x_l = 1 \quad \forall r \in R \tag{7}$$

$$\sum_{l \in L_d} \lambda_{r,l} x_l = 1 \quad \forall r \in R \tag{8}$$

$$x_l \in \{0, 1\} \quad \forall l \in L \tag{9}$$

The objective (6) is to minimize the cost of the selected legs, and constraints (7) (resp. (8)) ensure that each pickup point (resp. delivery point) is covered by exactly one leg.

Most MIP solvers use an *incumbent callback* to call an auxiliary subproblem after an integral solution has been found. To save time, we warm-start the MIP solver with the best solution so far ( $s^*$  in algorithm 1).

If the matching and scheduling subproblem is not feasible, one could add some lazy constraints (this is common in branch-and-check). For example, if, for a given pickup leg  $l$  and a given delivery leg  $l'$  with some requests in common, there is not enough time to perform the associated operations at the cross-dock whether or not they are packed together, we could add  $x_l + x_{l'} \leq 1$  as a lazy constraint in the SPP. We experimented with this technique, but it performed poorly. We identify two reasons for this: (1) the vast majority of the incumbent callbacks are successful, and (2) in CPLEX 12.6.1 (which we use) adding lazy-constraint callbacks disables dynamic search, which seems to reduce CPLEX's computational time on the SPP.

#### 4.2.3. Matching and scheduling subproblem

A solution of the SPP, with a set of pickup legs denoted  $\tilde{L}_p$  and a set of delivery legs denoted  $\tilde{L}_d$ , is a solution to the VRPCD iff there exists a matching of the pickup legs and delivery legs into routes that respect the time constraints. We can compute the earliest feasible arrival time,  $a_l$ , of each pickup leg  $l \in \tilde{L}_p$  at the cross-dock, and we can compute the latest feasible departure time,  $b_{l'}$ , of each delivery leg  $l' \in \tilde{L}_d$  from the cross-dock. Moreover, for each pickup leg  $l \in \tilde{L}_p$  we can determine the set of selected delivery legs  $T_l$  that deliver at least one item picked up by  $l$ . If we match a pickup leg  $l$  and a delivery leg  $l'$  to create a route, we can define an associated unloading task  $o_{ll'}^-$ , with a set of items  $R_{ll'}^-$  to unload, and a reloading task  $o_{ll'}^+$ , with a set of items  $R_{ll'}^+$  to reload. These tasks must be performed iff  $l$  and  $l'$  are in the same route.

The problem is modeled as a constraint satisfaction problem and represented using notation from OPL (Optimization Programming Language [43]). In particular, the model is based on the notion of interval variables [15]. An interval variable represents an unknown interval of time during which a task occurs. The interval has a start value, an end value, and a size, and the associated variable may be optional. We model alternative activities [4] using alternative constraints:

“An alternative constraint between an interval variable  $a$  and a set of interval variables  $b_1, \dots, b_n$  models an exclusive alternative between  $b_1, \dots, b_n$ . If interval  $a$  is present, then exactly one of intervals  $b_1, \dots, b_n$  is present and  $a$  starts and ends together with this specific interval. Interval  $a$  is absent if and only if all intervals in  $b_1, \dots, b_n$  are absent” [15].

We thus consider the following problem:

$$\text{Alternative}(t_l, \{o_{ll'}^-\}; \forall l' \in \tilde{L}_d) \quad \forall l \in \tilde{L}_p \quad (10)$$

$$\text{Alternative}(t_{l'}, \{o_{ll'}^+\}; \forall l \in \tilde{L}_p) \quad \forall l' \in \tilde{L}_d \quad (11)$$

$$o_{ll'}^- . \text{IsOptional} \leftarrow \text{True} \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \quad (12)$$

$$o_{ll'}^+ . \text{IsOptional} \leftarrow \text{True} \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \quad (13)$$

$$o_{ll'}^- . \text{IsPresent} \iff o_{ll'}^+ . \text{IsPresent} \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \quad (14)$$

$$t_{l'} . \text{Start} \geq t_l . \text{End} \quad \forall l \in \tilde{L}_p, l' \in T_l \quad (15)$$

$$o_{ll'}^+ . \text{Start} \geq o_{ll'}^- . \text{End} + \delta_r \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \text{ s.t. } R_{ll'}^+ \neq \emptyset \quad (16)$$

$$o_{ll'}^+ . \text{Start} \geq o_{ll'}^- . \text{End} \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \text{ s.t. } R_{ll'}^+ = \emptyset \quad (17)$$

$$o_{ll'}^- . \text{Start} \geq a_l + \delta_u \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \text{ s.t. } R_{ll'}^- \neq \emptyset \quad (18)$$

$$o_{ll'}^- . \text{Start} \geq a_l \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \text{ s.t. } R_{ll'}^- = \emptyset \quad (19)$$

$$o_{ll'}^+ . \text{End} \leq b_{l'} \quad \forall l \in \tilde{L}_p, \forall l' \in \tilde{L}_d \quad (20)$$

For each pickup leg  $l$  we create an interval variable  $t_l$  that represents the associated unloading task that takes place at the cross-dock. The alternative constraints (10) and (12) ensure that for each pickup leg  $l$  exactly one unloading task  $o_{ll'}$  is scheduled and that it is equal to  $t_l$ . The same holds for the delivery legs and the reloading operations through variables  $t_{l'}$  and constraints (11) and (13). Constraints (14) ensure that the unloading operation associated with the matching of pickup leg  $l$  and delivery leg  $l'$  in the same vehicle is present iff the corresponding reloading operation is present as well. Constraints (15) ensure that all the reloading operations that depend on a pickup leg  $l$  start no earlier than the end of the unloading task associated with  $l$ . Constraints (16) and (17) ensure that when two legs are packed together, the delay between the two tasks respects the model presented in Section 3.2. Constraints (18) and (19) ensure that the unloading of each pickup leg does not start before the earliest feasible arrival time at the cross-dock. Constraints (20) ensure that the reloading of each delivery leg is completed by its latest feasible departure time.

#### 4.2.4. Set partitioning and matching criterion and management of pool of legs

When solving the SPP, we consider only the legs in the pool of legs  $\mathcal{L}$  (from algorithm 1) that are nondominated. A pickup (resp. delivery) leg  $l_i$  is said to be dominated by a leg  $l_j$  iff  $l_i$  and  $l_j$  serve the same set of requests,  $c_j < c_i$ , and  $a_j \leq a_i$  (resp.  $b_j \geq b_i$ ). It is clear that for every solution of the SPP that contains a dominated leg, there exists a solution with a lower cost that does not contain it.

Regarding the management of the pool of legs, we make two observations. First, as the algorithm progresses, it is able to provide the MIP solver with better starting solutions, thus improving its cutoff ability. Therefore, the SPPs tend to be easier to solve at the end of the algorithm than at the beginning. Second, if the MIP has too many legs the solver may fail to improve the initial solution within the time limit. Therefore, every time the SPM is solved, if the solver is able to find an optimal solution and prove its optimality within the time limit, we retain the pool; otherwise we clear it. In practice, this policy tends to empty the pool more frequently at the beginning of the search.

The SPM procedure is called every  $k_{SPM}$  iterations; we discuss this value in Section 5.2.3.

## 5. Computational experiments

The algorithm is coded in C++. We use CPLEX and CP Optimizer from IBM ILOG Cplex Optimization Studio 12.6.1 as the MIP solver and CP solver respectively. The experiments were performed using Linux on an Intel Xeon X5675 @ 3.07 GHz. Just one core is used by our code and the third-party solvers.

### 5.1. Instances

There are two benchmarks for the VRPCD. The first, which we call the Wen set, was introduced by Wen et al. [45] and contains instances with 50 to 200 requests. It is based on data from a Danish logistics company. The second, which we call the Morais set, was introduced by Morais et al. [23] and contains instances with 200 to 500 requests. It is derived from Gehring and Homberger’s instances for the VRPTW. Neither set limits the number of vehicles.

### 5.2. Parameters

The stopping criterion is set to 20 000 iterations, which is a good compromise between solution quality and computational time. Based on Masson et al. [20] and our preliminary experiments, we choose the number  $\Phi$  of requests to remove in the repair phase of the LNS randomly from the interval  $[\min(30, 10\% \text{ of } |R|), \min(60, 20\% \text{ of } |R|)]$ .

In the following subsections, we describe the tuning experiments we conducted to set the acceptance criterion, the reduction of the transfer neighborhood, and the SPM period. For the training set, we selected the following instances: 50b, 100b, 150b, and 200b from the Wen set and R1-4-1, R1-6-1, R1-8-1, and R1-10-1 from the Morais set.

### 5.2.1. Acceptance criterion

We tested three acceptance criteria: descent (a solution is accepted iff it is better than the current best solution),  $\alpha$  threshold (a solution is accepted if it is less than  $\alpha\%$  more expensive than the current best solution), and simulated annealing (as implemented by Ropke and Pisinger [32]). Table 1 compares the performance of the following criteria: descent, 1% threshold, 3% threshold, 10% threshold, and simulated annealing, with descent as the reference. The table shows that all the criteria except the 10% threshold have similar performance at the end of the algorithm. We select descent because it has no parameters.

Accept. Crit.	Descent	1% thresh.	3% thresh.	10% thresh.	Sim. Ann.
Gap (%)	0.00	-0.01	0.04	1.55	0.00

Table 1: Average performance for five acceptance criteria; 10 runs were performed for each instance in the training set, and descent is taken as the reference.

### 5.2.2. Reduction of the transfer neighborhood

As mentioned in Section 4.1.3, for each request in the request bank we evaluate insertions with transfers between only the  $g$  most promising vehicles. We ran tests with a large value of  $g$  and we observed that in the best solutions found during the search no vehicle transferred items to more than 10 vehicles. Table 2 gives the average results over five runs on the Morais training instances for different values of  $g$ . The percentage gap is similar in each case; this confirms our intuition that only a subset of transfers are worth considering. Figure 2 shows the computational time for different values of  $g$ : for R1-10-1 (500 requests) there is a reduction of more than 50% as  $g$  decreases. Therefore, we set  $g$  to 5.

$g$	5	10	20	$\infty$
Gap (%)	-0.01	0.00	0.03	0.00

Table 2: Impact of  $g$  on solution quality for five runs for each instance in the Morais training set. Infinity is taken as the reference.

### 5.2.3. Set partitioning

Table 3 shows the influence of setting the SPM frequency  $k_{SPM}$  to 500, 1000, 2500, and 5000 iterations. We compare all settings, taking  $k_{SPM} = 500$  as a reference. To maintain a fair balance in the time budget given to the SPM, and thus to see the influence of  $k_{SPM}$  on the solution quality and computational time of LNS+SPM, we set the time limits for each call to the SPM to 45, 90, 225, and 450 seconds respectively. We see that frequent calls help to find better solutions, but calling the SPM too often can increase the computational time. We thus set  $k_{SPM} = 1000$  with a time limit of 90 seconds. It can be noted that calling SPM more often implies having smaller memory pools on average due to the pool management process presented in




Figure 2: Impact of  $g$  on the average computational time (in seconds) for the Morais training set; five runs were performed in each case.

Section 4.2.4. This explains why lower frequencies of SPM can lead to higher solving times. Indeed, the solver hits the time limit more often for large sets of legs.

$k_{SPM}$	500	1000	2500	5000
Average gap (%)	0	-0.05	0.18	0.70
Average computational time (%)	0	-1.7	10.3	14.3

Table 3: Impact of SPM frequency on solution quality and computational time for five runs for each instance in the training set. The reference is  $k_{SPM} = 500$ .

### 5.3. Efficiency of set partitioning and matching

The SPM component is the major contribution of this paper. To assess its efficiency, we compare LNS+SPM and LNS without SPM (we remove lines 16–23 and 23–24 of algorithm 1). Figure 3 shows the convergence curves of LNS and LNS+SPM. On the training set, for 20 000 iterations, LNS finds solutions that are 7.4% more expensive than those of LNS+SPM (4.2% on the Wen training set and 10.6% on the Morais training set). This difference is observed early in the search process. SPM accounts for an increase in the computational time of 20% on average. Thus, SPM is a key component that significantly improves the performance of LNS for the VRPCD. Figure 3 also shows that after a few thousand iterations, LNS alone is not able to improve the best solution (there are plateaus between calls to SPM). Thus, the LNS contribution is to find good legs that will then be matched by SPM.




Figure 3: Evolution of the average solution quality for LNS and LNS+SPM; 10 runs were performed for each instance in the training set. The results are normalized, with 100 representing the cost at the end for LNS+SPM.

#### 5.4. Set partitioning versus set covering formulation

In Section 4.2.2, one could formulate the MIP problem as a set covering instead of a set partitioning. The consequence being that a solution of this formulation could pickup and/or deliver some items multiple times. We modify our algorithm by solving a covering problem and removing the extra pickups and deliveries with a worst removal criterion as in Section 4.1.1. Overall, the average solution quality with the covering formulation is the same. There was no significant difference in quality during the optimization process either, only in the very first calls to the MIP is the solution not a partition. However, the runtimes are on average 12.9% longer with the covering formulation. Hence we have decided to keep the simpler set partitioning formulation.

#### 5.5. Results

In this section we recall the experiments conducted by other authors and report our results.

##### 5.5.1. Existing methods in the literature

In the literature, there are four methods for the VRPCD: Wen et al. [45], Tarantilis [40], Morais et al. [23] and Nikolopoulou et al. [24]. We evaluated all the methods on the Wen set and only the methods of [23] and [24] on the Morais set. Wen et al. [45] performed 25 runs for each instance; they report the average and best solutions found within 5 minutes and the best solution found without a time limit. Tarantilis [40] performed 3 runs for each instance with a time limit of 3000 seconds; he reports the best solution found for each instance. Morais et al. [23] performed 40 runs for each instance with a time limit of 3000 seconds for the Wen set. They report the best and average solutions for each instance for four variants of their algorithm. We report the results of their SP-ILS variant because it has the best performance. Nikolopoulou et al. [24] performed 10 runs for each instance; they report the average and best solutions found as well as the average runtime.

In what follows, we first compare results found with our methods with all other existing methods. Since we improve many best known solutions, we then make comparisons to match the results of existing methods taking into account the difference in performance in the computers used.

### 5.5.2. Best and average results

For each instance LNS+SPM was run ten times. Tables 4 and 5 give our average and best results for the Wen set and compare them to those of the existing methods. The lower bounds are those reported in [45]. We also give the percentage gap between the solution and the lower bound.

Instance	LB	Wen et al.		Morais et al.		Nikolopoulou et al.		LNS+SPM		
		Value	Gap (%)	Value	Gap (%)	Value	Gap (%)	Value	Gap (%)	Time (s)
50a	6340.90	6534.2	3.05	6477.72	2.16	6470.54	2.04	<b>6463.60</b>	<b>1.94</b>	209
50b	7201.89	7504.9	4.21	7443.92	3.36	7456.67	3.54	<b>7427.45</b>	<b>3.13</b>	545
50c	7241.05	7440.0	2.75	7441.64	2.77	7390.40	2.06	<b>7320.45</b>	<b>1.10</b>	261
50d	6887.93	7107.6	3.19	7063.17	2.54	7057.98	2.47	<b>7040.59</b>	<b>2.22</b>	705
50e	7347.54	7629.4	3.84	7514.02	2.27	7567.17	2.99	<b>7479.04</b>	<b>1.79</b>	272
100b	14200.48	14770.9	4.02	14498.69	2.10	14593.98	2.77	<b>14376.71</b>	<b>1.24</b>	778
100c	13631.24	14145.0	3.77	13993.00	2.65	14016.02	2.82	<b>13828.04</b>	<b>1.44</b>	1009
100d	13395.33	13949.6	4.14	13776.76	2.85	13778.57	2.86	<b>13600.80</b>	<b>1.53</b>	738
100e	13745.60	14396.1	4.73	14159.96	3.01	14210.54	3.38	<b>13958.75</b>	<b>1.55</b>	712
150a	19012.02	19871.3	4.52	19726.52	3.76	19705.74	3.65	<b>19401.77</b>	<b>2.05</b>	1911
150b	20371.08	21284.0	4.48	20986.64	3.02	20962.36	2.9	<b>20672.16</b>	<b>1.48</b>	1959
150c	19419.55	20320.5	4.64	20150.90	3.77	20042.94	3.21	<b>19771.90</b>	<b>1.81</b>	1862
150d	20013.37	20891.3	4.39	20656.44	3.21	20626.76	3.06	<b>20356.65</b>	<b>1.72</b>	1760
150e	19141.66	20034.6	4.66	19882.60	3.87	19800.79	3.44	<b>19493.53</b>	<b>1.84</b>	1525
200a	26538.53	27683.9	4.32	27391.74	3.22	27449.99	3.43	<b>26863.54</b>	<b>1.23</b>	2993
200b	26722.88	27989.1	4.74	27694.50	3.64	27663.96	3.52	<b>27295.45</b>	<b>2.14</b>	2986
200c	25607.31	26654.1	4.09	26490.33	3.45	26464.01	3.35	<b>26087.30</b>	<b>1.87</b>	2835
200d	26969.42	28088.2	4.15	27825.63	3.17	27767.50	2.96	<b>27394.04</b>	<b>1.57</b>	2774
200e	25776.01	26868.6	4.24	26753.12	3.79	26618.72	3.27	<b>26108.69</b>	<b>1.29</b>	2691

Table 4: Average values for the Wen set; LNS+SPM was run ten times for each instance with 20 000 iterations as the stopping criterion.

Instance	LB	Wen et al.		Tarantilis		Morais et al.		Nikolopoulou et al.		LNS+SPM		
		Value	Gap (%)	Value	Gap (%)	Value	Gap (%)	Value	Gap (%)	Value	Gap (%)	Time (s)
50a	6340.90	6471.9	2.07	<b>6450.28</b>	<b>1.73</b>	6453.08	1.77	<b>6450.28</b>	<b>1.72</b>	6455.77	1.81	149
50b	7201.89	7410.6	2.9	7428.54	3.15	7434.90	3.24	7428.54	3.15	<b>7320.77</b>	<b>1.65</b>	240
50c	7241.05	7330.6	1.24	<b>7311.77</b>	<b>0.98</b>	7317.35	1.05	<b>7311.77</b>	<b>0.98</b>	<b>7311.77</b>	<b>0.98</b>	120
50d	6887.97	7050.3	2.36	<b>7021.39</b>	<b>1.94</b>	7035.50	2.14	7028.22	2.04	7028.69	2.04	681
50e	7347.54	7516.8	2.30	<b>7451.42</b>	<b>1.41</b>	7482.01	1.83	<b>7451.42</b>	<b>1.41</b>	7452.83	1.43	168
100b	14200.48	14526.1	2.29	14405.52	1.44	14441.01	1.69	14398.17	1.39	<b>14349.60</b>	<b>1.05</b>	828
100c	13631.24	13967.8	2.47	13889.22	1.89	13932.78	2.21	13869.80	1.75	<b>13784.70</b>	<b>1.13</b>	688
100d	13395.33	13763.3	2.75	<b>13564.23</b>	<b>1.26</b>	13708.81	2.34	13603.03	1.55	13577.20	1.36	614
100e	13745.60	14212.7	3.4	14059.62	2.28	14122.32	2.74	14063.29	2.31	<b>13943.10</b>	<b>1.44</b>	601
150a	19012.02	19537.3	2.76	19638.04	3.29	19532.28	2.74	19391.16	1.99	<b>19358.90</b>	<b>1.82</b>	1886
150b	20371.08	20974.8	2.96	20922.27	2.71	20823.40	2.22	20764.50	1.93	<b>20581.50</b>	<b>1.03</b>	2093
150c	19419.55	20126.5	3.64	20019.50	3.09	19964.59	2.81	19864.86	2.29	<b>19726.80</b>	<b>1.58</b>	2187
150d	20013.37	20549.4	2.68	20600.33	2.93	20509.97	2.48	20355.27	1.71	<b>20318.80</b>	<b>1.53</b>	1660
150e	19141.66	19848.5	3.69	19782.00	3.35	19716.87	3.01	19634.47	2.57	<b>19449.50</b>	<b>1.61</b>	1546
200a	26538.53	27324.4	2.96	27397.31	3.24	27112.48	2.16	27073.57	2.02	<b>26816.50</b>	<b>1.05</b>	2988
200b	26722.88	27637.7	3.42	27582.87	3.22	27509.08	2.94	27337.49	2.3	<b>27215.10</b>	<b>1.84</b>	2804
200c	25607.31	26358.6	2.93	26425.29	3.19	26320.39	2.78	26181.73	2.24	<b>25926.00</b>	<b>1.24</b>	2376
200d	26969.42	27749.7	2.89	27818.77	3.15	27686.75	2.66	27439.50	1.74	<b>27328.70</b>	<b>1.33</b>	2457
200e	25776.01	26620.6	3.28	26704.81	3.60	26443.29	2.59	26305.30	2.05	<b>26063.50</b>	<b>1.12</b>	2240

Table 5: Best solutions found for the Wen set; LNS+SPM was run ten times for each instance with 20 000 iterations as the stopping criterion.

Instance	Morais et al.	Nikolopoulou et al.	LNS+SPM	
	Value	Value	Value	Time (min)
R1-4-1	15530.10	15282.80	<b>15211.2</b>	59.6
R1-4-2	14996.86	14730.59	<b>14666.2</b>	67.1
R1-4-3	14414.90	14194.88	<b>14192.0</b>	63.6
R1-4-4	15622.74	15371.78	<b>15336.4</b>	56.8
R1-6-1	33776.88	32864.91	<b>32748.1</b>	153.4
R1-6-2	33744.43	32893.13	<b>32726.5</b>	124.3
R1-6-3	33478.77	32841.05	<b>32658.6</b>	145.9
R1-6-4	33606.97	32946.17	<b>32850.6</b>	129.7
R1-8-1	60611.89	59178.75	<b>59046.5</b>	181.3
R1-8-2	58420.03	<b>57131.87</b>	57137.6	215.0
R1-8-3	58859.03	<b>57442.73</b>	57653.7	248.2
R1-8-4	60834.83	59494.36	<b>59427.9</b>	218.6
R1-10-1	94687.60	<b>92246.05</b>	92289.7	302.5
R1-10-2	93718.82	91465.13	<b>91360.2</b>	294.4
R1-10-3	94200.82	91536.84	<b>91504.6</b>	217.1
R1-10-4	94795.34	<b>92318.91</b>	92804.0	225.8

Table 6: Average values for the Morais set; LNS+SPM was run ten times for each instance with 20 000 iterations as the stopping criterion.

Instance	Morais et al.	Nikolopoulou et al.	LNS+SPM	
	Value	Value	Value	Time (min)
R1-4-1	15445.28	15208.84	<b>15170.0</b>	56.8
R1-4-2	14850.75	<b>14614.02</b>	14626.9	57.2
R1-4-3	14332.27	<b>14101.73</b>	14146.6	54.6
R1-4-4	15521.49	<b>15282.02</b>	15293.3	56.3
R1-6-1	33511.04	32696.90	<b>32598.1</b>	154.0
R1-6-2	33540.56	<b>32623.17</b>	32628.7	123.4
R1-6-3	33282.54	32624.50	<b>32571.5</b>	129.7
R1-6-4	33468.72	32748.70	<b>32746.3</b>	139.9
R1-8-1	60300.22	58961.82	<b>58831.1</b>	186.1
R1-8-2	58113.83	<b>56894.76</b>	56956.4	179.2
R1-8-3	58558.94	<b>57124.27</b>	57421.7	224.5
R1-8-4	60502.26	<b>59169.87</b>	59295.3	233.7
R1-10-1	94080.68	<b>91657.18</b>	91949.2	280.2
R1-10-2	92792.34	<b>91001.23</b>	91005.8	297.7
R1-10-3	93222.85	<b>91016.40</b>	91317.0	221.7
R1-10-4	94372.82	<b>92112.35</b>	92341.1	227.3

Table 7: Best solutions found for the Morais set; LNS+SPM was run ten times for each instance with 20 000 iterations as the stopping criterion.



Our method finds better solution than the existing methods: it has better average results for all but four of the instances, and it improves the best known solutions for 19 of the 35 instances, with an average improvement of 0.57%.

### 5.5.3. Comparison with existing methods

We observe that our method has often higher runtime than those of competing methods, but as Figure 3 shows, its convergence seems reasonably fast. Since we are able to improve many best known solutions, in Tables 8 to 14 we report the time taken by LNS+SPM to find solutions at least as good as those reported by other authors. For a fair comparison, we take into account the difference in the speed of the processors used. Following [17], we use CINT2000 and CINT2006<sup>1</sup> to normalize the performance of the processors: [45] used a processor about 4.7 times slower than ours, [40] used a processor about 2.4 times slower than ours, [23] used a processor about 2.0 times slower than ours, while Nikolopoulou et al. used a processor about 1.3 faster than ours. Note that these figures are only estimates, and should be interpreted with caution.

Instance	Value to match	Number of matching runs	Original time	Converted time	Runtime for LNS+SPM
50a	6534.2	10/10	17	3,6	47.0
50b	7504.9	10/10	19	4	59.0
50c	7440.0	10/10	20	4,3	23.0
50d	7107.6	10/10	20	4,3	72.0
50e	7629.4	10/10	16	3,4	41.0
100b	14770.9	10/10	56	11,9	64.0
100c	14145.0	10/10	57	12,1	69.0
100d	13949.6	10/10	57	12,1	108.0
100e	14396.1	10/10	63	13,4	79.0
150a	19871.3	10/10	139	29,6	139.0
150b	21284.0	10/10	125	26,6	104.0
150c	20320.5	10/10	139	29,6	160.0
150d	20891.3	10/10	123	26,2	177.0
150e	20034.6	10/10	140	29,8	142.0
200a	27683.9	10/10	273	58,1	217.0
200b	27989.1	10/10	278	59,1	174.0
200c	26654.1	10/10	282	60	304.0
200d	28088.2	10/10	296	63	229.0
200e	26868.6	10/10	275	58,5	298.0

Table 8: Number of runs and time taken by LNS+SPM to find solutions at least as good as the *average* results reported by Wen et al. [45]. Column *Converted time* corresponds to the estimated runtime of the Wen et al.’s algorithm on our computer.

Column *Runtime for LNS+SPM* is averaged over all matching runs

Among all the methods, that of Wen et al. is the only one designed with short runtimes in mind (at most five minutes). Indeed, it can quickly find good solutions (see Table 8), but this focus probably limited its capacity to find very good solutions (see Table 9). Table 10 shows that, except for small instances and with its proposed settings, the method of Tarantilis et al. is outperformed by LNS+SPM, as the later is able to find solution of similar quality significantly faster. Morais et al. [23] only reports their time limits, and not their runtimes, hence making the comparison less direct. It is interesting to point out that they use a time limit of 3000 seconds for the Wen et al.’s instances and only 1200 seconds for the Morais et al.’s instances that are yet larger. This probably explains the difference, shown in Tables 11 and 12, between the Wen et

<sup>1</sup>See the Standard Performance Evaluation Corporation website: <http://www.spec.org>

Instance	Value to match	Matching runs	Original time (s)	Converted time (s)	Runtime for LNS+SPM (s)
50a	6497.3	9/10	3865	822,3	19.0
50b	7466.3	1/10	3185	677,7	147.0
50c	7350.5	10/10	3269	695,5	11.0
50d	7074.0	9/10	3658	778,3	99.0
50e	7571.5	10/10	3159	672,1	11.0
100b	14646.8	10/10	9967	2120,6	52.0
100c	14056.4	10/10	10677	2271,7	84.0
100d	13844.4	10/10	11177	2378,1	65.0
100e	14300.4	10/10	10643	2264,5	101.0
150a	19784.0	10/10	24326	5175,7	199.0
150b	21098.1	10/10	24461	5204,5	107.0
150c	20166.2	10/10	23754	5054	205.0
150d	20747.2	10/10	24468	5206	239.0
150e	19888.5	10/10	23400	4978,7	154.0
200a	27537.4	10/10	46586	9911,9	270.0
200b	27851.7	10/10	43653	9287,9	283.0
200c	26472.5	10/10	46389	9870	418.0
200d	27935.3	10/10	46615	9918,1	278.0
200e	26703.4	10/10	45649	9712,6	279.0

Table 9: Number of runs and time taken by LNS+SPM to find solutions as least as good as the *best* solutions reported by Wen et al. [45]. Column *Converted time* corresponds to the estimated runtime of the Wen et al.'s algorithm on our computer. Column *Runtime for LNS+SPM* is the minimum runtime of LNS+SPM to find such solution

Instance	Value to match	Matching runs	Original time (s)	Converted time (s)	Runtime for LNS+SPM (s)
50a	6450.28	0/10	28	11,7	-
50b	7428.54	1/10	15	6,3	147.0
50c	7311.77	2/10	13	5,4	21.0
50d	7021.39	0/10	117	48,8	-
50e	7451.42	0/10	20	8,3	-
100b	14405.52	9/10	987	411,3	170.0
100c	13889.22	10/10	904	376,7	122.0
100d	13564.23	0/10	866	360,8	-
100e	14059.62	10/10	922	384,2	139.0
150a	19638.04	10/10	1302	542,5	106.0
150b	20922.27	10/10	1172	488,3	202.0
150c	20019.50	10/10	1004	418,3	259.0
150d	20600.33	10/10	673	280,4	207.0
150e	19782.00	10/10	877	365,4	154.0
200a	27397.31	10/10	1891	787,9	270.0
200b	27582.87	10/10	1665	693,8	287.0
200c	26425.29	10/10	1904	793,3	284.0
200d	27818.77	10/10	1789	745,4	278.0
200e	26704.81	10/10	1102	459,2	279.0

Table 10: Number of runs and time taken by LNS+SPM to find solutions as least as good as the *best* solutions reported by Tarantilis [40]

Instance	Value to match	Matching runs	Original time (s)	Converted time (s)	Runtime for LNS+SPM (s)
50a	6477.72	9/10	3000	1500	64.0
50b	7443.92	10/10	3000	1500	183.0
50c	7441.64	10/10	3000	1500	23.0
50d	7063.17	10/10	3000	1500	118.0
50e	7514.02	10/10	3000	1500	141.0
100b	14498.69	10/10	3000	1500	108.0
100c	13993.00	10/10	3000	1500	123.0
100d	13776.73	10/10	3000	1500	209.0
100e	14159.96	10/10	3000	1500	167.0
150a	19726.52	10/10	3000	1500	213.0
150b	20986.64	10/10	3000	1500	197.0
150c	20150.90	10/10	3000	1500	233.0
150d	20656.44	10/10	3000	1500	274.0
150e	19882.60	10/10	3000	1500	208.0
200a	27391.74	10/10	3000	1500	304.0
200b	27694.50	10/10	3000	1500	309.0
200c	26490.33	10/10	3000	1500	365.0
200d	27825.63	10/10	3000	1500	322.0
200e	26753.12	10/10	3000	1500	342.0
R1-4-1	15530.10	10/10	1200	600	792.0
R1-4-2	14996.86	10/10	1200	600	918.0
R1-4-3	14414.90	10/10	1200	600	1084.0
R1-4-4	15622.74	10/10	1200	600	796.0
R1-6-1	33776.88	10/10	1200	600	2155.0
R1-6-2	33744.43	10/10	1200	600	1701.0
R1-6-3	33478.77	10/10	1200	600	2031.0
R1-6-4	33606.97	10/10	1200	600	2084.0
R1-8-1	60611.89	10/10	1200	600	3247.0
R1-8-2	58420.03	10/10	1200	600	4225.0
R1-8-3	58859.03	10/10	1200	600	5157.0
R1-8-4	60834.83	10/10	1200	600	4090.0
R1-10-1	94687.60	10/10	1200	600	5918.0
R1-10-2	93718.82	10/10	1200	600	5553.0
R1-10-3	94200.82	10/10	1200	600	3962.0
R1-10-4	94795.34	10/10	1200	600	5099.0

Table 11: Number of runs and time taken by LNS+SPM to find solutions at least as good as the *average* results reported by Morais et al. [23]

Instance	Value to match	Matching runs	Original time (s)	Converted time (s)	Runtime for LNS+SPM (s)
50a	6453.08	0/10	3000	1500	-
50b	7434.90	3/10	3000	1500	109.0
50c	7317.35	2/10	3000	1500	16.0
50d	7035.50	2/10	3000	1500	262.0
50e	7482.01	8/10	3000	1500	31.0
100b	14441.01	10/10	3000	1500	98.0
100c	13932.78	10/10	3000	1500	100.0
100d	13708.81	10/10	3000	1500	65.0
100e	14122.32	10/10	3000	1500	124.0
150a	19532.28	10/10	3000	1500	307.0
150b	20823.40	10/10	3000	1500	298.0
150c	19964.59	10/10	3000	1500	314.0
150d	20509.97	10/10	3000	1500	310.0
150e	19716.87	10/10	3000	1500	208.0
200a	27112.48	10/10	3000	1500	295.0
200b	27509.98	10/10	3000	1500	292.0
200c	26320.39	10/10	3000	1500	423.0
200d	27686.75	10/10	3000	1500	392.0
200e	26443.29	10/10	3000	1500	417.0
R1-4-1	15445.28	10/10	1200	600	734.0
R1-4-2	14850.75	10/10	1200	600	1104.0
R1-4-3	14332.27	10/10	1200	600	959.0
R1-4-4	15521.49	10/10	1200	600	759.0
R1-6-1	33511.04	10/10	1200	600	1840.0
R1-6-2	33540.56	10/10	1200	600	1086.0
R1-6-3	33282.54	10/10	1200	600	1514.0
R1-6-4	33468.72	10/10	1200	600	1592.0
R1-8-1	60300.22	10/10	1200	600	2890.0
R1-8-2	58113.85	10/10	1200	600	3634.0
R1-8-3	58558.94	10/10	1200	600	4219.0
R1-8-4	60502.26	10/10	1200	600	3482.0
R1-10-1	94080.68	10/10	1200	600	5475.0
R1-10-2	92792.34	10/10	1200	600	4814.0
R1-10-3	93222.85	10/10	1200	600	4219.0
R1-10-4	94372.82	10/10	1200	600	3265.0

Table 12: Number of runs and time taken by LNS+SPM to find solutions as least as good as the *best* solutions reported by Morais et al. [23]

Instance	Value to match	Matching runs	Original time (s)	Converted time (s)	Runtime for LNS+SPM (s)
50a	6470.54	9/10	68	88,4	76.0
50b	7456.67	10/10	8	10,4	111.0
50c	7390.4	10/10	39	50,7	28.0
50d	7057.98	10/10	22	28,6	133.0
50e	7567.17	10/10	32	41,6	57.0
100b	14593.98	10/10	444	577,2	83.0
100c	14016.02	10/10	227	295,1	114.0
100d	13778.57	10/10	169	219,7	209.0
100e	14210.54	10/10	417	542,1	144.0
150a	19705.74	10/10	350	455	213.0
150b	20962.36	10/10	203	263,9	218.0
150c	20042.94	10/10	326	423,8	301.0
150d	20626.76	10/10	366	475,8	291.0
150e	19800.79	10/10	269	349,7	256.0
200a	27449.99	10/10	751	976,3	304.0
200b	27663.96	10/10	602	782,6	322.0
200c	26464.01	10/10	1024	1331,2	394.0
200d	27767.5	10/10	1631	2120,3	403.0
200e	26618.72	10/10	2033	2642,9	388.0
R1-4-1	15282.80	10/10	1574	2046,2	1807.0
R1-4-2	14730.59	10/10	1144	1487,2	2624.0
R1-4-3	14194.88	5/10	973	1264,9	2741.0
R1-4-4	15371.78	9/10	1009	1311,7	2155.0
R1-6-1	32864.91	10/10	4492	5839,6	6295.0
R1-6-2	32893.13	10/10	3502	4552,6	4358.0
R1-6-3	32841.05	10/10	2944	3827,2	4606.0
R1-6-4	32946.17	8/10	2800	3640	5228.0
R1-8-1	59178.75	9/10	3489	4535,7	8946.0
R1-8-2	57131.87	4/10	6259	8136,7	11418.0
R1-8-3	57442.73	1/10	5164	6713,2	13514.0
R1-8-4	59494.36	8/10	4175	5427,5	10915.0
R1-10-1	92246.05	5/10	5211	6774,3	16867.0
R1-10-2	91465.13	7/10	6670	8671	15581.0
R1-10-3	91536.84	7/10	6166	8015,8	12489.0
R1-10-4	92318.91	0/10	5662	7360,6	-

Table 13: Number of runs and time taken by LNS+SPM to find solutions at least as good as the *average* results reported by Nikolopoulou et al. [24]

Instance	Value to match	Matching runs	Original time (s)	Converted time (s)	Runtime for LNS+SPM (s)
50a	6450.28	0/10	68	88,4	-
50b	7428.54	1/10	8	10,4	147.0
50c	7311.77	2/10	39	50,7	21.0
50d	7028.22	0/10	22	28,6	-
50e	7451.42	0/10	32	41,6	-
100b	14398.17	9/10	444	577,2	170.0
100c	13869.80	9/10	227	295,1	122.0
100d	13603.03	8/10	169	219,7	359.0
100e	14063.29	10/10	417	542,1	139.0
150a	19391.16	4/10	350	455	714.0
150b	20764.50	9/10	203	263,9	298.0
150c	19864.86	10/10	326	423,8	403.0
150d	20355.27	6/10	366	475,8	673.0
150e	19634.47	10/10	269	349,7	322.0
200a	27073.57	10/10	751	976,3	295.0
200b	27337.49	7/10	602	782,6	702.0
200c	26181.73	10/10	1024	1331,2	699.0
200d	27439.50	9/10	1631	2120,3	871.0
200e	26305.30	10/10	2033	2642,9	548.0
R1-4-1	15208.84	4/10	1574	2046,2	2466.0
R1-4-2	14614.02	0/10	1144	1487,2	-
R1-4-3	14101.73	0/10	973	1264,9	-
R1-4-4	15282.02	0/10	1009	1311,7	-
R1-6-1	32696.90	3/10	4492	5839,6	7139.0
R1-6-2	32623.17	0/10	3502	4552,6	-
R1-6-3	32624.50	3/10	2944	3827,2	5499.0
R1-6-4	32748.70	2/10	2800	3640	7670.0
R1-8-1	58961.82	2/10	3489	4535,7	7910.0
R1-8-2	56894.76	0/10	6259	8136,7	-
R1-8-3	57124.27	0/10	5164	6713,2	-
R1-8-4	59169.87	0/10	4175	5427,5	-
R1-10-1	91657.18	0/10	5211	6774,3	-
R1-10-2	91001.23	0/10	6670	8671	-
R1-10-3	91016.40	0/10	6166	8015,8	-
R1-10-4	92112.35	0/10	5662	7360,6	-

Table 14: Number of runs and time taken by LNS+SPM to find solutions as least as good as the *best* solutions reported by Nikolopoulou et al. [24]

al.'s instances where LNS+SPM seems faster, and the Morais et al.'s instances where we can not find as good solution as quickly but where we can significantly improves their results. Finally, compared to Nikolopoulou et al.'s method (see Tables 13 and 14), LNS+SPM seems to have the advantage for Wen et al.'s instances, while being generally slower for the Morais et al.'s instances, especially for the larger R-8 and R-10 series. Many design decisions (stopping criteria, inherent tendency towards intensification or diversification of each method...) can account for such differences. Generally speaking, in line with the results of Section 5.5.2, LNS+SPM is very good for a wide range of medium size instances (Wen-100 to Wen-200, Morais R-4 and R-6), while being outperformed on small (Wen-50) and large instances (Morais-R8 and R-10). A reasonable explanation seems to be the memory and set partitioning component. In small size instances, the time spent managing these components would probably be better invested in simply exploring the solution space. On the other hand, for large size instances, even the last calls to the SPM reach the time limit defined in CPLEX (which is not the case for medium size instances). This suggest that a memory management procedure would be useful for large size instances.

## 6. Concluding remarks

This paper presents a new method for the VRPCD based on LNS and regular calls to an SPP. The SPP component is solved using both a MIP solver and a CP solver. This component helps to find solutions that are significantly better than those obtained by LNS alone. The proposed method has been tested on two benchmarks. It improves many of the previously best known results and average results. Compared to existing methods, it compares favorably for a wide range of medium size instances.

Solving a CP subproblem provides a simple and efficient way to integrate precedence constraints into the SPP. It would be interesting to investigate whether this method could be adapted to solve other VRPs with synchronization-related constraints, and VRPs with multiple trips. In particular, the case where requests are not constrained to go through the cross-dock between their pickup and delivery locations is also an interesting perspective. This problem has been addressed in previous studies under the name *Pickup and Delivery Problem with Cross-Docking Opportunity* [27], or *Pickup and Delivery Problem with Cross-Docking* [34].

## Acknowledgements

This work was partially supported by the Canadian Natural Science and Engineering Research Council (RGPIN-2015-04696) and by the Fonds de recherche du Québec - Nature et technologies through its Team research Program.

- [1] D. Agustina, C. K. M. Lee, R. Piplani, A review: Mathematical models for cross docking planning, *International Journal of Engineering Business Management* 2 (2) (2010) 47–54.

- [2] C. Archetti, M. G. Speranza, A survey on matheuristics for routing problem, *EURO Journal on Computational Optimization* 2 (2014) 223–246.
- [3] C. Archetti, M. G. Speranza, M. W. P. Savelsbergh, An optimization-based heuristic for the split delivery vehicle routing problem, *Transportation Science* 42 (1) (2008) 22–31.
- [4] J. C. Beck, M. S. Fox, Scheduling alternative activities, *Proceedings of the Sixteenth National Conference on Artificial Intelligence and Eleventh Conference on Innovative Applications of Artificial Intelligence* (1999) 680–687.
- [5] N. Boysen, M. Fliedner, Cross dock scheduling: Classification, literature review and research agenda, *Omega* 38 (6) (2010) 413–422.
- [6] P. Buijs, I. F. Vis, H. J. Carlo, Synchronization in cross-docking networks: A research classification and framework, *European Journal of Operational Research* 239 (3) (2014) 593–608.
- [7] C. E. Cortés, M. Matamala, C. Contardo, The pickup and delivery problem with transfers: Formulation and a branch-and-cut solution method, *European Journal of Operational Research* 200 (3) (2010) 711–724.
- [8] R. Cuda, G. Guastaroba, M. G. Speranza, A survey on two-echelon routing problems, *Computers & Operations Research* 55 (2015) 185–199.
- [9] K. F. Doerner, V. Schmid, Survey: Matheuristics for rich vehicle routing problems, in: *Hybrid Metaheuristics*, vol. 6373 of LNCS, Springer, 2010, pp. 206–221.
- [10] R. Dondo, J. Cerdá, A monolithic approach to vehicle routing and operations scheduling of a cross-dock system with multiple dock doors, *Computers & Chemical Engineering* 63 (2014) 184–205.
- [11] M. Drexl, Synchronization in vehicle routing—A survey of VRPs with multiple synchronization constraints, *Transportation Science* 46 (3) (2012) 297–316.
- [12] F. Enderer, Integrating dock-door assignment and vehicle routing in cross-docking, Ph.D. thesis, Concordia University (2014).
- [13] P. Grangier, M. Gendreau, F. Lehuédé, L.-M. Rousseau, An adaptive large neighborhood search for the two-echelon multiple-trip vehicle routing problem with satellite synchronization, *Tech. rep., CIRRELT* 2014-33 (2014).
- [14] A. Grimault, N. Bostel, F. Lehuédé, An adaptive large neighborhood search for the full truckload pickup and delivery problem with resource synchronization, *Tech. rep., Ecole des Mines de Nantes* 15/4/AUTO (2015).
- [15] IBM Corporation, IBM ILOG CPLEX Optimization Studio V12.6.1 documentation (2014).



- [16] E. Lam, P. V. Hentenryck, A branch-and-price-and-check model for the vehicle routing problem with location congestion, *Constraints* (2016) 1–19.
- [17] G. Laporte, S. Ropke, T. Vidal, Heuristics for the vehicle routing problem, in: P. Toth, D. Vigo (eds.), *Vehicle Routing Problems: Problems, Methods, and Applications*, 2nd ed., chap. 4, MOS-SIAM, 2014, pp. 87–116.
- [18] Y. H. Lee, J. W. Jung, K. M. Lee, Vehicle routing scheduling for cross-docking in the supply chain, *Computers & Industrial Engineering* 51 (2) (2006) 247–256.
- [19] C.-J. Liao, Y. Lin, S. C. Shih, Vehicle routing with cross-docking in the supply chain, *Expert Systems with Applications* 37 (10) (2010) 6868–6873.
- [20] R. Masson, F. Lehuédé, O. Péton, An adaptive large neighborhood search for the pickup and delivery problem with transfers, *Transportation Science* 47 (3) (2013) 344–355.
- [21] R. Masson, F. Lehuédé, O. Péton, Efficient feasibility testing for request insertion in the pickup and delivery problem with transfers, *Operations Research Letters* 41 (3) (2013) 211–215.
- [22] J. Mendoza, L.-M. Rousseau, J. G. Villegas, A hybrid metaheuristic for the vehicle routing problem with stochastic demand and duration constraint, *Journal of Heuristics*.
- [23] V. W. Morais, G. R. Mateus, T. F. Noronha, Iterated local search heuristics for the vehicle routing problem with cross-docking, *Expert Systems with Applications* 41 (16) (2014) 7495–7506.
- [24] A. I. Nikolopoulou, P. P. Repoussis, C. D. Tarantilis, E. E. Zachariadis, Moving products between location pairs: Cross-Docking versus Direct-Shipping, *European Journal of Operational Research*.
- [25] S. N. Parragh, V. Schmid, Hybrid column generation and large neighborhood search for the dial-a-ride problem, *Computers and Operations Research* 40 (1) (2013) 490–497.
- [26] G. Perboli, R. Tadei, D. Vigo, The two-echelon capacitated vehicle routing problem: models and math-based heuristics, *Transportation Science* 45 (3) (2011) 364–380.
- [27] H. L. Petersen, S. Ropke, The pickup and delivery problem with cross-docking opportunity, in: *Computational Logistics*, Springer, 2011, pp. 101–113.
- [28] V. Pillac, C. Guéret, A. Medaglia, A parallel matheuristic for the technician routing and scheduling problem, *Optimization Letters* 7 (7) (2013) 1525–1535.
- [29] D. Pisinger, S. Ropke, A general heuristic for vehicle routing problems, *Computers & Operations Research* 34 (8) (2007) 2403–2435.
- [30] Y. Qu, J. F. Bard, A GRASP with adaptive large neighborhood search for pickup and delivery problems with transshipment, *Computers & Operations Research* 39 (10) (2012) 2439–2456.

- [31] Y. Rochat, É. D. Taillard, Probability diversification and intensification in local search for vehicle routing, *Journal of Heuristics* 1 (1995) 147–167.
- [32] S. Ropke, D. Pisinger, An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows, *Transportation Science* 40 (4) (2006) 455–472.
- [33] F. A. Santos, G. R. Mateus, A. S. da Cunha, A novel column generation algorithm for the vehicle routing problem with cross-docking, in: *Network Optimization - INOC 2011*, vol. 6701 of LNCS, 2011, pp. 412–425.
- [34] F. A. Santos, G. R. Mateus, A. S. da Cunha, The pickup and delivery problem with cross-docking, *Computers & Operations Research* 40 (4) (2013) 1085–1093.
- [35] F. A. Santos, G. R. Mateus, A. Salles da Cunha, A branch-and-price algorithm for a vehicle routing problem with cross-docking, *Electronic Notes in Discrete Mathematics* 37 (2011) 249–254.
- [36] M. W. P. Savelsbergh, Local search in routing problems with time windows, *Annals of Operations Research* 4 (1) (1985) 285–305.
- [37] P. Shaw, Using constraint programming and local search methods to solve vehicle routing problems, in: *Principles and Practice of Constraint Programming CP98*, vol. 1520 of LNCS, 1998, pp. 417–431.
- [38] G. Stalk, P. Evans, L. E. Schulman, Competing on capabilities: The new rules of corporate strategy, *Harvard Business Review* 70 (2) (1992) 57–69.
- [39] A. Subramanian, E. Uchoa, L. S. Ochi, A hybrid algorithm for a class of vehicle routing problems, *Computers and Operations Research* 40 (10) (2013) 2519–2531.
- [40] C. D. Tarantilis, Adaptive multi-restart tabu search algorithm for the vehicle routing problem with cross-docking, *Optimization Letters* 7 (7) (2012) 1583–1596.
- [41] E. S. Thorsteinsson, Branch-and-check: A hybrid framework integrating mixed integer programming and constraint logic programming, in: T. Walsh (ed.), *Principles and Practice of Constraint Programming CP 2001*, vol. 2239 of LNCS, 2001, pp. 16–30.
- [42] J. Van Belle, P. Valckenaers, D. Cattrysse, Cross-docking: state of the art, *Omega* 40 (6) (2012) 827–846.
- [43] P. Van Hentenryck, *The OPL Optimization Programming Language*, MIT Press, 1999.
- [44] J. G. Villegas, C. Prins, C. Prodhon, A. Medaglia, N. Velasco, A matheuristic for the truck and trailer routing problem, *European Journal of Operational Research* 230 (2) (2013) 231–244.
- [45] M. Wen, J. Larsen, J. Clausen, J.-F. Cordeau, G. Laporte, Vehicle routing with cross-docking, *Journal of the Operational Research Society* 60 (12) (2008) 1708–1718.