On efficient testing of capacity constraints in pickup-and-delivery problems with trailers

Michael Drex1^{1,2}

Received: 14 June 2019 / Revised: 4 May 2020 / Published online: 23 July 2020
© The Author(s) 2020

Abstract

Efficient feasibility tests are important in many heuristics for routing problems. This paper considers several variants of pickup-and-delivery problems with trailers. Its contribution consists in the description of constant-time procedures for testing observance of capacity constraints when inserting tasks into routes. It is demonstrated that the presence of vehicles with detachable trailers makes capacity feasibility tests considerably more involved.

Keywords Routing · Pickup-and-delivery · Trailers · Constant-time feasibility test

Mathematics Subject Classification 90B06: Transportation, logistics · 90C59: Approximation methods and heuristics

1 Introduction

Pickup-and-delivery problems (PDPs) are concerned with the transport of goods or persons from different origins to different destinations. PDPs come in several variants and have received a lot of attention in the literature (see the surveys by Parragh et al. (2008a), Parragh et al. (2008b), Doerner and Salazar-González (2014), and Battarra et al. (2014)). Whereas vehicle routing problems (VRPs, the special case of the PDP where either all pickups or all deliveries occur at a central depot) with trailers are rather well examined (see the surveys by Prodhon and Prins (2014), Section 3.3, Cuda et al. (2015), Section 4, and the more recent works by Parragh and Cordeau (2017) and Rothenbächer et al. (2018)), PDPs with trailers (PDPTs) are rarely studied in the

✉ Michael Drex1
michael.drex1@th-deg.de

¹ Faculty of Applied Natural Sciences and Industrial Engineering, Deggendorf Institute of Technology, Deggendorf, Germany

² Chair of Logistics Management, Gutenberg School of Management and Economics, Johannes Gutenberg University, Mainz, Germany

literature, despite their practical relevance. Trailers are advantageous as they increase the overall vehicle capacity. However, some locations may be accessible only by a lorry without a trailer, e.g., because of insufficient manoeuvring space. Therefore, special parking and transshipment locations (PTLs) are available where trailers can be parked while lorries visit accessibility-constrained locations. This induces a non-trivial trade-off between an enlarged vehicle capacity and the necessity of making detours to park and reattach trailers, and, as will be demonstrated in this paper, makes capacity feasibility tests considerably more involved.

There are several works on PDPs where vehicles consisting of a tractor and a semi-trailer are employed to perform full-load tasks, i.e., where a vehicle can transport only one task at a time (for example, Cheung et al. (2008), Xue et al. (2014), Tilk et al. (2018)), but the author of the present work is aware of only two papers (Bürckert et al. (2000), Drexl (2018)) dealing with less-than-truckload tasks, i.e., the case where more than one task can be on a vehicle at the same time. The present paper is concerned with PDPTs of the latter type.

Heuristics based on neighbourhood search are widely used approaches for solving many different kinds of vehicle routing problem, including the various types of PDPs. A decisive aspect for the performance of such methods is their ability to quickly test the feasibility of an insertion of a customer or task into a route. There are several theoretical papers on efficient feasibility testing, for example, Hunsaker and Savelsbergh (2002), Haugland and Ho (2010), Firat and Woeginger (2011) (dial-a-ride problems), and Masson et al. (2013) (PDP with time windows and transfers).

In a similar vein, the present work deals with efficient testing whether the insertion of a task into a feasible PDPT route *at specified positions* for the pickup and for the delivery part of the task maintains capacity feasibility of the route. ‘Efficient’ means that, given auxiliary data which is computed in a preprocessing step that requires quadratic time in the number of tasks on the route, the test itself takes constant time for any pair of positions at which to insert the pickup and the delivery of the task. For testing whether a certain task can be feasibly inserted *at all* into a certain route, in the worst case, insertion of the pickup must be tried at each position (after the first and before the last one), and insertion of the delivery must be tried at all subsequent positions (ditto). With a naïve procedure that passes linearly over each such tentative enlargement of a route, a test whether a certain task can be feasibly inserted at all takes cubic time in the number of tasks on that route. Consequently, a constant-time test for insertion at specified positions can reduce the one-round time (a term coined by Yagiura and Ibaraki (1999)), i.e., the time for determining the task to insert in one iteration of an insertion heuristic or a local search procedure, by one order of magnitude. The contribution of the present paper consists precisely in the description of such a test for several variants of PDPTs.

An efficient test for time window feasibility that can be used for all PDP(T) types considered in this paper is presented in Drexl (2018), based on earlier work by Savelsbergh (1992) and Masson et al. (2013). Other seminal papers on feasibility tests for routing problems are Savelsbergh (1985), Savelsbergh (1990), Kindervater and Savelsbergh (1997), Funke et al. (2005), Irnich et al. (2006), Irnich (2008a), Irnich (2008b), Vidal et al. (2014), Grangier et al. (2016), and Gschwind and Drexl (2019), but none of these considers routing problems with trailers.

The rest of the paper is structured as follows. In the next section, the PDP(T) types of interest in this work are specified. Then, the notation used in the paper and the formal modelling of the different PDPT variants are explained in Sect. 3. Afterwards, Sect. 4 describes how capacities can be tested in linear and in constant time. Section 5 concludes the paper with a summary and an outlook.

2 PDP variants

Following the PDP classification of Battarra et al. (2014), the problem variants covered in this paper are:

- One-to-one problems: each task consists in the transport of a particular *commodity* from a given pickup location to a given delivery location. The transported commodities are not interchangeable. This is obvious for passenger transport. In goods transport, applications arise, e.g., in letter mail and parcel services as well as full and less-than-truckload forwarding.
- One-to-many-to-one problems: each task consists in either the transport of a commodity from a central depot to a delivery location (linehaul tasks) or the transport from a pickup location to the depot (backhaul tasks). Practical applications are common in supplying supermarkets, beverage stores, apparel stores and, above all, in the less-than-truckload business. Several subtypes of one-to-many-to-one problems can be distinguished:
 - The VRP with mixed backhauls: linehaul and backhaul tasks can occur anywhere on a route.
 - The VRP with backhauls: all linehaul tasks on a route must be fulfilled before a pickup location of a backhaul task can be visited. In this and the previous subtype, it is possible that a route fulfils only linehaul or only backhaul tasks.
 - The VRP: the special case where all tasks are either linehaul or backhaul tasks.
 - The simultaneous PDP: any task location may require a delivery of a good from the depot and a pickup of another good for transport back to the depot; the delivery and the pickup must be performed during one visit.

Note that the goods to be transported in one-to-many-to-one problems may be interchangeable or not. For linehaul tasks, this makes a difference. For example, in soft drink distribution, crates with full bottles are delivered to households, and empty crates are taken back. It is irrelevant whether household A receives crate 1 of a particular soft drink brand and household B receives crate 2 of the same brand or the other way round. By contrast, in freight forwarding of consignments on pallets, each loaded pallet of a linehaul task has a specific destination. For the algorithmic treatment of such problems, these issues do not matter. What counts is that the capacity requirements of each task are taken into account correctly. In reality, however, if goods are not interchangeable, the driver must ensure that the right items are delivered to linehaul customers.

- Many-to-many PDPs: goods picked up at a location can be used to fulfil demand at several other locations (e.g., bike-sharing systems, empty container movement,

raw milk bulk transports). Mostly, only a single commodity is considered in this problem type. Many-to-many PDPs come in four sub-variants: it may be allowed or forbidden that a vehicle leave or enter the depot with some load.

For all these problem variants, it is sensible to consider the use of vehicles with detachable trailers.

3 Notation and modelling

The following subsections describe how each of the above PDPT variants is represented by a suitable digraph $D = (V, A)$. This is helpful to understand which insertions of tasks into routes are possible and how capacity feasibility can be tested. In all variants, each task t has a pickup capacity requirement $q_p^t \geq 0$ and a delivery capacity requirement $q_d^t \leq 0$. The vertex set V of D contains a start and an end depot vertex, s and e respectively, one pickup and one delivery vertex for each task, and one vertex for each PTL. The pickup and the delivery vertices are referred to as *task vertices*. Each vertex $v \in V$ has an associated capacity requirement $q_v \in \mathbb{R}$, which is zero for the depot and the PTL vertices. Moreover, for modelling some of the above PDPT variants, ‘artificial’ task vertices are used. Each task vertex is visited exactly once in a feasible solution, whereas PTL vertices can be visited more than once. All artificial vertices are reachable with a trailer. The arc set A , in principle, contains one arc for each pair of vertices, and, to model the possibility of consecutive *subroutes* (definition in next paragraph), there is a loop on each PTL vertex. Some exceptions apply: there is no arc entering s , leaving e , leading from s to a delivery vertex, from a pickup vertex to e , or from a delivery vertex to the associated pickup vertex. Moreover, there are additional restrictions on the arc sets of the different problem variants as specified in the respective subsections below.

Associated with D is a vehicle fleet made up of single lorries and lorry-trailer combinations (LTCs). All vehicles start their routes at s and end them at e . Each vehicle k has a lorry capacity of Q_k^l and a trailer capacity of Q_k^t . For single lorries k , $Q_k^t = 0$. The route of an LTC consists of the *main route*, where the lorry pulls its trailer, and zero or more *subroutes* starting and ending at a PTL where the trailer is parked while the lorry performs one or more pickups and/or deliveries. Several consecutive subroutes can start and end at the same PTL before the trailer is re-coupled and pulled away. If the pickup of a task whose delivery is performed on a subroute has been performed before this subroute, the entire load of the task must be on the lorry at the start of the subroute. This may require a load transfer from a trailer to its lorry when decoupling. In the present paper, it is assumed that such load transfers are possible without restrictions, and that, at pickup locations reachable by trailer, the load to be picked up can be split arbitrarily between a lorry and its trailer.

Although costs and times are not relevant for capacity tests, note that the locations associated with any two vertices u and v determine the distance-dependent costs and the travel times of the arcs (u, v) and (v, u) , if these exist, and are therefore relevant for the definition of the objective function and for time-window feasibility tests. With regard to the PDPTs considered here, for non-artificial (‘real’) vertices, travel costs

and times are computed on the basis of the respective physical locations, and artificial task vertices are located at the start or the end depot, as indicated in each case.

3.1 The one-to-one PDPT

The one-to-one PDPT is the simplest case. For any task t , the associated pickup vertex has a capacity requirement of $q_p^t > 0$, and the delivery vertex requires $q_d^t < 0$ capacity units. There is an arc between each pair of vertices, with the exceptions mentioned above.

3.2 The VRPT with mixed backhauls

Extending the approach by Ropke and Pisinger (2006) to trailers, this problem type is modelled as follows. Each linehaul task is represented by an artificial pickup vertex located at the central depot and a delivery vertex at the physical delivery location. Similarly, for each backhaul task, there is a pickup vertex at the physical pickup location and an artificial delivery vertex at the depot. For a linehaul or backhaul task t , the capacity requirements of linehaul and backhaul pickup and delivery vertices are $q_p^t > 0$ and $q_d^t < 0$. The only arcs leading to an artificial pickup vertex come from the start depot vertex or other artificial pickup vertices, the only arcs emanating from an artificial delivery vertex lead to other artificial delivery vertices or to the end depot. There is no arc from an artificial pickup to an artificial delivery vertex.

3.3 The VRPT with backhauls

When all linehauls must be performed before any backhaul, the setup described in the previous subsection can be modified as follows. No arc exists from an artificial pickup vertex to a real pickup vertex or from a real pickup vertex to a real delivery vertex. Moreover, it must be ensured that the precedence requirements are not undermined by visiting PTLs; i.e., sequences such as $\dots \rightarrow \text{linehaul delivery} \rightarrow \text{decoupling} \rightarrow \text{backhaul pickup} \rightarrow \text{coupling} \rightarrow \text{linehaul delivery} \rightarrow \dots$ must not be allowed. This can be achieved by storing, for each route, the position of the last linehaul delivery. When an insertion is tested and the pickup vertex is a pickup of a backhaul, i.e., a non-artificial pickup, insertion of this vertex is considered only after the last linehaul delivery position. Similarly, the last insertion position to consider for a non-artificial delivery of a linehaul is after the last non-artificial delivery. Ropke and Pisinger (2006) use a different concept, that of precedences, which is applicable in the presence of trailers as well.

3.4 The VRPT

Vehicle routing problems with trailers can be represented as instances of VRPTs with mixed backhauls without requiring further modelling.

3.5 The simultaneous PDPT

Ropke and Pisinger (2006) propose the following modelling approach for the simultaneous PDP: each task location requiring a delivery and a pickup is represented by a ‘delivery task’ and a ‘pickup task’. The delivery (pickup) vertex of each pickup (delivery) task is located at the depot. Moreover, the costs of all arcs emanating from a delivery vertex of a delivery task are set to a very high value, except for the arc leading to the pickup vertex of the partner task, which has cost zero. These high-cost arcs are necessary, because otherwise, inserting a single task would be impossible. After inserting a single task, whether a delivery or a pickup task, the next task that is inserted will always be the partner task, because this will remove a high-cost arc and will thus dramatically improve the current solution.

An alternative approach that avoids having to create two tasks for each original task, i.e., avoids doubling the instance size, and prevents difficulties with precedences in the presence of PTLs works as follows. Each task t is represented by an artificial pickup vertex located at the start depot and a delivery vertex located where the simultaneous delivery and pickup must occur. The amount to be picked up at the artificial pickup vertex equals q_d^t , the original amount to be *delivered*. The actual amount to be delivered at the delivery vertex is set to $q_p^t - q_d^t$, i.e., to (original pickup amount – original delivery amount). This means that delivery *vertices* with *positive* capacity requirement are possible (note that the delivery amount of a *task* is always non-positive), and that there can be different absolute values for the pickup and the delivery amount of a task; see the example in Sect. 4.2.1. No arc exists from a delivery vertex or from a PTL vertex to an artificial pickup vertex.

3.6 The many-to-many PDPT

Similar to the modelling of backhauling problems, a task is represented by an artificial vertex located at the depot, and similar to the approach for the simultaneous PDP(T), the absolute values of the pickup and the delivery amount of a task may differ. If the task consists of a pickup (delivery) at a vertex, the artificial vertex is located at the end (start) depot. The capacity requirement at an artificial vertex is zero; its partner vertex has the original positive (if it is a pickup) or negative (if it is a delivery) capacity requirement. The only arcs leading to an artificial pickup vertex come from the start depot or other artificial pickup vertices, the only arcs emanating from an artificial delivery vertex lead to other artificial delivery vertices or the end depot. Moreover, there is no arc from an artificial pickup to an artificial delivery vertex.

Many-to-many PDP(T)s are special in (at least) two respects. First, as mentioned, four sub-variants are possible: it may be allowed or forbidden that a vehicle leave or enter the depot with some load.

When no loading at the start depot is allowed but load may be brought to the end depot, it should be tested in a preprocessing step whether the total load to be picked up is at least as much as the total load to be delivered. If this is not the case, no feasible solution exists. Moreover, on each route, sufficiently many pickups must be performed before any deliveries are possible, and it must be ensured in the capacity

tests that the total load of a single lorry or an LTC is always non-negative. Likewise, if it is forbidden to bring load to the end depot, but loading at the start depot is allowed, it should be tested in advance whether the total load to be delivered is at least as much as the total load to be picked up, and one or more delivery tasks must be inserted into an empty route before any pickup task. If it is forbidden that a loaded vehicle leave or enter the depot, a route with one task is necessarily infeasible, and the insertion of a task into a feasible route also makes the route infeasible. This can be handled by allowing (penalized) infeasible solutions in a surrounding algorithm that uses the capacity test. Then, the return value of the test should not be simply ‘true’ or ‘false’, but the amount by which the vehicle capacity and/or the zero load balance at the start and/or end depot is/are violated.

Second, many-to-many PDP(T)s are special because they possess the property, rarely encountered in vehicle routing problems, that *removing* a task from a feasible route can make the route infeasible. This holds true irrespective of whether or not it is allowed to pick up or deliver load at the depot. For example, consider the following route, performed by a single lorry with capacity 100 (the artificial pickup and delivery vertices at the depot are omitted for simplicity):

Vertex	s	1	2	3	4	5	6	e
Capacity requirement	0	+50	+50	-50	-50	+50	-30	0

Removing vertex 5 makes the route infeasible even if loading at the depot is allowed. In the present paper, however, the focus is on (feasible) insertions, so issues arising when removing tasks are not discussed further here.

4 Feasibility tests

Generally speaking, for a route to be feasible regarding capacity constraints, it must be ensured that the vehicle capacity is not exceeded at any vertex, and that enough load can be aboard the LTC (when on the main route) or the lorry (when on a subroute) to satisfy any subsequent negative capacity requirements. Put differently, it must be ensured on the main route and on all subroutes that, when reaching a vertex with positive capacity requirement, the load in the LTC or the lorry plus this positive capacity requirement does not exceed the LTC or the lorry capacity; likewise, when reaching a vertex with negative capacity requirement, enough load must be in the LTC or the lorry to satisfy this capacity requirement. It will become clear in the following that in particular the latter aspect is non-trivial on subroutes and, for the many-to-many PDPT, also on the main route.

The algorithms and data structures presented in this section are extensions of those described in Drexl (2018) for the one-to-one PDPT. The algorithmic descriptions assume that the feasibility of an insertion of a task t with associated pickup vertex p and delivery vertex d , written as $t = (p, d)$, into a feasible route $r = (0, 1, \dots, n-1, n)$ which is performed by single lorry or LTC k is to be tested. Vertex p is to be inserted directly after position (zero-based index of the route) h ; d is to be inserted directly

after position/index i . If p cannot be reached with a trailer, r is performed by an LTC, and the trailer is attached when leaving h , a location triple $\tilde{p} = ptl_p \rightarrow p \rightarrow ptl_p$ corresponding to a new subroute must be inserted after h ; similar for i and d . ptl_p is a PTL vertex; similar for d . Note that p , d , ptl_p , and ptl_d are vertices, whereas h and i are indices on a route. However, to simplify notation, when referring to a vertex visited at a certain position on a route, only the index is used: for example, the capacity requirement of the vertex at index i is denoted by q_i .

Indices h and i indicate positions in the route *before* p and d are inserted. Hence, if $h == i$, then d is to be inserted directly after p , or, if a triple $\tilde{p} = ptl_p \rightarrow p \rightarrow ptl_p$ must be inserted, directly after the triple. If, however, d cannot be reached with a trailer and p is left with the trailer attached or a triple $\tilde{p}$ is to be inserted, then a triple $\tilde{d} = ptl_d \rightarrow d \rightarrow ptl_d$ is inserted. If $h == i$ and p and/or d must be surrounded by a decouple-couple pair, it would also be possible to surround both p and d by one such pair, but this makes no difference for capacity feasibility. Several consecutive subroutes by one LTC lorry at the same PTL are modelled by inserting a decouple-couple pair for each subroute. For simplicity, the vehicle index k is omitted: Q^l and Q^t are used instead of Q_k^l and Q_k^t to denote the lorry and the trailer capacity. The capacity requirements of the task $t = (p, d)$ to be inserted are indicated by q_p and q_d . For each problem type, q_p and q_d are determined from the original capacity requirements of task t as described in Sects. 3.1–3.6. In the pseudocode, the symbol ‘==’ serves as equality operator, ‘=’ is the assignment operator, and ‘ $x += y$ ’ is used as shortcut for ‘ $x = x + y$ ’.

4.1 Testing capacities in linear time

Testing capacity in linear time for single-lorry routes is simple: the to-be-inserted task is tentatively inserted, one pass over the route is performed, and the capacity requirement at each visited vertex is added to the total load and compared with the lorry capacity, cf. lines 16–20 in Algorithm 1. By contrast, testing capacity in the presence of trailers is not entirely straightforward. The procedure detailed in Algorithm 1 can be used. The procedure covers all problem types considered in this paper, because large parts of the routine apply to all types. As indicated in the pseudocode, lines 2–15 can be omitted for problem types other than many-to-many, and lines 23–30 can be ignored for many-to-many-problems.

For problem types other than many-to-many, it must be known at the start of a subroute how much load must be in the lorry to be able to perform the deliveries whose pickups are not on this subroute. This information is gathered in one forward pass over the route (lines 23–30) and stored in an array `LoadDeliveredButNotPickedUpOnSubroute`. Note that, in practice, it is not sufficient to have the correct *amount* of load in the lorry at the start of a subroute. It is also necessary to have the right *commodities* on the lorry, those that must be delivered on this subroute, but this must be ensured by the driver.

For many-to-many problems, one forward pass over the route is performed to determine how much load must be picked up at the depot. This amount is determined by the most negative load balance at any vertex on the route. If loading at the start depot

is not allowed and a negative load balance occurs at a vertex, the insertion is infeasible (lines 2–8). If load may be picked up at the start depot, a second pass over the route (lines 17–20 or 34–52) tests whether there is sufficient capacity taking this load into account. If no load may be brought to the end depot, but the load balance from any index to the end of the route is positive, the route is infeasible as well (lines 9–14). Moreover, for many-to-many problems, it must be ensured for all subroutes that at each non-artificial delivery vertex, sufficient load is available on the lorry to perform the delivery. This load may have been picked up earlier on the subroute or before the subroute. It is possible that not enough load has been picked up earlier on the subroute and that, at the same time, not enough load has been picked up before the subroute. This is tested in lines 50–52.

4.2 Testing capacities in constant time

In the following, it is described how capacity feasibility can be tested in constant time. Algorithm 2 is for the one-to-one and the one-to-many-to-one problem types. The constant-time test for many-to-many PDPTs differs significantly from the tests for the other problem types. Thus, it is presented separately, in Algorithm 3.

4.2.1 One-to-one and one-to-many-to-one PDPTs

The following data structures are used in the test for the one-to-one and the one-to-many-to-one PDPT variants. These were introduced by Drexler (2018) for the one-to-one PDPT.

1. $\text{MaxTotalLoadOfSegment}[i][\text{offset}]$ is the maximal load balance from the start of the route at any index from i up to and including $i + \text{offset}$. In particular, $\text{MaxTotalLoadOfSegment}[i][0]$ is the overall load picked up but not delivered yet from the start depot to and including the location at index i .

For example, consider the following route in a one-to-one PDPT:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13
Capacity requirement	0	+30	+10	0	+20	-30	+10	+15	-10	-10	0	-20	-15	0

This route contains a subroute that starts at index 3 and ends at index 10. The load balances at indices 1 to 6 are +30, +40, +40, +60, +30, and +40; thus, $\text{MaxTotalLoadOfSegment}[1][5] = +60$. Moreover, $\text{MaxTotalLoadOfSegment}[7][0] = +55$.

2. $\text{TrailerAttached}[i]$ is true if the trailer is attached when leaving i , false otherwise.
3. $\text{IndexOfLastPrecedingDecouple}[i]$ stores the index of the last decoupling that precedes i .
4. If i is the start of a subroute, $\text{LoadDeliveredButNotPickedUpOnSubroute}[i]$ is the sum of the pickup amounts of those tasks on the subroute whose pickups lie before i .

Algorithm 1 TestCapacityLinear(r, k)

Input: Route $r = (0, 1, 2, \dots, n)$ with to-be-inserted task tentatively inserted, including decoupling and coupling vertices where necessary, and capacity requirements $q_v \in \mathbb{R}$, $v = 0, 1, 2, \dots, n$

Vehicle k (single lorry or LTC) with capacities Q^l and Q^t ; for single lorries, $Q^t = 0$

Result: Returns **true** iff lorry and, if applicable, trailer capacity of k are maintained at each vertex of r ; **false** otherwise

```

1 TotalLoad = 0
2 if Problem type is many-to-many
3   LoadBalance = 0
4   for  $v = 1, 2, 3, \dots, n - 1$ 
5     LoadBalance  $+= q_v$ 
6     TotalLoad = min(TotalLoad, LoadBalance)
7     if Loading at start depot is not allowed and TotalLoad < 0
8       return false
9   if Unloading at end depot is not allowed
10    LoadBalance = 0
11    for  $v = n - 1, n - 2, n - 3, \dots, 1$ 
12      LoadBalance  $+= q_v$ 
13      if LoadBalance > 0
14        return false
15  TotalLoad =  $(-1) \cdot$  TotalLoad
16 if  $Q^t = 0$  // Test for single lorries
17   for  $v = 1, 2, 3, \dots, n - 1$ 
18     TotalLoad  $+= q_v$ 
19     if TotalLoad >  $Q^l$ 
20       return false
21 else // Test for LTCs
22   LoadDeliveredButNotPickedUpOnSubroute = array of integers of length  $n + 1$ , initialized to 0
23   if Problem type is not many-to-many
24     IndexOfLastDecouple = 0
25     for  $v = 1, 2, 3, \dots, n - 1$ 
26       if  $v$  is a decoupling vertex
27         IndexOfLastDecouple =  $v$ 
28       if Trailer is not attached upon leaving  $v$ 
29         if  $v$  is a delivery vertex and corresponding pickup  $v' < \text{IndexOfLastDecouple}$ 
30           LoadDeliveredButNotPickedUpOnSubroute[IndexOfLastDecouple]  $+= q_{v'}$ 
31   MinLorryLoadSinceLastDecouple = 0
32   MaxLorryLoad = 0
33   MaxLorryLoadSinceLastDecouple = 0
34   for  $v = 1, 2, 3, \dots, n - 1$ 
35     TotalLoad  $+= q_v$ 
36     if TotalLoad >  $Q^l + Q^t$ 
37       return false
38     MaxLorryLoad = min(TotalLoad,  $Q^l$ )
39     if  $v$  is a decoupling vertex
40       MinLorryLoadSinceLastDecouple =
41         max(TotalLoad -  $Q^t$ , LoadDeliveredButNotPickedUpOnSubroute[ $v$ ])
42       if MinLorryLoadSinceLastDecouple >  $Q^l$ 
43         return false
44       MinLorryLoadSinceLastDecouple = max(MinLorryLoadSinceLastDecouple, 0)
45       MaxLorryLoadSinceLastDecouple = MaxLorryLoad
46   if Trailer is not attached upon leaving  $v$ 
47     MinLorryLoadSinceLastDecouple = max(MinLorryLoadSinceLastDecouple +  $q_v$ , 0)
48     if MinLorryLoadSinceLastDecouple >  $Q^l$ 
49       return false
50     MaxLorryLoadSinceLastDecouple = min(MaxLorryLoadSinceLastDecouple +  $q_v$ ,  $Q^l$ )
51     if MaxLorryLoadSinceLastDecouple < 0
52       return false
53 return true

```

5. If i is on a subroute, $\text{LoadBalanceFromStartOfSubroute}[i]$ is the positive, negative or zero load balance from the start of the subroute up to and including i .
In the above example route, $\text{LoadBalanceFromStartOfSubroute}[9] = -5 = 20 - 30 + 10 + 15 - 10 - 10$.
6. $\text{OffsetOfNextCoupling}[i]$ stores the number of positions from i to the next index of a coupling process.
7. If i is on a subroute, $\text{MaxLoadBalanceFromStartOfSubroute}[i][\text{offset}]$ stores the maximum of zero and the largest load balance from the start of the subroute to any index from i up to and including $i + \text{offset}$.
In the example, $\text{MaxLoadBalanceFromStartOfSubroute}[5][3] = +15 = \max(0, -10, 0, +15, +5) = \max(0, \text{LoadBalanceFromStartOfSubroute}[7])$.

As an example of the simultaneous PDPT, consider the following five tasks:

Task	t_1	t_2	t_3	t_4	t_5
Original amount to be delivered	+5	+20	+10	+12	+16
Original amount to be picked up	+10	+15	+10	+18	+13
q_p^t	+5	+20	+10	+12	+16
q_d^t	+5	-5	0	+6	-3

Assume these five tasks are fulfilled by the following route:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	13
Task	—	t_1	t_2	t_3	t_4	t_5	t_5	—	t_1	t_3	t_2	—	t_4	—
Vertex type	s	P	P	P	P	P	D	Dec	D	D	D	Co	D	e
Artificial?	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No
Capacity requirement	0	+5	+20	+10	+12	+16	-3	0	+5	0	-5	0	+6	0

This route contains a subroute that starts at index 7 and ends at index 11. The load balances at indices 2 to 10 are +25, +35, +47, +63, +60, +60, +65, +65, and +60; thus, $\text{MaxTotalLoadOfSegment}[2][8] = +65$. Moreover, $\text{MaxTotalLoadOfSegment}[12][0] = +66$, $\text{LoadBalanceFromStartOfSubroute}[10] = 0$, and $\text{MaxLoadBalanceFromStartOfSubroute}[10] = +5 = \max(0, +5, +5, 0) = \text{LoadBalanceFromStartOfSubroute}[8] = \text{LoadBalanceFromStartOfSubroute}[9]$.

The decisive aspect for the feasibility test of simultaneous PDPTs is that the delivery amount is not necessarily the negative of the pickup amount. In Algorithm 2, this is taken into account in lines 6–8 and 32–36.

$\text{MaxTotalLoadOfSegment}$ and $\text{MaxLoadBalanceFromStartOfSubroute}$ can be initialized and updated with a nested loop, by iterating over all indices $j \geq i$ for each index i . The other data structures can be determined in a single loop. Hence, computing the preprocessing data is possible in quadratic time in the number of tasks on the route.

Algorithm 2 uses these data to test the capacity feasibility of an insertion of a task $t = (p, d)$ into an existing route r , with p to be inserted directly after position

(zero-based index of the route) h and d to be inserted directly after position i . It is evident that the running time of the algorithm is constant, independent of the number of vertices visited on the route.

Algorithm 2 returns false in line 4 if the total capacity is less than the maximal total load between the pickup and the delivery plus the pickup amount. For the simultaneous PDPT, false is returned in line 8 if the former condition is fulfilled or the total capacity is less than the maximal load between the delivery and the end of the route plus the pickup plus the delivery amount. The latter condition must be tested because, as explained in Sect. 3.5, for the simultaneous PDPT the sum of the pickup and the delivery amount indicates the change of load at deliveries.

The algorithm returns false in line 23 if p is to be inserted on a subroute and the minimal load that must be in the lorry when leaving the decoupling vertex plus the load balance from the start of the subroute up to h plus q_p plus the maximal load balance, counted from the beginning of the subroute, from $h + 1$ to i or to the end of the subroute, whichever vertex is visited earlier, exceeds the lorry capacity.

False is also returned if d is to be inserted on a subroute and the minimal load that must be in the lorry when leaving the decoupling vertex plus the maximal load balance from the start of the subroute up to i plus q_p exceeds the lorry capacity (line 31).

Moreover, for the simultaneous PDPT, similar to the situation on the main route, the algorithm returns false in line 36 if the minimal load aboard the lorry when leaving i plus q_d plus the maximal load balance from $i + 1$ to the end of the subroute exceeds the lorry capacity.

Finally, Algorithm 2 returns false in line 43 if d is to be inserted directly after p on a subroute and the total load upon leaving h plus q_p exceeds the lorry capacity.

If Algorithm 2 returns false from line 4, 8, or 23, further potential insertion positions for d need not be tested with the current insertion position of p . Instead, the next potential position for inserting p can be considered. Hence, it is sufficient to execute lines 2–23 of Algorithm 2 only once for each h .

Note also that lines 14–23 and 37–43 in Algorithm 2 are never reached for the simultaneous PDPT, as a pickup is never on a subroute.

4.2.2 The many-to-many PDPT

In the capacity test for the many-to-many PDPT, all data structures from Sect. 4.2.1 are employed, except for `LoadDeliveredButNotPickedUpOnSubroute`. In addition, the following data structures are used:

For testing the total vehicle capacity:

1. `OffsetsToWhereActualTotalLoadIsMaximal[i]` indicates, for each index i on the route, the non-negative offset from i to the index where the overall load in the vehicle, not taking into account what might have to be picked up at the depot, is maximal.
2. `MaxActualTotalLoad[i]` stores the maximal physical load from each index i until the end of the route, not taking into account any load that might have to be picked up at the depot. This load is aboard the vehicle at index $i + \text{OffsetsToWhereActualTotalLoadIsMaximal}[i]$.

Algorithm 2 TestCapacityConstantOneToOneAndOneToManyToOnePDPT (p, d, r, h, i, k)

Input: Pickup-and-delivery task $t = (p, d)$ with capacity requirements $q_p \geq 0$ and $q_d \in \mathbb{R}$
Route $r = (0, 1, 2, \dots, n)$
Indices of insertion positions h, i with $0 \leq h \leq i \leq n - 1$
Vehicle k (single lorry or LTC) with lorry and trailer capacities Q^l and Q^t ; for single lorries, $Q^t = 0$

Result: Returns **true** iff inserting p or a triple $\tilde{p} = p t l_p \rightarrow p \rightarrow p t l_p$ into r after index h and d or a triple $\tilde{d} = p t l_d \rightarrow d \rightarrow p t l_d$ after i (or, iff $h = i$, after p or $\tilde{p}$) is feasible regarding lorry and trailer capacity, **false** otherwise

```

1 // Evaluate feasibility of insertion regarding total capacity
2 if Problem type is not simultaneous
3   if  $Q^l + Q^t < \text{MaxTotalLoadOfSegment}[h][i - h] + q_p$ 
4     return false
5 else
6   if  $Q^l + Q^t < \max(\text{MaxTotalLoadOfSegment}[h][i - h] + q_p,$ 
7      $\text{MaxTotalLoadOfSegment}[i][n - i] + q_p + q_d)$ 
8     return false
9 if  $Q^t = 0$ 
10  return true
11 // Evaluate feasibility of insertion regarding lorry capacity
12 if  $i > h$  // Delivery not directly after pickup
13   // Evaluate feasibility of insertion of pickup
14   if TrailerAttached[ $h$ ] == false // Trailer not attached when leaving  $h$ 
15     ind = IndexOfLastPrecedingDecouple[ $h$ ]
16     MinLoadAtDecouple =  $\max(\text{MaxTotalLoadOfSegment}[\text{ind}][0] - Q^t,$ 
17        $\text{LoadDeliveredButNotPickedUpOnSubroute}[\text{ind}])$ 
18     LoadAfterPickup =  $\text{MinLoadAtDecouple} + \text{LoadBalanceFromStartOfSubroute}[h] + q_p$ 
19     offset =  $\text{OffsetOfNextCoupling}[h + 1]$ 
20     if  $h + \text{OffsetOfNextCoupling}[h] \geq i$ 
21       offset =  $i - h$ 
22     if  $\text{LoadAfterPickup} + \text{MaxLoadBalanceFromStartOfSubroute}[h + 1][\max(0, \text{offset} - 1)] > Q^l$ 
23       return false
24   // Evaluate feasibility of insertion of delivery
25   if TrailerAttached[ $i$ ] == false
26     if  $i - h \geq \text{OffsetOfNextCoupling}[h]$  // Delivery not on same subroute as pickup
27       ind = IndexOfLastPrecedingDecouple[ $i$ ]
28       MinLoadAtDecouple =  $\max(\text{MaxTotalLoadOfSegment}[\text{ind}][0] - Q^t,$ 
29          $\text{LoadDeliveredButNotPickedUpOnSubroute}[\text{ind}])$ 
30       if  $\text{MinLoadAtDecouple} + \text{MaxLoadBalanceFromStartOfSubroute}[\text{ind}][i - \text{ind}] + q_p > Q^l$ 
31         return false
32     if Problem type is simultaneous
33       offset =  $\text{OffsetOfNextCoupling}[i + 1]$ 
34       if  $\text{MinLoadAtDecouple} + q_p + \text{LoadBalanceFromStartOfSubroute}[i]$ 
35          $+ q_d + \text{MaxLoadBalanceFromStartOfSubroute}[i + 1][\text{offset}] > Q^t$ 
36         return false
37 else //  $i = h$ , i.e., delivery directly after pickup
38   if TrailerAttached[ $h$ ] == false
39     ind = IndexOfLastPrecedingDecouple[ $h$ ]
40     MinLoadAtDecouple =  $\max(\text{MaxTotalLoadOfSegment}[\text{ind}][0] - Q^t,$ 
41        $\text{LoadDeliveredButNotPickedUpOnSubroute}[\text{ind}])$ 
42     if  $\text{MinLoadAtDecouple} + \text{LoadBalanceFromStartOfSubroute}[i] + q_p > Q^l$ 
43       return false
44 return true

```

As an example, consider the following route:

Index	0	1	2	3	4	5	6	7	8	9
Capacity requirement	0	+50	+10	-70	+30	+25	-5	+15	-25	0

For this route, $\text{MaxActualTotalLoad}[3] = +65 = 0 + 30 + 25 - 5 + 15$, and $\text{OffsetsToWhereActualTotalLoadIsMaximal}[3] = 4$.

3. $\text{MostNegativeLoadBalanceFromStart}[i]$ indicates, for each index i , the most negative load balance (or zero if there are only pickups up to i) from the start of the route to i . The term ‘most negative’ is used instead of ‘smallest’ or ‘minimal’ to point out that all values are less than or equal to zero. A ‘minimal’ or ‘smallest’ element of an arbitrary set of numbers may be positive.

In the example, $\text{MostNegativeLoadBalanceFromStart}[2] = 0$, and $\text{MostNegativeLoadBalanceFromStart}[3] = -10$.

4. $\text{MostNegativeLoadBalancesFromPosToEndOfRoute}[i]$ specifies, for each index i , the most negative load balance (or zero if there are only pickups) from i until the end of the route.

In the example, for indices 3, 4, and 6, $\text{MostNegativeLoadBalancesFromPosToEndOfRoute}$ equals -70 , 0 , and -15 .

5. $\text{LoadBalanceFromTo}[i][\text{offset}]$ is the load balance from i to $i + \text{offset}$.

In the example, $\text{LoadBalanceFromTo}[2][2] = -30$, $\text{LoadBalanceFromTo}[3][4] = -5$, and $\text{LoadBalanceFromTo}[4][1] = +55$.

For testing lorry capacities on subroutes:

1. $\text{MostNegativeLoadBalanceOnSubrouteFromTo}[i][\text{offset}]$ is the most negative load balance from i until $i + \text{offset}$ on the subroute where i is visited, counted from the beginning of the subroute.

As an example, consider the following route:

Index	0	1	2	3	4	5	6	7	8	9
Capacity requirement	0	+50	0	-20	+30	-25	-5	0	-25	0

This route contains a subroute that starts at index 2 and ends at index 7. $\text{MostNegativeLoadBalanceOnSubrouteFromTo}[2][\text{offset}]$ equals zero for $\text{offset} = 0$ and -20 for $\text{offset} = 1, \dots, 4$. For index 4 and $\text{offset} = 0, 1$, and 2 , the values are 0 , -15 , and -20 .

2. $\text{MostNegativeLoadBalanceOnSubrouteFromAnyVertexBeforeUntilPos}[i]$ specifies the most negative load balance from any vertex that precedes i on the subroute, including i , up to i .

In the example, the values for indices 2–6 are 0 , -20 , 0 , -25 , -30 .

3. $\text{MostNegativeLoadBalanceFromIUntilEndOfSubroute}[i]$ indicates the most negative load balance on the segment from any index i until the end of the subroute, counted from i .

In the above example, the values of $\text{MostNegativeLoadBalanceFromIUntilEndOfSubroute}$, for indices 2–6, are -20 , -20 , 0 , -30 , and -5 .

`LoadBalanceFromTo` and `MostNegativeLoadBalanceOnSubrouteFromTo` can be filled using a nested forward pass, i.e., by iterating over all indices $j \geq i$ for each index i on the route; the other data structures can be filled or updated in one pass. As explained in Sect. 4.2, `MaxTotalLoadOfSegment[i][offset]` is the maximal load balance from the start of the route at any index from i up to and including $i + \text{offset}$. For the many-to-many problem where loading at the start depot is allowed, `MaxTotalLoadOfSegment` also takes into account the load that must be picked up at the depot, by adding this load to each component of `MaxTotalLoadOfSegment`. Computing the load to be picked up at the depot can be done in one additional forward pass over the route, which does not change the time complexity of determining the preprocessing data. Hence, the preprocessing data for a many-to-many-PDPT route can be computed in quadratic time in the number of tasks on the route. With these data structures, capacity feasibility can be tested as described in Algorithm 3. It is easy to see that the algorithm itself runs in constant time.

Algorithm 3 returns false in line 8 if the maximal load in the vehicle, disregarding any load brought from the depot, at any index from and including h until the end of the route (i.e., `MaxActualTotalLoad[h]`) plus q_p plus the rest of the load that still needs to be picked up at the depot even after inserting p exceeds the total vehicle capacity. To exemplify this, consider the following route and assume $h = 3$ and $q_p = 40$:

Index	0	1	2	3	4	5	6	7
Capacity requirement	0	-10	+10	+10	-30	+30	-100	0

It is easy to see that 90 units of load must be picked up at the depot for the route to be feasible, 10 units of which are needed at index 1, 10 at index 4, and 70 at index 6. The maximal load in the vehicle anywhere after h , not taking into account load brought from the depot, is 30 at index 5. Up to this point, 20 units of load picked up at the depot have been unloaded. Hence, the expression in line 5 equals $(-1) \cdot (-90) + (-20) = +70$. The value to be subtracted from q_p in line 6 is $(-1) \cdot (-90) + (-10) = +80$. This value is greater than q_p , so that the overall value of line 6 is negative, and this means that the entire amount q_p can be used to substitute load brought from the depot. This, in turn, means that inserting p does not increase the load at any index, which implies that the insertion is feasible. Now assume that the capacity requirement at index 6 is -50 . Then, 40 units must be picked up at the depot, and line 5 equals $(-1) \cdot (-40) + (-20) = +20$. In line 6, $+30$ must be subtracted from q_p , so that the value of line 6 is $+40 - 30 = +10$. This means that 30 units of q_p can be used to help satisfy the negative capacity requirements at indices 4 and 6, so that the 30 units picked up at the depot to this end are no longer necessary. This, in turn, means that inserting p after h increases the load from $h + 1$ onwards by 10 units, which may or may not violate the total vehicle capacity.

The return false in line 11 is self-explanatory.

The return false in line 19 has the following rationale: $q_p \leq 0$ implies that $|q_d| > 0$. To satisfy this negative capacity requirement, additional load can be picked up at the depot if this is allowed and if there is free capacity up to index i . In any case, for the insertion of d to be feasible, it must still be possible to satisfy all negative capacity

requirements after i , which, in other words, means that the maximal possible load after d must be at least as much as the absolute value of the most negative load balance computed from i until the end of the route.

The return false in line 35 can be explained as follows. The minimal load aboard the lorry when leaving a decoupling vertex must be the maximum of the following quantities: (i) the absolute value of the most negative load balance on the subroute up to and including h , (ii) the absolute value of the most negative load balance on the subroute from $h + 1$ until the end of the subroute minus q_p , and (iii) the total load in the vehicle when leaving the decoupling vertex minus the trailer capacity. Then, when p is inserted, the load after p is the sum of this minimal load, the load balance from the start of the subroute up to h , and q_p . If this sum exceeds the lorry capacity, the insertion is infeasible.

The algorithm returns false in line 40 if the absolute value of the most negative load balance on the subroute on which d is to be inserted, computed from any index before i up to i , plus $|q_d|$ exceeds the lorry capacity, as this quantity is a lower bound for the amount of load that must be aboard the lorry after decoupling when d shall be inserted after i .

Finally, Algorithm 3 returns false in line 51 when, after inserting d directly behind i , not enough load can be in the lorry to satisfy the most negative load balance on the subroute after i . This is tested as follows. The load that can be in the lorry at the decoupling vertex equals the load balance from the start of the route up to the decoupling vertex plus the maximal additional load that can be picked up at the depot. The load at i is the minimum of the lorry capacity and the load at the decoupling vertex plus the most positive load balance from the start of the subroute up to i plus the most negative load balance on the subroute from any vertex before up to i . The load after visiting d is then equal to the load at i plus q_d , and this load must be non-negative and greater than or equal to the absolute value of the most negative load balance from $i + 1$ to the end of the subroute.

Note that lines 21 ff. are valid whether or not loading at the start depot is allowed.

An alternative feasibility testing procedure for the many-to-many PDP without trailers is described by Kindervater and Savelsbergh (1997).

5 Conclusion and outlook

This paper has studied several variants of the PDPT, a routing problem where transport tasks between pickup and delivery locations must be fulfilled by a fleet of capacity-constrained single lorries and lorry-trailer combinations subject to accessibility restrictions at locations. Procedures to test the capacitive feasibility of inserting a task into an existing route have been presented. The procedures run in constant time, given appropriate auxiliary data that can be computed in a preprocessing step.

In their survey, Battarra et al. (2014) write (p. 181): ‘When looking at the pickup and delivery literature as a whole, one cannot fail to notice that there exists a very large number of problem variants which differ in their structure but nevertheless share many similarities. One can thus hope to see the development of general modelling and

Algorithm 3 TestCapacityConstantManyToManyPDPT(p, d, r, h, i, k)

Input: Pickup-and-delivery task (p, d) with capacity requirements $q_p \geq 0$ and $q_d \leq 0$
Route $r = (0, 1, 2, \dots, n)$
Indices of insertion positions h, i with $0 \leq h \leq i \leq n-1$
Vehicle k (single lorry or LTC) with lorry and trailer capacities Q^l and Q^t ; for single lorries, $Q^t = 0$

Result: Returns **true** iff inserting p or a triple $\tilde{p} = ptl_p \rightarrow p \rightarrow ptl_p$ into r after index h and d or a triple $\tilde{d} = ptl_d \rightarrow d \rightarrow ptl_d$ after i (or, iff $h = i$, after p or $\tilde{p}$) is feasible regarding lorry and trailer capacity, **false** otherwise

```

1 // Evaluate feasibility of insertion regarding total capacity
2 if  $q_p > 0$ 
3    $offset = \text{OffsetsToWhereActualTotalLoadIsMaximal}[h]$ 
4   if  $\text{MaxActualTotalLoad}[h]$ 
5      $+((-1) \cdot \text{MostNegativeLoadBalanceFromStart}[n] + \text{MostNegativeLoadBalanceFromStart}[h + offset])$ 
6      $+ (q_p - ((-1) \cdot \text{MostNegativeLoadBalanceFromStart}[n] + \text{MostNegativeLoadBalanceFromStart}[h]))$ 
7      $> Q^l + Q^t$ 
8     return false
9   if Unloading at end depot is not allowed
10    if  $\max(0, \text{LoadBalanceFromTo}[0][h]) + q_p + \text{LoadBalanceFromTo}[h+1][n - (h+1)] > 0$ 
11      return false
12 else
13    $\text{MaxAdditionalLoadThatCanBePickedUpAtDepot} = 0$ 
14   if Loading at start depot is allowed
15      $\text{MaxAdditionalLoadThatCanBePickedUpAtDepot} = Q^l + Q^t - \text{MaxTotalLoadOfSegment}[0][i]$ 
16    $\text{MaxTotalLoadAfterDelivery} =$ 
17      $\text{MaxTotalLoadOfSegment}[i][0] + \text{MaxAdditionalLoadThatCanBePickedUpAtDepot} + q_d$ 
18   if  $\text{MaxTotalLoadAfterDelivery} < \max(0, (-1) \cdot \text{MostNegativeLoadBalancesFromPosToEndOfRoute}[i+1])$ 
19     return false
20 // Evaluate feasibility of insertion regarding lorry capacity
21 if  $Q^t > 0$  and  $i > h$  // Vehicle is an LTC and delivery is not directly after pickup
22 // Evaluate feasibility of insertion of pickup
23 if  $q_p > 0$  and  $\text{TrailerAttached}[h] == \text{false}$ 
24    $ind = \text{IndexOfLastPrecedingDecouple}[h]$ 
25    $\text{MinLoadAtDecouple} = (-1) \cdot \text{MostNegativeLoadBalanceOnSubrouteFromTo}[ind][h - ind]$ 
26    $offset = \text{OffsetOfNextCoupling}[h+1]$ 
27    $\text{MinLoadAtDecouple} =$ 
28      $\max(\text{MinLoadAtDecouple}, (-1) \cdot \text{MostNegativeLoadBalanceOnSubrouteFromTo}[h+1][offset] - q_p)$ 
29    $\text{MinLoadAtDecouple} = \max(\text{MinLoadAtDecouple}, \text{MaxTotalLoadOfSegment}[ind][0] - Q^t)$ 
30    $\text{LoadAfterPickup} = \text{MinLoadAtDecouple} + \text{LoadBalanceFromStartOfSubroute}[h] + q_p$ 
31    $offset = \text{OffsetOfNextCoupling}[h+1]$ 
32   if  $h + \text{OffsetOfNextCoupling}[h] \geq i$ 
33     |  $offset = i - h$ 
34   if  $\text{LoadAfterPickup} + \text{MaxLoadBalanceFromStartOfSubroute}[h+1][\max(0, offset - 1)] > Q^l$ 
35     return false
36 // Evaluate feasibility of insertion of delivery
37 if  $q_d < 0$  and  $\text{TrailerAttached}[i] == \text{false}$ 
38    $ind = \text{IndexOfLastPrecedingDecouple}[i]$ 
39   if  $(-1) \cdot \text{MostNegativeLoadBalanceOnSubrouteFromAnyVertexBeforeUntilPos}[i] + (-1) \cdot q_d > Q^t$ 
40     return false
41    $\text{MaxAdditionalLoadThatCanBePickedUpAtDepot} = 0$ 
42   if Loading at start depot is allowed
43      $\text{MaxAdditionalLoadThatCanBePickedUpAtDepot} = Q^l + Q^t - \text{MaxTotalLoadOfSegment}[0][ind]$ 
44    $\text{PossibleLorryLoadAtDecouple} =$ 
45      $\text{MaxTotalLoadOfSegment}[ind][0] + \text{MaxAdditionalLoadThatCanBePickedUpAtDepot}$ 
46    $\text{LoadAtPredOfDelivery} =$ 
47      $\min(Q^l, \text{PossibleLorryLoadAtDecouple} + \text{MaxLoadBalanceFromStartOfSubroute}[ind][i - ind])$ 
48      $+ \text{MostNegativeLoadBalanceOnSubrouteFromAnyVertexBeforeUntilPos}[i]$ 
49    $\text{LoadAfterDelivery} = \text{LoadAtPredOfDelivery} + q_d$ 
50   if  $\text{LoadAfterDelivery} < \max(0, (-1) \cdot \text{MostNegativeLoadBalanceFromUntilEndOfSubroute}[i+1])$ 
51     return false
52 return true

```

solution techniques capable of handling multiple variants with a unified framework.’ The author hopes that the present paper constitutes a step in this direction.

In any case, there is ample opportunity for further research: throughout the paper at hand, it was assumed that the goods to be transported are homogeneous and load can be split arbitrarily between a lorry and its trailer. This is true for many practical situations, most notably, for the transport of liquids, bulk cargo, and palletized goods.

However, in all problem variants, the transported goods might just as well consist of a number of distinct items with different sizes, so that load splitting between lorry and trailer is possible only in discrete amounts. The transfer of load between a lorry and its trailer can even be completely forbidden or impossible for technical reasons. Taking this into account would require major changes to the procedures described above. It is thus beyond the scope of the present work and left for future research.

Acknowledgements Open Access funding provided by Projekt DEAL.

Compliance with ethical standards

Conflict of interest The author declares that he has no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Battarra M, Cordeau JF, Iori M (2014) Pickup-and-delivery problems for goods transportation. In: Toth P, Vigo D (eds) *Vehicle routing: problems, methods, and applications*. Society for industrial and applied mathematics, Philadelphia, pp 161–191
- Bürckert H, Fischer K, Vierke G (2000) Holonic transport scheduling with TELETRUCK. *Appl Artif Intell* 14:697–725
- Cheung R, Shi N, Powell W, Simão H (2008) An attribute-decision model for cross-border drayage problem. *Transp Res Part E* 44(2):217–234
- Cuda R, Guastaroba G, Speranza M (2015) A survey on two-echelon routing problems. *Comput Oper Res* 55:185–199
- Doerner K, Salazar-González J (2014) Pickup-and-delivery problems for people transportation. In: Vigo D, Toth P (eds) *Vehicle routing: problems, methods, and applications*. Society for industrial and applied mathematics, Philadelphia, pp 193–212
- Drexl M (2018) On the one-to-one pickup-and-delivery problem with time windows and trailers. Technical Report 1816, Gutenberg School of Management and Economics, Johannes Gutenberg University Mainz
- Firat M, Woeginger G (2011) Analysis of the dial-a-ride problem of Hunsaker and Savelsbergh. *Oper Res Lett* 39(1):32–35
- Funke B, Grünert T, Irnich S (2005) Local search for vehicle routing and scheduling problems: review and conceptual integration. *J Heuristics* 11(4):267–306
- Grangier P, Gendreau M, Lehuédé F, Rousseau LM (2016) An adaptive large neighborhood search for the two-echelon multiple-trip vehicle routing problem with satellite synchronization. *Eur J Oper Res* 254(1):80–91
- Gschwind T, Drexl M (2019) Adaptive large neighborhood search with a constant-time feasibility test for the dial-a-ride problem. *Transp Sci* 53(2):480–491
- Haugland D, Ho S (2010) Feasibility testing for dial-a-ride problems. In: Chen B (ed) *Algorithmic aspects in information and management. AAIM 2010. Lecture notes in computer science* 6124. Springer, Berlin, pp 170–179
- Hunsaker B, Savelsbergh M (2002) Efficient feasibility testing for dial-a-ride problems. *Oper Res Lett* 30(3):169–173

- Irnich S (2008a) Resource extension functions: properties, inversion, and generalization to segments. *OR Spectrum* 30(1):113–148
- Irnich S (2008b) A unified modeling and solution framework for vehicle routing and local search-based metaheuristics. *INFORMS J Comput* 20(2):270–287
- Irnich S, Funke B, Grünert T (2006) Sequential search and its application to vehicle-routing problems. *Comput Oper Res* 33(8):2405–2429
- Kindervater G, Savelsbergh M (1997) Vehicle routing: handling edge exchanges. In: Aarts E, Lenstra J (eds) *Local search in combinatorial optimization*. Wiley, Chichester, pp 337–360
- Masson R, Lehuédé F, Péton O (2013) Efficient feasibility testing for request insertion in the pickup and delivery problem with transfers. *Oper Res Lett* 41(3):211–215
- Parragh S, Cordeau JF (2017) Branch-and-price and adaptive large neighborhood search for the truck and trailer routing problem with time windows. *Comput Oper Res* 83:28–44
- Parragh S, Doerner K, Hartl R (2008a) A survey on pickup and delivery models part I: transportation between customers and depot. *J Betriebswirtschaft* 58(1):21–51
- Parragh S, Doerner K, Hartl R (2008b) A survey on pickup and delivery models part II: transportation between pickup and delivery locations. *J Betriebswirtschaft* 58(2):81–117
- Prodhon C, Prins C (2014) A survey of recent research on location-routing problems. *Eur J Oper Res* 238(1):1–17
- Ropke S, Pisinger D (2006) A Unified Heuristic for a Large Class of Vehicle Routing Problems with Backhauls. *Eur J Oper Res* 171(3):750–775
- Rothenbächer AK, Drexl M, Irnich S (2018) Branch-and-price-and-cut for the truck-and-trailer routing problem with time windows. *Transp Sci* 52(5):1174–1190
- Savelsbergh M (1985) Local search in routing problems with time windows. *Ann Oper Res* 4(1):285–305
- Savelsbergh M (1990) An efficient implementation of local search algorithms for constrained routing problems. *Eur J Oper Res* 47(1):75–85
- Savelsbergh M (1992) The vehicle routing problem with time windows: minimizing route duration. *ORSA J Comput* 4(2):146–154
- Tilk C, Bianchessi N, Drexl M, Irnich S, Meisel F (2018) Branch-and-price-and-cut for the active-passive vehicle-routing problem. *Transp Sci* 52(2):300–319
- Vidal T, Crainic TG, Gendreau M, Prins C (2014) A unified solution framework for multi-attribute vehicle routing problems. *Eur J Oper Res* 234(3):658–673
- Xue Z, Zhang C, Lin W, Miao L, Yang P (2014) A Tabu search heuristic for the local container drayage problem under a new operation mode. *Transp Res Part E* 62:136–150
- Yagiura M, Ibaraki T (1999) Analyses on the 2 and 3-flip neighborhoods for the MAX SAT. *J Comb Opt* 3(1):95–114