Gutenberg School of Management and Economics

& Research Unit "Interdisciplinary Public Policy"

Discussion Paper Series

Route Feasibility Testing and Forward Time Slack for the Synchronized Pickup and Delivery Problem

Timo Gschwind

May 2015

Discussion paper number 1503

Contact details

Timo Gschwind
Lehrstuhl für BWL insb. Logistikmanagement
Johannes-Gutenberg-Universität Mainz
Jakob-Welder-Weg 9
55128 Mainz

gschwind@uni-mainz.de

Route Feasibility Testing and Forward Time Slack for the Synchronized Pickup and Delivery Problem

Timo Gschwind[a]

[a]*Chair of Logistics Management, Gutenberg School of Management and Economics,*
Johannes Gutenberg University Mainz, Jakob-Welder-Weg 9, D-55128 Mainz, Germany.

Abstract

The temporal constraints of the Synchronized Pickup and Delivery Problem (SPDP) impose a complex scheduling problem for the service times at the customer locations. This makes the efficient feasibility checking of routes intricate. We present two different route feasibility checks for the SPDP and compare their practical runtime on a huge number of randomly generated routes. Furthermore, we generalize the concept of forward time slack, which has proven a versatile tool for feasibility testing of VRP variants, to the SPDP.

Key words: Vehicle routing, Temporal synchronization, Feasibility testing, Forward time slack

1. Introduction

The Synchronized Pickup and Delivery Problem (SPDP) (Gschwind, 2014) is the prototypical Vehicle Routing Problem (VRP) with *temporal intra-route synchronization* constraints. It seeks to find a set of minimum-cost routes servicing n user-specified transportation requests from origin (or pickup) to destination (or delivery) points subject to pairing and precedence, capacity, and time-window constraints. Moreover, the service times at the pickup and delivery locations of the customer requests are synchronized in the following way: After completing the service at a pickup point, the corresponding delivery has to be performed within prespecified *minimum* and *maximum time lags* (called *Ride Times*, RT). From a modeling point of view, the SPDP generalizes the Dial-A-Ride Problem (DARP, see Cordeau and Laporte, 2007, for a survey) in which no minimum RT are present.

Because of the complex temporal constraints of the SPDP, deciding whether or not a given route is feasible is a non-trivial task. The efficient feasibility testing of routes, however, is a crucial part in many exact and heuristic algorithms for VRPs. Sophisticated feasibility tests for a DARP with an additional constraint on the maximum waiting time at the customer nodes have been proposed in (Tang *et al.*, 2010) and (Firat and Woeginger, 2011). The time complexity of these approaches is $\mathcal{O}(n^2)$ and $\mathcal{O}(n)$, respectively. In the SPDP, the additional presence of minimum ride-time constraints further complicates the route feasibility problem.

Another crucial aspect for some solution approaches to VRP variants is the ability to quickly evaluate the feasibility of insertions of single nodes or requests into a given (feasible) route. The concept of *Forward Time Slack* (FTS) originally introduced in (Savelsbergh, 1992) for the VRP with Time Windows (VRPTW) can be a useful tool for this kind of evaluation. Generalized versions of the FTS have been used to asses the feasibility of insertions in heuristic algorithms, e.g., for the DARP (Cordeau and Laporte, 2003) and the Pickup and Delivery Problem with Transfers (PDPT, Masson *et al.*, 2013). In a companion paper (Gschwind, 2014), the FTS principle is used within a dynamic-programming labeling algorithm for the solution of the column-generation subproblem of the SPDP.

The contribution of this note is twofold. First, we derive two different route feasibility checks for the SPDP by adapting the approaches for the DARP presented in (Firat and Woeginger, 2011) and (Tang *et al.*,

2010) and conduct a computational study over a large number of randomly generated routes to compare the practical runtime of the two procedures. Our results indicate that the algorithm with the inferior worst-case complexity shows on average a better practical performance even on very long routes. Second, we generalize the concept of FTS to the SPDP and demonstrate why the definition of the FTS is not unique for problems with maximum RT constraints.

The remainder of the paper is structured as follows: Section 2 defines the SPDP. Section 3 derives two feasibility tests for individual routes and presents aggregated results comparing their practical runtime. The adaptation of the FTS concept to the SPDP is described in Section 4. Short conclusions are drawn in Section 5.

2. Problem Definition

The SPDP is defined on a complete digraph $G = (V, A)$ with node set V and arc set A. The node set V comprises the origin and destination depots 0 and $2n+1$, the set of pickup nodes $P = \{1, ..., n\}$, and the set of delivery nodes $D = \{n+1, ..., 2n\}$. For each node $i \in V$ a time window $[a_i, b_i]$, a service duration s_i, and a demand q_i with $q_i = -q_{i+n}$ are given. A travel time t_{ij} and a routing cost c_{ij} are associated with each arc $(i, j) \in A$. Both travel times and routing costs are assumed to satisfy the triangle inequality.

Each transportation request $i = 1, ..., n$ consists of transporting a specific good from the pickup node $i \in P$ to the delivery node $i+n \in D$. A minimum RT $\underline{L}_i$ and a maximum RT $\overline{L}_i$ are associated with each transportation request i.

A fleet K of homogeneous vehicles each with a capacity of Q is located at the origin depot 0 to serve the transportation requests. The task of the SPDP is to find a set of $|K|$ routes starting and ending at the depot nodes 0 and $2n+1$ such that each transportation request is performed exactly once. Thereby, vehicle capacities have to be respected and the service at each node has to be started within its time window. If a vehicle arrives prior to a_i at node i, it has to wait until the time window opens. Moreover, waiting, i.e., voluntarily delaying the start of service, is allowed at any node at any time. For each transportation request i, the pickup and delivery nodes have to be served on the same route and the pickup has to be serviced before. Furthermore, the service at the delivery node $i+n$ has to be started at least $\underline{L}_i$ and at most $\overline{L}_i$ units of time after the service at the pickup has been completed.

3. Route Feasibility Testing

Consider a given route $R = (h_1, ..., h_q)$ with $h_1 = 0$ and $h_q = 2n+1$. We use the notation $h_i \in R$ to indicate that a node h_i is part of route R. Furthermore, given a pickup (delivery) node $h_i \in P$ ($h_i \in D$), we indicate by h_{i-} and i^- (h_{i+} and i^+) the corresponding delivery (pickup) node and its associated notation.

Testing the feasibility of a route R means checking its consistency with all constraints of the SPDP that relate to individual routes. The verification of pairing and precedence and the capacity constraint is independent of each other and of the temporal constraints and can be done straightforwardly in linear time. We assume in the following that these constraints are respected. Checking whether or not the temporal constraints are satisfied is intricate. Indeed, one has to verify if there exists a time schedule $\mathcal{T}_R = (\tau_1, ..., \tau_q)$ satisfying

$$a_{h_i} \leq \tau_i \leq b_{h_i} \quad \forall h_i \in R, \tag{1}$$

$$\tau_i + s_{h_i} + t_{h_i h_{i+1}} \leq \tau_{i+1} \quad \forall h_i \in R, \tag{2}$$

$$\tau_i + s_{h_i} + \overline{L}_{h_i} \geq \tau_{i^-} \quad \forall h_i \in R \cap P, \tag{3}$$

$$\tau_i + s_{h_i} + \underline{L}_{h_i} \leq \tau_{i^-} \quad \forall h_i \in R \cap P, \tag{4}$$

where τ_i denotes the start of service at node h_i. This is called the *scheduling problem* of the SPDP. Constraints (1) are time-window constraints. Consistency of the service times along the route is ensured by constraints (2), while (3) and (4) are maximum and minimum RT constraints, respectively. For ease

of notation, we will assume service durations of zero and omit them in all formulae in the following. The extension of all concepts to non-zero service durations is straightforward.

In (Tang *et al.*, 2010) and (Firat and Woeginger, 2011), feasibility tests for a similar scheduling problem, i.e., with constraints (1)–(3) and an additional constraint on the waiting time at each node $h_i \in R$, were presented. In the following, we sketch both algorithms and extend them to solve the scheduling problem of the SPDP. Also, we highlight the increased complexity coming from the additional presence of the minimum RT constraints that leads to increased worst-case running times of the adapted procedures.

3.1. Adapted Feasibility Test of Firat and Woeginger

The basic idea of Firat and Woeginger (2011) is to rewrite the considered scheduling problem as a system of difference constraints. It is well-known (see, e.g., Section 24.4 in Cormen *et al.*, 2001) that such a system has a solution if and only if an associated digraph (called *constraint graph*) has no negative-weight cycle. Moreover, they formulate the difference-constraint system over an appropriate set of variables allowing the cycle-detection test to be performed in linear time.

Before we sketch their algorithm, some additional notation is necessary. For each node $h_i \in R$, the constant T_i denotes the sum of the travel times along the route R up to node h_i. The total waiting time up to node $h_i \in R$ is given by x_i.

Original Algorithm for the DARP. The first step of Firat and Woeginger is to rewrite constraints (1)–(3) of the considered scheduling problem in terms of T_i and x_i yielding

$$a_{h_i} \leq T_i + x_i \leq b_{h_i} \quad \forall h_i \in R, \tag{5}$$
$$x_i - x_{i+1} \leq 0 \quad \forall h_i \in R, \tag{6}$$
$$x_{i^-} - x_i \leq T_i + \overline{L}_{h_i} - T_{i^-} \quad \forall h_i \in R \cap P. \tag{7}$$

Next, they introduce two additional dummy variables x_0 and x_{q+1} representing values $x_0 = 0$ and $x_{q+1} = b_{2n+1} - a_0 = K$, i.e., they constitute lower and upper bounds for all $x_i, i = 1, ..., q$. Using x_0 and x_{q+1}, constraints (5) can be written as the difference constraints

$$x_{q+1} - x_i \leq K + T_i - a_{h_i} \quad \forall h_i \in R, \tag{8}$$
$$x_i - x_0 \leq b_{h_i} - T_i \quad \forall h_i \in R. \tag{9}$$

The final set of difference constraints is (6)–(9) together with

$$x_{q+1} - x_0 \leq K, \tag{10}$$
$$x_0 - x_{q+1} \leq -K, \tag{11}$$

to enforce $x_{q+1} - x_0 = K$. The constraint graph associated with this system has a node for each $x_i, i = 0, ..., q+1$ and an arc with weight d_{ij} from node x_j to node x_i for each constraint $x_i - x_j \leq d_{ij}$ of the system (6)–(11). It has $\mathcal{O}(n)$ nodes and $\mathcal{O}(n)$ arcs. Thus, negative-weight cycles can be detected in $\mathcal{O}(n^2)$ using the Bellman-Ford algorithm.

The key observation of Firat and Woeginger to obtain a linear-time feasibility check is the following: They call an arc from x_j to x_i a forward arc if $j < i$, otherwise it is a backward arc. All arcs in the constraint graph are either forward arcs with non-negative weights (otherwise the scheduling problem would be trivially infeasible) or backward arcs with weight zero. The only exception is the arc corresponding to constraint (11) which is a backward arc with negative weight. Using this structure, they are able to transform the graph into a specific interval graph for which the cycle-detection test can be done in linear time.

Adapted Algorithm for the SPDP. The adaptation of the approach to the SPDP requires to perform the same transformation described above also to constraints (4) yielding:

$$x_i - x_{i^-} \leq T_{i^-} - T_i - \underline{L}_{h_i} \quad \forall h_i \in R \cap P. \tag{12}$$

The additional difference constraints (12) correspond to additional arcs in the constraint graph. Still, the total number of arcs is $\mathcal{O}(n)$ and the application of the Bellman-Ford algorithm gives a $\mathcal{O}(n^2)$ feasibility test.

To show that feasibility testing in linear time is not possible for the SPDP (with the technique in Firat and Woeginger, 2011), we have to analyze the new arcs: They are backward arcs and have non-negative weight if $T_{i^-} - T_i \geq \underline{L}_{h_i}$, i.e., if the minimum RT of request h_i is trivially satisfied on the given route because the travel time from h_i to h_{i^-} is already larger. Otherwise, the arc weight is negative. As a result, the specific structure of the constraint graph that was exploited to obtain a linear time cycle-detection test is lost.

Summing up, the adaptation of the approach of Firat and Woeginger (2011) yields a $\mathcal{O}(n^2)$ feasibility test for the SPDP. However, the algorithm requires to build an extra network for each route which might negatively affect the practical runtime.

3.2. Adapted Feasibility Test of Tang et al.

Tang *et al.* (2010) proposed a different feasibility test for the same scheduling problem as in (Firat and Woeginger, 2011), but with weaker, quadratic worst-case runtime. However, their algorithm seems more intuitive as it gradually constructs the schedule $\mathcal{T}_R$ directly on the original network. This also means that some information that is needed might already be available within the exact or heuristic approach in which the feasibility test is necessary. Therefore, their algorithm might be sufficiently fast (or even faster than the algorithm in Firat and Woeginger, 2011) in practice, especially if the given routes are not too long (see, e.g., Gschwind and Irnich, 2014).

Again, we first sketch the original algorithm before we present our adaptations to the SPDP.

Original Algorithm for the DARP. The algorithm of Tang *et al.* tries to construct a feasible schedule $\mathcal{T}_R$ satisfying (1)–(3) by traversing the route twice: once forward and once backward. If no feasible schedule can be found the route is infeasible. Note that the third traversal of the original algorithm is redundant.

Algorithm 1: Algorithm of Tang *et al.*

Result: **true** if feasible schedule exists, **false** otherwise

```
// Pass 1 (forward)
```
1   $\tau_1 := a_{h_1}$
2   **for** $i = 2, ..., q$ **do**
3     $\tau_i := \max\left\{\tau_{i-1} + t_{h_i h_{i-1}}, a_{h_i}\right\}$
4   **if** $\tau_i > b_{h_i}$ **then return** false

```
// Pass 2 (backward)
```
5   **for** $i = q-1, ..., 1$ **do**
6   **if** $h_i \in P$ **then**
7       $\Delta := \tau_{i^-} - \tau_i - \overline{L}_{h_i}$
8   **if** $\Delta > 0$ **then**
9         $\tau_i := \tau_i + \Delta$
10   **if** $\tau_i > b_{h_i}$ **then return** false
11   **for** $k = i+1, ..., q$ **do**
12          $\tau_k := \max\left\{\tau_{k-1} + t_{h_k h_{k-1}}, a_{h_k}\right\}$
13   **if** $\tau_k > b_{h_k}$ **then return** false
14   **if** $\tau_{i^-} - \tau_i - \overline{L}_{h_i} > 0$ **then return** false

15 **return** true

The whole procedure is described in Algorithm 1. The forward pass (Steps 1–4) builds a schedule of service times τ_i that satisfy constraints (1) and (2). Thereby, all times are scheduled as early as possible.

The backward pass (Steps 5–14) checks for consistency with the maximum RT constraints (3) adjusting some values τ_i if necessary: At each pickup node $h_i \in P$ it is checked if the current schedule satisfies the maximum RT of request h_i. If not, the algorithm tries to shift waiting time that occurs between pickup and delivery of request h_i before the pickup node h_i in order to decrease the ride time of request h_i. Thereto, the service at node h_i is delayed by as much as necessary to meet the maximum RT. This requires shifting the service times τ_k of all succeeding nodes $h_k, k = i+1, ..., q$ forward in time and is done in the same fashion as in the first pass, i.e., accounting for constraints (1) and (2). If there was not enough waiting time between pickup and delivery of request h_i, the maximum RT of h_i cannot be satisfied and the route is infeasible (Step 14). Note that this adjustment of the service times can never increase the RT of a request that is picked up later than h_i meaning that the maximum RT constraint of such a request never gets violated. Thus, after traversing a pickup node h_i in the backward pass we have a schedule that respects constraints (1) and (2), and the maximum RT constraints (3) of those requests for which the pickup node is not before h_i in the route. Moreover, all service times are still scheduled as early as possible with respect to the constraints that are satisfied at that point of the algorithm. Consequently, whenever a rescheduling results in a service time $\tau_j > b_{h_j}$ no feasible schedule exists and the route is infeasible.

Both forward and backward pass traverse the route once. Because of the inner forward loop (Steps 11–13), the backward pass has a quadratic worst-case running time and, hence, the overall algorithm also has time complexity $\mathcal{O}(n^2)$.

Adapted Algorithm for the SPDP. Solving the scheduling problem of the SPDP with the technique of Tang *et al.* (2010) requires the integration of the minimum RT constraints into the scheduling process. This can be done as follows: In the forward pass, change Step 3 to

$$\tau_i := \begin{cases} \max\left\{\tau_{i-1} + t_{h_i h_{i-1}}, a_{h_i}, \tau_{i^+} + \underline{L}_{h_{i^+}}\right\} & \text{if } h_i \in D, \\ \max\left\{\tau_{i-1} + t_{h_i h_{i-1}}, a_{h_i}\right\} & \text{otherwise.} \end{cases} \tag{13}$$

The resulting service times after the first pass satisfy constraints (1), (2), and (4) and they are scheduled in an early-as-possible fashion.

In the backward pass, whenever a shifting of service times is necessary because of some violated maximum RT (Steps 8–14), we need to change Step 12 in the same way as Step 3 in order to maintain feasibility with respect to constraints (1), (2), and (4). In contrast to the original algorithm, however, the shifting of waiting times can increase the RT of requests that are picked up later in the route. Decisive is that we might be forced to re-introduce waiting time somewhere in the route due to the minimum RT constraints of other requests. As a result, the property that after traversing a pickup node h_i in the backward pass all maximum RTs of requests which are not picked before h_i are respected, is lost.

Consider the example given in Table 1. The travel times between all nodes are assumed to be 10. The maximum RTs of requests i and m are 30. The minimum RT of request j is 31. The time window of each node is specified in Table 1. All other constraints are assumed to be never binding.

Table 1 gives the service times at each node at different stages of the algorithm. Waiting times at nodes are in brackets. After the first pass of the algorithm, the maximum RT constraint of request i is violated by 1 unit of time. Therefore, the second pass delays the service τ_i at i by 1 unit trying to shift waiting time that occurs between i and $i + n$ before i. However, the removed waiting time at node k decreases the RT of request j so that it has to be re-introduced before starting the service at node $j + n$. Otherwise, the minimum RT constraint of j would be violated. As a consequence, the RT of request m increases and the schedule after the backward pass is infeasible which, however, does not imply that the route is infeasible.

A straightforward way to fix this defect of the algorithm is to loop over the second pass as long as an adjustment of the service times was necessary because of a violated maximum RT. For the example above, one additional backward pass identifies a feasible schedule (see Table 1). Clearly, this change of the algorithm leads to a significant deterioration of the worst-case complexity, especially if the number of iterations of the backward pass cannot be bound.

5

Node		0	i	j	k	i^-	m	j^-	k^-	m^-	$2n{+}1$
$[a., b.]$		$[0, 100]$	$[10, 20]$	$[20, 30]$	$[31, 41]$	$[41, 51]$	$[51, 61]$	$[61, 71]$	$[71, 81]$	$[81, 91]$	$[0, 100]$
After pass 1	$\tau.$	0	10	20	(1) 31	41	51	61	71	81	91
Pass 2, at node i	$\tau.$	0	(1) 11	21	31	41	51	(1) 62	72	82	92
After add. pass 2	$\tau.$	0	(1) 11	21	31	41	(1) 52	62	72	82	92

Table 1: Example showing the necessity to loop over second pass in Algorithm 1

Node		0	i	j	m	i^-	j^-	m^-	$2n{+}1$
$[a., b.]$		$[0, 100]$	$[10, 20]$	$[20, 30]$	$[31, 41]$	$[41, 51]$	$[51, 61]$	$[61, 71]$	$[0, 100]$
After pass 1	$\tau.$	0	10	20	(1) 31	41	51	61	71
Pass 2, at node i	$\tau.$	0	(1) 11	21	31	41	(1) 52	62	72
2nd pass 2, at node m	$\tau.$	0	(1) 11	21	(1) 32	42	52	62	72
2nd pass 2, at node i	$\tau.$	0	(2) 12	22	32	42	(1) 53	63	73

Table 2: Example showing nested requests causing a worst-case complexity greater than $\mathcal{O}\left(n^2\right)$

Consider the example as given in Table 2. The time windows of the nodes are specified in the table. All other data is equivalent to the previous example. Apparently, the three requests i, j, and m are nested in a way that the adapted Algorithm 1 alternatively shifts waiting times that occur in between pickup and delivery of the requests i and m. In general, this can lead to an unbounded number of iterations and, hence, a superlinear runtime of the algorithm. In case of integer inputs for the temporal data, e.g., the total number of iterations that can be caused by this behavior can be bounded by $M = \max_{i \in P \cup D}(b_i - a_i)$. The runtime of the algorithm is then given by $\mathcal{O}\left(Mn^2\right)$. Overall, the worst-case complexity of the adapted algorithm seems to be prohibitively large for frequent calls. However, in the computational tests of a companion paper (Gschwind, 2014), we found that the described algorithm can solve the scheduling problems arising within a proposed branch-and-cut-and-price approach to the SPDP sufficiently fast. Note that the paths arising in the instances considered in (Gschwind, 2014) are typically very short (not more than 10–15 requests). We, therefore, conduct a more comprehensive computational study with much longer routes in the following.

3.3. Computational Comparison

To compare the practical performance of the approaches of Sections 3.1 and 3.2, we evaluate their runtime over a huge number of randomly generated paths with lengths reaching from 15 to 200 requests. A separate analysis is conducted for integer and double precision input parameters. Details on the generation of the paths and more detailed results can be found in Section A of the appendix. Note that in our analysis we do not include instances that are identified as infeasible already in the first pass of the adapted Algorithm 1.

Table 3 summarizes our results. Each row aggregates over a total of 50 000 paths. *Firat* and *Tang* denote the adapted feasibility checks of Firat and Woeginger and Tang *et al.*, respectively. Table 3 reports over all instances the average, maximum, and minimum times in milliseconds needed by the algorithms to solve an instance 10 000 consecutive times. Furthermore, we give the average, maximum, and minimum number of restarts of the second pass of the adapted Algorithm 1.

Table 3 reveals that there are paths, for which a huge number of iterations of the second pass of the adapted Algorithm 1 is necessary, especially for instances with double precision parameters. For these instances the runtime of *Tang* is very long. However, this seems to happen only very rarely so that the average number of iterations is small. Moreover and in contrast to the theoretical worst-case complexity, the average runtime of *Tang* is significantly better than that of *Firat*.

4. Forward Time Slack

The concept of FTS was originally introduced by Savelsbergh (1992) in the context of the VRPTW. Let $\mathcal{T}_R = (\tau_1, ..., \tau_q)$ be a feasible schedule for route R. Savelsbergh defines the FTS F_i for a node $h_i \in R$ as

Param.	Algo.	Solution time			# Restarts		
		avg.	max	min	avg.	max	min
integer	*Firat*	2 160.3	19 081	15	-	-	-
	Tang	95.6	21 590	<1	3.1	420	0
double	*Firat*	1 937.4	21 684	15	-	-	-
	Tang	111.6	176 016	<1	6.0	12 425	0

Table 3: Aggregated computational results.

the maximum value by which the start of service τ_i at h_i can be increased without causing the route to become infeasible. Different authors have generalized the concept of FTS to other problems, e.g., Cordeau and Laporte (2003) for the DARP or Masson *et al.* (2013) for the PDPT.

VRPTW. The slack at a node $h_j, j \geq i$ with respect to node h_i is given by the cumulative waiting time between h_i and h_j and the difference between the end of the time window and the start of service at node h_j. The FTS F_i at node h_i is the minimum of all slacks at nodes h_j with $i \leq j \leq q$. Denote by W_i the waiting time at node $h_i \in R$. Then (Savelsbergh, 1992):

$$F_i = \min_{i \leq k \leq q} \left\{ \sum_{i < p \leq k} W_p + (b_{h_k} - \tau_k) \right\}. \tag{14}$$

Note that this definition of F_i is independent of the service times $\tau_k, k < i$, i.e., it is unique with respect to the route segment preceding h_i. This can be an important property for feasibility testing when inserting pairs of nodes (as, e.g., in Masson *et al.*, 2013, for the PDPT).

DARP. For problems with RTs, shifting the start of service at some node may cause infeasibilities also with respect to the RT constraints. This has to be incorporated in the definition of the FTS for such problems. For example, postponing the service at node h_i may increase the RT of a request h_k with $k < i$ and $k^- > i$. Consequently, Cordeau and Laporte (2003) define the following generalization of the FTS for the DARP:

$$F_i = \min_{i \leq k \leq q} \left\{ \sum_{i < p \leq k} W_p + \min \left(b_{h_k} - \tau_k, S_{ik} \right) \right\}, \tag{15}$$

where S_{ik} is given by $\overline{L}_{h_{k^+}} - (\tau_k - \tau_{k^+})$ if $h_k \in D; i > k^+$ and $+\infty$ otherwise.

In contrast to (14), the FTS (15) of a node h_i for the DARP (and similarly for other problems with maximum RTs) is not unique with respect to the route segment preceding h_i. Indeed, it may be possible for some $h_k \in D$ with $k^+ < i$ and $k > i$ to increase τ_{k^+} such that τ_i stays constant while S_{ik} and F_i also increase. For that reason, feasibility checking of request insertions based on FTS (as, e.g., proposed in Masson *et al.*, 2013) is problematic.

Consider the route given in Table 4. Feasibility of the insertion of another request k is to be evaluated. Assume a travel time of ten between all nodes, $\overline{L}_i = 40$, and time window $[0, 20]$ at the pickup node k. Inserting node k before i and node k^- before j^- results in a feasible route. When using the FTS-values of the earliest-as-possible schedule, however, the insertion appears to be infeasible because τ_{j^-} is increased by $10 > F_{j^-} = 0$. Considering FTS-values that allow delaying τ_i to time 20 correctly determines feasibility of this route. On the other hand, using these FTS-values results in misjudging the insertion of k before j and k^- before j^- as feasible because $F_{j^-} = 10$ relies on $\tau_i = 20$ which is not possible in the new route.

SPDP. The presence of minimum RT raises two additional issues compared to the DARP. First, shifting the service time at a pickup node is constrained by the minimum RT of the request. Second, not necessarily all of the cumulative waiting time between two nodes h_i and h_k is slack. Decisive is that some waiting time may have to be included between them because of the minimum RT of some requests. Denote by UW_{ik} the

Node		0	i	j	j^-	i^-	$2n+1$
$[a., b.]$		$[0, 100]$	$[10, 30]$	$[30, 50]$	$[40, 60]$	$[50, 70]$	$[0, 100]$
Earliest as possible schedule	$\tau.$	0	10	(10) 30	40	50	60
FTS for nodes after h_i with $\tau_i = 10$	$F.$	–	–	0	0	0	40
FTS for nodes after h_i with $\tau_i = 20$	$F.$	–	–	10	10	10	40

Table 4: Example instance for which feasibility checking of request insertions based on FTS is problematic

usable waiting time between h_i and h_k, i.e., the waiting time between h_i and h_k that can be incorporated into the slack. Then, the FTS for a node h_i is given by

$$F_i = \min_{i \leq k \leq q} \{UW_{ik} + \min(b_{h_k} - \tau_k, S_{ik})\}, \tag{16}$$

with

$$S_{ik} = \begin{cases} b_{h_{k^-}} - (\underline{L}_{h_k} + \tau_k) & \text{if } h_k \in P, \\ \overline{L}_{h_k} - (\tau_k - \tau_{k^+}) & \text{if } h_k \in D; i > k^+, \\ +\infty & \text{otherwise.} \end{cases} \tag{17}$$

To determine UW_{ik}, we first define the *needed waiting time* $NW_k = (\underline{L}_k - \sum_{k < j \leq k^-} t_{j-1,j})^+$ for request h_k which is the amount of waiting time that inevitably occurs between pickup and delivery of request $h_k \in R$ and, hence, can never be part of the slack between nodes h_i and h_{k^-} if $k > i$. When there is another request h_j with positive NW_j in between h_i and h_k, care has to be taken not to consider the overlapping needed waiting times NW_j and NW_k twice in the definition of UW_{ik}. Therefore, we define for each delivery node h_j after h_i the already *included waiting time* between them as

$$IW_{ij} = \sum_{\substack{i < p < j \\ h_p \in D, p^+ \geq i, j^+ < p}} (NW_{p^+} - IW_{ip}). \tag{18}$$

Then, the usable waiting time UW_{ik} is given by

$$UW_{ik} = \sum_{i < p \leq k} W_p - \sum_{\substack{i < p \leq k \\ h_p \in D, p^+ \geq i}} (NW_{p^+} - IW_{ip})^+. \tag{19}$$

Note that a direct adaptation of the technique of Masson *et al.* (2013) using the FTS (16) does not lead to a valid feasibility check of request insertions for the SPDP because of the non-uniqueness of the FTS as described above. Still, the technique can be used to obtain a constant-time verifiable, sufficient condition. A feasibility check of request insertions for the SPDP and likewise for other problems with maximum RT that is more efficient than checking the resulting routes from scratch, e.g., with the algorithms of Section 3, is an open question.

5. Conclusions

In this note, we derived two route feasibility checks for the SPDP and extended the concept of FTS to this problem. Because of the prototypical character of the SPDP, we believe that the presented concepts are also relevant for other routing problems with temporal synchronization constraints. For example, the feasibility test of Section 3.1 can be adapted to a SPDP with transfer possibilities between routes in a straightforward way. Likewise, the proposed definition of the FTS can be generalized to the problem with transfer possibilities using the technique in (Masson *et al.*, 2013).

n	Cust. loc.	Horizon	TW	Min RT	Max RT
15	$[10, 10] \times [10, 10]$	450	60	2	10
25	$[10, 10] \times [10, 10]$	750	100	2	10
50	$[10, 10] \times [10, 10]$	1500	200	2	10
100	$[10, 10] \times [10, 10]$	3000	400	2	15
200	$[10, 10] \times [10, 10]$	6000	500	2	30

Table 5: Details on the parameters used for instance generation

References

Cordeau, J.-F. and Laporte, G. (2003). A tabu search heuristic for the static multi-vehicle dial-a-ride problem. *Transportation Research Part B*, **37**(6), 579–594.

Cordeau, J.-F. and Laporte, G. (2007). The dial-a-ride problem: models and algorithms. *Annals of Operations Research*, **153**(1), 29–46.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2001). *Introduction To Algorithms*. MIT Press.

Firat, M. and Woeginger, G. J. (2011). Analysis of the dial-a-ride problem of Hunsaker and Savelsbergh. *Operations Research Letters*, **39**(1), 32–35.

Gschwind, T. (2014). A comparison of column-generation approaches to the synchronized pickup and delivery problem. Technical Report LM-2014-01, Chair of Logistics Management, Gutenberg School of Management and Economics, Johannes Gutenberg University Mainz, Mainz, Germany.

Gschwind, T. and Irnich, S. (2014). Effective handling of dynamic time windows and its application to solving the dial-a-ride problem. *Transportation Science*.

Masson, R., Lehuédé, F., and Péton, O. (2013). Efficient feasibility testing for request insertion in the pickup and delivery problem with transfers. *Operations Research Letters*, **41**(3), 211–215.

Savelsbergh, M. W. P. (1992). The vehicle routing problem with time windows: Minimizing route duration. *INFORMS Journal on Computing*, **4**(2), 146–154.

Tang, J., Kong, Y., Lau, H., and Ip, A. W. H. (2010). A note on "efficient feasibility testing for dial-a-ride problems". *Operations Research Letters*, **38**(5), 405–407.

Appendix

A. Detailed Computational Results

In this section, we give additional details on the computational comparison of the two feasibility checks of Sections 3.1 and 3.2. The algorithms were tested on a huge number of randomly generated paths using Euclidean distances (rounded down in the case of instances with integer inputs) as travel times between the nodes. Minimum and maximum RTs are specified proportional to the Euclidean distances between pickup and delivery location. Details on the parameters used for the generation of the paths are given in Table 5. The columns have the following meaning:

n	Number of requests
Cust. loc.	Intervals from which the customer locations are drawn. The depots are located at $(0, 0)$.
Horizon	Time horizon
TW	Mean value for the length of the time windows
Min RT	Mean value for the factor specifying the minimum RT
Max RT	Mean value for the factor specifying the maximum RT

For each value of n we generate paths with different characteristics regarding the number of requests that are open at the nodes in the following way: At each node another request is picked up with probability p, otherwise one of the open requests is delivered. The considered values of p reach from 0.15 to 0.50 in steps of 0.05. For each probability p, a pretest filters 625 feasible and 625 infeasible instances so that a total of 10 000 paths is considered for each value of n. Note that we consider only instances that are not identified as infeasible already in the first pass of the adapted Algorithm 1.

Table 6 summarizes our results. The columns have the following meaning:

| | | Max # open | | | Firat | | | Tang | | | | | |
| | | | | | Solution time | | | Solution time | | | # Restarts | | |
Param.	n	avg.	max	min	avg.	max	min	avg.	max	min	avg.	max	min
integer	15	2.7	6	1	57.2	146	15	5.8	115	<1	1.0	45	0
	25	3.3	8	1	137.9	312	31	11.4	319	<1	1.6	76	0
	50	4.2	11	2	513.9	1 046	78	33.0	1 810	<1	2.2	162	0
	100	5.5	19	2	2 077.7	4 477	281	125.0	10 936	15	4.4	364	0
	200	6.8	23	2	8 014.9	19 081	726	302.7	21 590	31	6.4	420	0
double	15	2.7	6	1	49.8	125	15	8.9	15 459	<1	3.3	7 596	0
	25	3.3	6	1	115.8	312	31	18.2	14 492	<1	4.2	3 628	0
	50	4.2	11	1	438.4	1 123	78	37.8	15 116	<1	4.1	4 979	0
	100	5.5	19	2	1 787.5	4 556	203	161.6	140 994	<1	9.0	12 425	0
	200	6.8	23	2	7 295.4	21 684	624	331.5	176 016	15	9.7	9 548	0

Table 6: Detailed Computational Results

Param.	Type of input data (integer or double precision) of the instances
n	Number of requests
Max # open	The maximum number of open requests at a node; we give average, maximum, and minimum values over the instances
Firat	The adapted feasibility check of Firat and Woeginger (Section 3.1)
Tang	The adapted feasibility check of Tang *et al.* (Section 3.2)
Solution time	The time in milliseconds needed by the respective algorithm to solve an instance 10 000 consecutive times; we present average, maximum, and minimum values over the instances
# Restarts	The number of restarts of the second pass of Algorithm 1