

The general pickup and delivery problem

Citation for published version (APA):

Sol, M., & Savelsbergh, M. W. P. (1992). *The general pickup and delivery problem*. (Memorandum COSOR; Vol. 9244). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/1992

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

244

ARC
01
COS

EINDHOVEN UNIVERSITY OF TECHNOLOGY
Department of Mathematics and Computing Science

Memorandum COSOR 92-44

**The General Pickup and
Delivery Problem**

M. Sol
M.W.P. Savelsbergh

Eindhoven, November 1992
The Netherlands

TU/e

BSA

Eindhoven University of Technology
Department of Mathematics and Computing Science
Probability theory, statistics, operations research and systems theory
P.O. Box 513
5600 MB Eindhoven - The Netherlands

Secretariate: Dommelbuilding 0.03
Telephone: 040-47 3130

ISSN 0926 4493

The General Pickup and Delivery Problem

M. Sol

M.W.P. Savelsbergh

Department of Mathematics and Computer Science

Eindhoven University of Technology

P.O. Box 513

5600 MB Eindhoven

The Netherlands

Abstract

In pickup and delivery problems vehicles have to transport loads from origins to destinations without transshipment at intermediate locations. In this paper, we discuss several characteristics that distinguish them from standard vehicle routing problems and present a survey of the problem types and solution methods found in the literature.

1 Introduction

In the *General Pickup and Delivery Problem* (GPDP) a set of routes has to be constructed in order to satisfy transportation requests. A fleet of vehicles is available to operate the routes. Each vehicle has a given capacity, a start location and an end location. Each transportation request specifies the size of the load to be transported, the locations where it is to be picked up (the origins) and the locations where it is to be delivered (the destinations). Each load has to be transported by one vehicle from its set of origins to its set of destinations without any transshipment at other locations.

Three well-known and extensively studied routing problems are special cases of the GPDP. In the *Pickup and Delivery Problem* (PDP), each transportation request specifies a single origin and a single destination and all vehicles depart from and return to a central depot. The *Dial-a-Ride Problem* (DARP) is a PDP in which the loads to be transported represent people. Therefore, we usually speak of clients or customers instead of transportation requests and all load sizes are equal to one. The *Vehicle Routing Problem* (VRP) is a PDP in which either all the origins or all the destinations are located at the depot.

The GPDP is introduced in order to be able to deal with various complicating characteristics found in many practical pickup and delivery problems, such as transportation requests specifying a set of origins associated with a single destination or a single origin associated with a set of destinations, vehicles with different start and end locations, and transportation requests evolving in real time.

Many practical pickup and delivery situations are demand responsive, i.e., new transportation requests become available in real-time and are immediately eligible for consideration. As a consequence, the set of routes has to be reoptimized at some point to include the new transportation requests. Observe that at the time of the reoptimization, vehicles are on the road and the notion of depots becomes void.

The purpose of this paper is to present a general model that can handle the practical complexities mentioned above, to isolate and discuss some of the characteristics that differentiate pickup and delivery problems from traditional vehicle routing problems, and to give an overview of the literature on pickup and delivery problems.

With respect to the literature on routing and scheduling problems, it is interesting to observe that although pickup and delivery problems are as important from a practical point of view and as interesting from a theoretical point of view as vehicle routing problems, they have received far less attention. We hope that by identifying and discussing important issues regarding pickup and delivery problems and by presenting a survey of the literature, we stimulate further research in the area.

2 Problem Formulation

Let N be the set of transportation requests. For each transportation request $i \in N$, a load of size $\bar{q}_i \in \mathbb{N}$ has to be transported from a set of origins N_i^+ to a set of destinations N_i^- . Each load is subdivided as follows: $\bar{q}_i = \sum_{j \in N_i^+} q_j = -\sum_{j \in N_i^-} q_j$, i.e., positive quantities for pickups and negative quantities for deliveries. Define $N^+ := \cup_{i \in N} N_i^+$ as the set of all origins and $N^- := \cup_{i \in N} N_i^-$ as the set of all destinations. Let $V := N^+ \cup N^-$. Furthermore, let M be the set of vehicles. Each vehicle $k \in M$ has a capacity $Q_k \in \mathbb{N}$, a start location k^+ , and an end location k^- . Define $M^+ := \{k^+ | k \in M\}$ as the set of start locations and

$M^- := \{k^- | k \in M\}$ as the set of end locations. Let $W := M^+ \cup M^-$.

For all $i, j \in V \cup W$ let $d_{i,j}$ denote the travel distance, $t_{i,j}$ the travel time, and $c_{i,j}$ the travel cost.

Definition 1 A pickup and delivery route R_k for vehicle k is a directed route through a subset $V_k \subset V$ such that:

1. R_k starts in k^+ .
2. $(N_i^+ \cup N_i^-) \cap V_k = \emptyset$ or $(N_i^+ \cup N_i^-) \cap V_k = N_i^+ \cup N_i^-$ for all $i \in N$.
3. If $N_i^+ \cup N_i^- \subseteq V_k$, then all locations in N_i^+ are visited before all locations in N_i^- .
4. Vehicle k visits each location in V_k exactly once.
5. The vehicle load never exceeds Q_k .
6. R_k ends in k^- .

Definition 2 A pickup and delivery plan is a set of routes $\mathcal{R} := \{R_k | k \in M\}$ such that:

1. R_k is a pickup and delivery route for vehicle k , for each $k \in M$.
2. $\{V_k | k \in M\}$ is a partition of V .

Define $f(\mathcal{R})$ as the cost of plan \mathcal{R} corresponding to a certain objective function f . We can now define the general pickup and delivery problem as the problem:

$$\min\{f(\mathcal{R}) | \mathcal{R} \text{ is a pickup and delivery plan.}\}$$

The special cases of the GPDP mentioned in the introduction can be characterized as follows:

The pickup and delivery problem: $|W| = 1$ and $|N_i^+| = |N_i^-| = 1$ for all $i \in N$. In this case we define i^+ as the unique element of N_i^+ and i^- as the unique element of N_i^- .

The dial-a-ride problem: $|W| = 1$ and $|N_i^+| = |N_i^-| = 1$ and $\bar{q}_i = 1$ for all $i \in N$.

The vehicle routing problem: $|W| = 1$, $|N_i^+| = |N_i^-| = 1$ for all $i \in N$, and $N^+ \subseteq W$.

To formulate the GPDP as a mathematical program, we introduce four types of variables: z_i^k ($i \in N, k \in M$) equal to 1 if transportation request i is assigned to vehicle k and 0 otherwise,

x_{ij}^k ($((i, j) \in (V \times V) \cup \{(k^+, j) | j \in V\} \cup \{(j, k^-) | j \in V\}, k \in M)$ equal to 1 if vehicle k travels from location i to location j and 0 otherwise,

D_i ($i \in V \cup W$), specifying the departure time at vertex i , and

y_i ($i \in V \cup W$), specifying the load of the vehicle arriving at vertex i .

Define $q_{k^+} = 0$ for all $k \in M$. The problem is now to

minimize $f(x)$

subject to

$$\begin{aligned}
\sum_{k \in M} z_i^k &= 1 && \text{for all } i \in N, && (1) \\
\sum_{j \in V \cup W} x_{i,j}^k &= \sum_{j \in V \cup W} x_{j,i}^k = z_i^k && \text{for all } i \in N, l \in N_i^+ \cup N_i^-, k \in M, && (2) \\
\sum_{l \in N_i^+ \cup N_i^-} \sum_{j \in V \cup W} x_{i,j}^k &= |N_i^+ \cup N_i^-| z_i^k && \text{for all } i \in N, k \in M, && (3) \\
\sum_{j \in V \cup \{k-\}} x_{k+,j}^k &= 1 && \text{for all } k \in M, && (4) \\
\sum_{i \in V \cup \{k+\}} x_{i,k-}^k &= 1 && \text{for all } k \in M, && (5) \\
D_{k+} &= 0 && \text{for all } k \in M, && (6) \\
D_p &\leq D_q && \text{for all } i \in N, p \in N_i^+, q \in N_i^-, && (7) \\
x_{i,j}^k = 1 &\Rightarrow D_i + t_{i,j} \leq D_j && \text{for all } i, j \in V \cup W, k \in M, && (8) \\
y_{k+} &= 0 && \text{for all } k \in M, && (9) \\
y_l &\leq \sum_{k \in M} Q_k z_i^k && \text{for all } i \in N, l \in N_i^+ \cup N_i^-, && (10) \\
x_{i,j}^k = 1 &\Rightarrow y_i + q_i = y_j && \text{for all } i, j \in V \cup W, k \in M, && (11) \\
x_{i,j}^k &\in \{0, 1\} && \text{for all } i, j \in V \cup W, k \in M, && (12) \\
z_i^k &\in \{0, 1\} && \text{for all } i \in N, k \in M, && (13) \\
D_i &\geq 0 && \text{for all } i \in V \cup W, && (14) \\
y_i &\geq 0 && \text{for all } i \in V \cup W. && (15)
\end{aligned}$$

Constraint (1) ensures that each transportation request is assigned to exactly one vehicle. By constraint (2) a vehicle only enters or leaves a location l if it is an origin or a destination of a transportation request assigned to that vehicle. By constraint (3) the origins and destinations of a transportation request are only visited by that vehicle the transportation request is assigned to. Constraints (4) and (5) make sure that each vehicle starts and ends at the correct place. Constraints (6), (7), (8) and (14) together form the precedence constraints. Constraints (9), (10), (11) and (15) together form the capacity constraints.

3 Problem Characteristics

3.1 Transportation requests

A very important characteristic of routing problems is the way in which transportation requests become available. In a static situation, all requests are known at the time the routes have to be constructed. In a dynamic situation, some of the requests are known at the time the routes have to be constructed and the other requests become available in real time during execution of the routes. Consequently, in a dynamic situation, when a new transportation request becomes available, at least one route has to be changed in order to serve this new request. Most vehicle routing problems are static, whereas most pickup and delivery problems are dynamic.

In practice, a dynamic problem is often solved as a sequence of static problems. Each time a new request becomes available and the current set of routes has to be updated, only the known transportation requests are considered. In doing so, one ignores any information that may be known about the spatial or time distribution of future requests.

Another important concept in routing problems is that of a depot. In the routing literature, the depot is usually the place where vehicles start and end their routes. Since most pickup and delivery problems are dynamic, often with a long planning horizon, the concept of a depot vanishes. Drivers sleep at the last location they visited or at the first location

they have to visit the next day. Even for problems with a short planning horizon, such as a single day, where vehicles start and end at a central depot, a demand responsive situation leads to problems without depots. When new transportation requests become available and the current set of routes has to be updated the vehicles are spread out over the planning area.

The general pickup and delivery model is well suited for dealing with the subproblems that occur in dynamic demand responsive routing problems.

3.2 Time constraints

Apart from the vehicle capacity constraints and the intrinsic precedence constraints related to pickup and delivery, side constraints related to time arise in almost every practical pickup and delivery situation. Although time constraints have become an integral part of models for vehicle routing problems (for recent surveys on the vehicle routing problem with time windows see Desrochers, Lenstra, Savelsbergh, and Soumis [8] and Solomon and Desrosiers [38]), they play an even more prominent role in pickup and delivery problems. Among other reasons, because the most studied pickup and delivery problem is the dial-a-ride problem, which deals with the transportation of people who specify desired pickup or delivery times.

The presence of time constraints complicates the problem considerably. If there are no time constraints, finding a feasible pickup and delivery plan is trivial: arbitrarily assign transportation requests to vehicles, arbitrarily order the transportation requests assigned to a vehicle and process each transportation request separately. In the presence of time constraints the problem of finding a feasible pickup and delivery plan is \mathcal{NP} -hard. Consequently, it may be difficult to construct a feasible plan, especially when time constraints are restrictive. On the other hand, an optimization method may benefit from the presence of time constraints, since the solution space may be much smaller.

3.2.1 Time constraints related to transportation requests

For each $i \in V \cup W$ a time window $[e_i, l_i]$ is introduced denoting the time interval in which service at location i must take place. Given a pickup and delivery plan and departure times of the vehicles, the time windows define for each $i \in V$ the arrival time A_i and the departure time D_i . Note that $D_i = \max\{A_i, e_i\}$. If $A_i < e_i$, then the vehicle has to wait at location i .

Next to the explicit time windows mentioned above there also exist implicit time windows. Implicit time windows originate from controlling customer inconvenience. In dial-a-ride systems, people usually specify either a desired delivery time \bar{A}_i^- or a desired pickup time \bar{D}_i^+ . Because people do not want to be late at their destination, in a pickup and delivery plan the actual delivery time A_i^- must satisfy $A_i^- \leq \bar{A}_i^-$. Analogously, we require $D_i^+ \geq \bar{D}_i^+$. This defines half open time windows. To prevent clients from being served long before (after) their desired delivery (pickup) time, we can either construct closed time windows or take an objective function that penalizes deviations from the desired service time. Closed time windows can, for example, be constructed by defining a maximum deviation from the desired pickup or delivery time.

Another problem in which customer inconvenience can be controlled by time windows emerges from a demand responsive immediate request dial-a-ride system. In such a system, clients that request service want to be picked up as soon as possible. If client i has a pickup location that is geographically far away, this client will be scheduled last on a vehicle route. This situation will not change even when new clients request service. To prevent the client

from suffering indefinite deferment, a closed pickup time window $[0, l_{i+}]$ is defined where l_{i+} is an input to the system.

Other time constraints that originate from customer inconvenience restrictions in dial-a-ride systems are maximum ride time restrictions, i.e., a bound on the time a client is in the vehicle, and deadhead restrictions, i.e., no waiting time is allowed when a client is in the vehicle. Deadhead restrictions introduce the concept of schedule blocks, i.e., working periods between two successive slack periods.

3.2.2 Time constraints related to vehicles

Usually, vehicles are not available all day. Drivers have to eat and sleep and vehicles are subjected to service plans. These constraints can be modeled as time windows for vehicles. Typically a vehicle has multiple time windows defining all the periods in which it is available.

3.3 Objective functions

A wide variety of objective functions is found in pickup and delivery problems. The most common ones are discussed below.

First, we present objective functions related to single-vehicle pickup and delivery problems.

Minimize duration. The duration of a route is the total time a vehicle needs to execute the route. Route duration includes travel times, waiting times, loading and unloading times, and break times.

Minimize completion time. The completion time of a route is the time service at the last location is completed. In case the start time of the vehicle is fixed at time zero, the completion time coincides with the route duration.

Minimize travel time. The travel time of a route refers to the total time spent on actual traveling between different locations.

Minimize client inconvenience. In dial-a-ride systems, client inconvenience is measured in terms of pickup time deviation, i.e., the difference between the actual pickup time and the desired pickup time, delivery time deviation, i.e., the difference between the desired delivery time and the actual delivery time, and excess ride time, i.e. the difference between the realized ride time and the direct ride time. In demand responsive situations where clients request immediate service, i.e., as soon as possible, the difference between the time of pickup and the time of request placement may also contribute to the definition of client inconvenience. Different kinds of functions, linear as well as nonlinear, have been proposed to model client inconvenience.

Second, we present objective functions related to multiple-vehicle pickup and delivery problems.

Minimize the number of vehicles. This function is almost always used in dial-a-ride systems, combined with one of the above functions to optimize the single-vehicle subproblems. Dial-a-

ride systems are normally highly subsidized systems for the transportation of the elderly and handicapped. Therefore the objective is to minimize cost (mostly together with customer inconvenience). Because drivers and vehicles are the most expensive parts in a dial-a-ride system, minimizing the number of vehicles to serve all requests is usually the main objective.

Maximize profit. This function, which can use all of the above functions, can be used in a system where the dispatcher has the possibility of rejecting a transportation request when it is unfavorable to transport the corresponding load. Note that, for example, in a dial-a-ride system, it is not allowed to reject a transportation request. A model based on this objective function should not only incorporate the costs, but also the revenues associated with the transportation of loads.

For dynamic pickup and delivery problems it is not clear what kind of objective functions should be used. In a demand responsive environment pickup and delivery routes may be open ended. Therefore, objectives such as duration, completion time, and travel time have no clear meaning. Intuitively, an objective function for dynamic problems should emphasize decisions that affect the near future more than decisions regarding the remote future.

4 Solution approaches

In this section, we review the literature on pickup and delivery problems. The survey is organized as follows. The class of pickup and delivery problems has been divided into static and dynamic problems, since their characteristics as well as their solution approaches differ considerably. Within either of these classes, we distinguish single-vehicle and multiple-vehicle problems. Obviously, in the single-vehicle PDP all transportation requests are handled by the same vehicle, whereas in the multiple-vehicle PDP the transportation requests have to be divided over the set of vehicles. Assigning transportation requests to vehicles in the PDP is much more difficult than assigning transportation requests to vehicles in the VRP. In the VRP, all the origins of transportation requests are located at the depot. Therefore, transportation requests with geographically close destinations are likely to be served by the same vehicle. In the PDP, geographically close destinations may have origins that are geographically far apart and we cannot conclude that they are likely to be served by the same vehicle.

For each of the resulting subclasses, one or more papers are discussed. The level of detail depends on the originality, viability and importance of the described solution approach. Furthermore, we do not cover iterative improvement methods in any detail. For these the interested reader is referred to Kindervater and Savelsbergh [22].

4.1 The static pickup and delivery problem

4.1.1 The static single-vehicle pickup and delivery problem

The static single-vehicle pickup and delivery problem is probably the most studied variant of the PDP. First of all because the dial-a-ride problem belongs to this class, but also because it appears as a subproblem in multiple-vehicle pickup and delivery problems. We discuss solution approaches for problems with and without time windows separately.

(A) The static 1-PDP without time windows

Optimization

Psaraftis [27] considers immediate request dial-a-ride problems. In these problems, every client requesting service wishes to be served as soon as possible. The objective is to minimize a weighted combination of the time needed to serve all clients and the total degree of 'dissatisfaction' clients experience until their delivery. Dissatisfaction is assumed to be a linear function of the time each client waits to be picked up and of the time he spends riding in the vehicle until his delivery. This leads to the following objective function:

$$\min w_1 T + w_2 \sum_{i \in N} (\alpha W_i + (2 - \alpha) R_i),$$

where T denotes the route length, W_i the waiting time of client i from the departure time of the vehicle until the time of pickup, R_i the riding time of client i from pickup to delivery, and $0 \leq \alpha \leq 2$.

The problem is solved using a straightforward dynamic programming algorithm. The state space consists of vectors $(L, k_1, k_2, \dots, k_n)$, where L denotes the location currently being visited ($L = 0$ at the starting location, $L = i$ at the origin of client i and $L = i + n$ at the destination of client i) and k_i denotes the status of client i ($k_i = 3$ if client i has not been picked up, $k_i = 2$ if client i has been picked up but has not been delivered and $k_i = 1$ if client i has been delivered). Starting with state vector $(0, 3, 3, \dots, 3)$ the algorithm explores new states preserving feasibility of the partial routes constructed so far. The algorithm has a complexity of $O(n^2 3^n)$, where $n = |N|$, and can solve problems with up to 9 or 10 clients, i.e., 18 or 20 locations.

Kalantari, Hill and Arora [21] propose a method based on the branch and bound algorithm for the TSP developed by Little, Murty, Sweeney and Karel [23]. The method handles precedence constraints by precluding in each branch all the arcs that violate the active precedence constraints.

Fischetti and Toth [17] develop an additive bounding procedure that can be used in a branch-and-bound algorithm for the TSP with precedence constraints, i.e., a single-vehicle dial-a-ride problem without capacity constraints. The idea behind additive bounding is to use several bounds for a problem type additively rather than separately. Let $\mathcal{L}^{(k)}$ ($k = 1, \dots, r$) be a procedure which, when applied to problem P with cost vector c , returns a lower bound $\delta^{(k)}$ and a residual cost vector $c^{(k)}$, such that $\delta^{(k)} + c^{(k)}x \leq cx$ for all feasible x . Apply $\mathcal{L}^{(1)}$ to problem P with cost vector c and subsequently apply $\mathcal{L}^{(k)}$ to problem P with cost vector $c^{(k-1)}$ ($k = 2, \dots, r$). Then $\sum_{k=1}^r \delta^{(k)} + c^{(k)}x \leq cx$ for all feasible x . So $\sum_{k=1}^r \delta^{(k)}$ is a lower bound for P .

The bounds used by Fischetti and Toth for the TSP with precedence constraints are based on the assignment relaxation, the shortest spanning 1-arborescence relaxation, disjunctions, and on variable decomposition, in that order.

Approximation

Stein [40] presents a probabilistic analysis of a simple approximation algorithm for the single-vehicle dial-a-ride problem without capacity constraints. The algorithm constructs a TSP

tour through all origins and a TSP tour through all the destinations and then concatenates them.

Stein shows that if n origin-destination pairs are randomly chosen in a region of area a , using a uniform distribution, there exists a constant b such that the length y'_n of the tour constructed by the algorithm satisfies

$$\lim_{n \rightarrow \infty} \frac{y'_n}{\sqrt{n}} = 2b\sqrt{a} \text{ with probability 1,}$$

where b is the constant that appears in the theorem of Beardwood, Halton and Hammersley [2] on the length of an optimal traveling salesman tour in the Euclidean plane. Recent experiments by Johnson [18] indicate that $b \approx 0.713$. Stein also proves that if n origin-destination pairs are randomly chosen in a region of area a , using a uniform distribution, the length y_n^* of the optimal tour satisfies

$$\lim_{n \rightarrow \infty} \frac{y_n^*}{\sqrt{n}} = 1.89b\sqrt{a} \text{ with probability 1.}$$

This implies that the algorithm has an asymptotic performance bound of 1.06 with probability 1.

Psaraftis [30] presents a worst-case analysis of a simple two-phase approximation algorithm for the single-vehicle dial-a-ride problem in the plane. In the first phase, a TSP tour through all locations is constructed. In the second phase, this tour is traversed clock-wise, beginning at the start location of the vehicle and skipping any location that has been visited before, any origin where the pickup of a client would violate the capacity constraints, and any destination whose origin has not yet been visited, until a feasible dial-a-ride solution has been obtained.

Psaraftis shows that if there are no capacity constraints, the TSP tour has to be traversed at most twice. Furthermore, if in addition to the absence of capacity constraints the triangle inequality holds and Christofides' heuristic is used to construct the TSP tour, the algorithm has a worst case ratio of 3.

Fiala Timlin and Pulleyblank [16] consider a slightly different problem. There are groups of customers, each group with its own priority. Customers have to be served in order of increasing priority. Fiala Timlin and Pulleyblank develop an algorithm based on insertion techniques and iterative improvement methods.

(B) The static 1-PDP with time windows

Optimization

Psaraftis [28] modifies the dynamic programming algorithm discussed above to solve the static single-vehicle dial-a-ride problems with time windows. The major difference between the new and the original algorithm is the use of forward instead of backward recursion. Time windows are handled by the elimination of time infeasible states. The new algorithm still has a complexity of $O(n^2 3^n)$.

Desrosiers, Dumas and Soumis [10] also present a dynamic programming approach. Their algorithm uses states (S, i) with $S \subseteq V$ and $i \in V$. State (S, i) is defined only if there exists a feasible path that visits all nodes in S and ends in i . Elaborate state elimination criteria,

based not only on S , but also on the state (S, i) , are used to reduce the state space. These elimination criteria are very effective when the time windows are tight and the vehicle capacity is small. The algorithm has been tested successfully on problems with up to 40 transportation requests. When capacity constraints are rather tight, then, despite the fact that a dynamic programming algorithm is exponential, the running time of this algorithm seems to increase only linearly with problem size.

Approximation

Sexton and Bodin [35, 36] consider the dial-a-ride problem with desired delivery times specified by the clients. The objective is to minimize client inconvenience, which is defined as a weighted combination of delivery time deviation and excess ride time.

The presented solution approach applies Benders decomposition to a mixed 0-1 nonlinear programming formulation, which separates the routing and scheduling component.

The concept of space-time separation indicators between tasks plays an important role in the algorithm. Such an indicator measures the travel time between locations at which the tasks are performed, i.e., their spatial separation, and the difference between the latest feasible times at which they can be performed, i.e., their temporal separation.

Let \bar{A}_{k-} be the desired delivery time of client k . The latest possible pickup time \bar{D}_{k+} is then defined as $\bar{D}_{k+} := \bar{A}_{k-} - t_{k+,k-}$. The space-time separation indicators between tasks i and j are now defined as follows:

$$\sigma_{i+j+} = t_{i+j+} + \bar{D}_{j+} - \bar{D}_{i+}$$

$$\sigma_{i+j-} = t_{i+j-} + \bar{A}_{j-} - \bar{D}_{i+}$$

$$\sigma_{i-j+} = t_{i-j+} + \bar{D}_{j+} - \bar{A}_{i-}$$

$$\sigma_{i-j-} = t_{i-j-} + \bar{A}_{j-} - \bar{A}_{i-}$$

Van der Bruggen, Lenstra and Schuur [5] develop a solution procedure based on the Lin-Kernighan variable-depth local search method for the TSP. They use the techniques introduced by Savelsbergh [33, 34] to handle time windows and precedence constraints efficiently. The algorithm has a construction and an improvement phase. An initial route is obtained by visiting the locations in order of increasing centers of their time window, i.e., $(e_i + l_i)/2$ for location i , taking precedence and capacity constraints into account. The resulting route may be time infeasible. Using an objective function measuring total infeasibility, the route is made feasible by iterative improvement methods. The same procedures, with a different objective function, are then applied to find a better feasible route.

4.1.2 The static multiple-vehicle pickup and delivery problem

(A) The static m-PDP without time windows

Approximation

Cullen, Jarvis and Ratliff [6] propose an interactive approach for the multiple-vehicle dial-a-ride problem with a homogeneous fleet, i.e., equal vehicle capacities. The problem is decomposed in a clustering part and a chaining part. Both parts are solved in an interactive setting,

i.e., man and machine cooperate to obtain high quality solutions. The algorithmic approach in both parts is based on set partitioning and column generation.

A cluster consists of a seed arc and a set of clients assigned to this seed arc. The total number of clients assigned to a seed arc may not exceed the vehicle capacity Q . Let (u^+, u^-) denote the seed arc of the cluster and let S denote the set of clients assigned to this seed arc. The cost c of serving this cluster is approximated by $c = 2 \sum_{i \in S} d_{u^+, i^+} + d_{u^+, u^-} + 2 \sum_{i \in S} d_{u^-, i^-}$, i.e., it is assumed that a vehicle serving a cluster starts in u^+ , makes a round trip to each pickup location in the cluster, travels to u^- and makes a round trip to each delivery location in the cluster.

The clustering problem, i.e., the problem of constructing and selecting clusters to serve all the clients, can be formulated as a set partitioning problem. Let J be the set of all possible clusters, i.e., seed arcs and assignments of clients to seed arcs. For each $j \in J$, let c_j denote the approximate cost of serving the cluster, and for each $i \in N, j \in J$ let a_{ij} be a binary constant indicating whether client i is a member of cluster j or not. Furthermore, introduce a binary decision variable y_j to indicate whether a cluster is selected or not. The clustering problem is now to

$$\begin{aligned} & \text{minimize} && \sum_{j \in J} c_j y_j \\ & \text{subject to} && \sum_{j \in J} a_{ij} y_j = 1 && \text{for all } i \in N, \\ & && y_j \in \{0, 1\} && \text{for all } j \in J. \end{aligned}$$

Because the set of all possible clusters is extremely large, a column generation scheme is used to solve the linear programming relaxation of this set partitioning problem.

The master problem is initialized with all columns corresponding to clusters consisting of a single client. The row prices $(\pi_1, \pi_2, \dots, \pi_n)$ are computed and used to define a subproblem to generate columns with negative reduced costs, i.e., clusters that correspond to attractive columns for the master problem. The master problem now heuristically tries to improve the current solution by using (some of) the new columns. Then new row prices are calculated and the subproblem is solved again.

The subproblem that has to be solved is a location-allocation problem. Let $c_{ij} = 2d_{u_j^+, i^+} + 2d_{u_j^-, i^-}$ and $f_j = d_{u_j^+, u_j^-}$. The subproblem has the following mathematical programming formulation:

$$\begin{aligned} & \text{minimize} && \sum_{j \in J} \sum_{i \in N} (c_{ij} - \pi_i) z_{ij} + \sum_{j \in J} f_j x_j \\ & \text{subject to} && \sum_{i \in N} z_{i,j} \leq Q x_j && \text{for all } j \in J, \\ & && \sum_{j \in J} z_{i,j} \leq 1 && \text{for all } i \in N, \\ & && z_{i,j} \in \{0, 1\} && \text{for all } i \in N, j \in J, \\ & && x_j \in \{0, 1\} && \text{for all } j \in J. \end{aligned}$$

Because of the size of J , the subproblem is still very difficult to solve. Therefore, it is solved approximately: Choose a set of clients to form the seed arcs. With these seed arcs fixed, solve the resulting assignment problem. With these assignments fixed, solve the resulting location problem. Continue alternating between the assignment problem and the location problem for

some specified number of iterations, or until no further improvement in the objective function is found.

The clusters in the solution to the clustering problem, together with some promising clusters that did not appear in the final solution, are input for the chaining problem. In the chaining problem a subset of these clusters, partitioning the set of clients, is selected and linked to form pickup and delivery routes. In the set partitioning formulation of this problem the rows correspond to clients again, but columns now correspond to vehicle routes. The column generation procedure links a subset of the seed arcs of the selected clusters. Each linking of seed arcs is then translated into a column for the set partitioning problem by placing a 1 in row i if client i is part of one of the clusters in the linking.

(B) The static m-PDP with time windows

Optimization

Dumas, Desrosiers and Soumis [15] present a set partitioning formulation for the static pickup and delivery problem with time windows and a column generation scheme to solve it to optimality. The approach is very robust in the sense that it can be adapted easily to handle different objective functions and variants with multiple depots and an inhomogeneous fleet of vehicles.

Let Ω be the set of all feasible pickup and delivery routes. For each route $r \in \Omega$, we denote the cost of the route by c_r . For all $i \in N$, $r \in \Omega$ let a_{ir} be a binary constant indicating whether transportation request i is served on route r ($a_{ir} = 1$) or not ($a_{ir} = 0$). Furthermore, introduce a binary variable indicating whether route r is used in the optimal solution ($x_r = 1$) or not ($x_r = 0$). The static PDP can now be formulated as follows:

$$\begin{aligned} & \text{minimize} && \sum_{r \in \Omega} c_r x_r \\ & \text{subject to} && \sum_{r \in \Omega} a_{ir} x_r = 1 && \text{for all } i \in N, \\ & && \sum_{r \in \Omega} x_r = |M|, \\ & && x_r \in \{0, 1\} && \text{for all } r \in \Omega. \end{aligned}$$

The cardinality of Ω is much too large to allow an exhaustive enumeration. Therefore, a column generation method is used to solve the linear relaxation of the set partitioning problem. Columns, defining admissible routes, are generated as needed by solving a constrained shortest path problem on a perturbed distance matrix. The perturbation of the distance matrix depends on the dual variables of the set partitioning problem. The lower bound found in this way is usually excellent and constitutes a good starting point for a branch and bound approach to obtain an integral solution. The constrained shortest path problem is solved by dynamic programming. The dynamic programming algorithm employed is essentially the same as the one in Desrosiers, Dumas, and Soumis [10] (see Section 4.1.1), except for the fact that not all transportation requests have to be processed.

Approximation

Dumas, Desrosiers and Soumis [14] develop an approximation algorithm for the dial-a-ride problem based on their optimization algorithm discussed above. The basic idea is to create route segments for small groups of clients, called *mini-clusters*. A minicluster is a segment of a route starting and ending with an empty vehicle. Each minicluster is then treated as a

transportation request that entirely fills a vehicle. The optimization algorithm is now applied to this set of transportation requests. This reduces the number of rows in the setpartitioning matrix. The subproblem is now much easier to solve because each transportation request corresponds to a full truck load. Note that in the approach of Cullen, Jarvis and Ratliff [6] clusters cannot be identified with rows because they do not partition the set of clients.

Miniclusters are constructed simply by taking a known pickup and delivery plan, constructed using any existing algorithm, and cutting it into pieces, such that each piece starts and ends with an empty vehicle. In this setup, the approach can best be viewed as an improvement method.

An alternative method to construct mini-clusters is presented by Desrosiers, Dumas, Soumis, Taillefer and Villeneuve [11]. A parallel insertion heuristic creates mini-clusters. The insertion criteria are based on the concept of *neighboring requests*. Two requests are considered to be neighbors if they satisfy certain temporal and spatial restrictions. To be more precise, two requests i and j are considered to be neighbors if the time intervals $[e_{i+}, l_{i-}]$ and $[e_{j+}, l_{j-}]$ overlap, if $t_{i+;+} + t_{j+;-} \leq \alpha t_{i+;-}$ or $t_{j+;+} + t_{i+;-} \leq \alpha t_{j+;-}$ for some constant α , if the angle between the arcs (i^+, i^-) and (j^+, j^-) is less than some constant β , and if the difference in cost between serving i and j together and serving i and j separately is larger than some constant γ .

Jaw, Odoni, Psaraftis and Wilson [19, 20] present two different approaches to the multiple-vehicle dial-a-ride problem, in which clients that are to be picked up and delivered have the following types of service constraints: Each client i specifies either a desired pickup time \bar{D}_{i+} or a desired delivery time \bar{A}_{i-} , and a maximum ride time \bar{T}_i . The objective is to minimize a combination of customer dissatisfaction and resource usage.

Jaw, Odoni, Psaraftis and Wilson [19] present a three-phase algorithm. The first phase, the grouping phase, decomposes the problem by dividing the time horizon into intervals and then assigning clients to groups according to the time interval into which their desired pickup or delivery time falls. The time intervals are chosen in such a way that it is possible to fully serve a client in two consecutive time intervals. The second phase, the clustering phase, partitions the set of clients in each time group into clusters and assigns a vehicle to each cluster. The number of clusters created is equal to the number of available vehicles. The third phase, the routing phase, constructs a route for each vehicle using a standard approximation algorithm for the single-vehicle dial-a-ride problem.

Jaw, Odoni, Psaraftis and Wilson [20] present an insertion algorithm. Time windows on both the pickup time and delivery time of a client are defined based on a prescribed tolerance U and the specified desired pickup or delivery time. If client i has specified a desired pickup time, the actual pickup time D_{i+} should fall within the time window $[\bar{D}_{i+}, \bar{D}_{i+} + U]$; if he has specified a desired delivery time, the actual delivery time A_{i-} should fall within the window $[\bar{A}_{i-} - U, \bar{A}_{i-}]$. Moreover, his actual travel time should not exceed his maximum ride time: $A_{i-} - D_{i+} \leq \bar{T}_i$. Note that this information suffices to determine two time windows $[e_{i+}, l_{i+}]$ and $[e_{i-}, l_{i-}]$ for each customer i . Finally, waiting time is not allowed when the vehicle is carrying passengers.

The selection criterion is simple: customers are selected for insertion in order of increasing e_{i+} . The insertion criterion is as follows: among all feasible points of insertion of the customer into the vehicle schedules, choose the cheapest; if no feasible point exists, introduce an additional vehicle.

For the identification of feasible insertions, the notion of an *active period* is introduced. This is a period of time a vehicle is active between two successive periods of slack time. For convenience, we drop the superscript indicating pickup or delivery. For each visit to an address i during an active period, we define the following possible forward and backward shifts:

$$P_i^- = \min\{\min_{j \leq i}\{A_j - e_i\}, \rho\},$$

$$P_i^+ = \min_{j \leq i}\{l_j - A_i\},$$

$$S_i^- = \min_{j \geq i}\{A_j - e_i\},$$

$$S_i^+ = \min\{\min_{j \geq i}\{l_j - A_i\}, \sigma\},$$

where ρ and σ are the durations of the slack periods immediately preceding and following the active period in question. P_i^- (P_i^+) denotes the maximum amount of time by which every stop preceding but not including i can be advanced (delayed) without violating the time windows, and S_i^- (S_i^+) denotes the maximum amount of time by which every stop following but not including i can be advanced (delayed). These quantities thus indicate how much each segment of an active period can be displaced to accommodate an additional customer. Once it is established that some way of inserting the pickup and delivery of customer i satisfies the time window constraints, it must be ascertained that it satisfies the maximum travel time constraints.

Psaraftis [31] compares and tests these two approaches and concludes that the three-phase algorithm can be best used as a fast planning tool or as a device to produce good starting solutions in an operational situation, whereas the insertion algorithm can form the basis of an operational scheduling system that would assist the dispatcher in the actual execution of the schedule.

Bodin and Sexton [4] extend their algorithm for the single-vehicle dial-a-ride problem with desired delivery times to a two-phase algorithm for the multiple-vehicle dial-a-ride problem with desired delivery times. In the first phase, the clients are assigned to vehicles and a single-vehicle dial-a-ride problem is solved to construct an initial feasible solution. In the second phase, the algorithm attempts to reassign clients with bad service to other vehicles in order to find a better solution. This swapping of clients is done in the following way: the algorithm passes repeatedly through the client list, each time reassigning clients with a client inconvenience larger than some threshold. This threshold depends on the pass the algorithm is on.

4.1.3 The static full-truck-load pickup and delivery problem

The full-truck-load problem is the special case of the GPDP in which each load has to be transported directly from its origin to its destination, i.e., $|N_i^+| = |N_i^-| = 1$ for all $i \in N$, $\bar{q}_i = 1$ for all $i \in N$, and $Q_k = 1$ for all $k \in M$. In these problems, a transportation request is usually called a trip.

The full-truck-load is an interesting special case since it can be easily transformed into an asymmetric TSP. The set of vertices of this asymmetric TSP corresponds to the set of transportation requests, and the distance between two vertices i, j is taken to be $d_{i-} j+$.

Ball, Golden, Assad and Bodin [3] present two different route first-cluster second approximation procedures for the static full-truck-load problem. Both procedures construct one giant route servicing all transportation requests, and then divide this route into feasible vehicle routes. The first procedure constructs the giant route by solving a rural postman problem. The objective in this problem is to find a minimum cost cycle that traverses each arc in a given subset of arcs, in this case all arcs corresponding to traveling from an origin to a destination, at least once. The second procedure applies the above suggested transformation and constructs a giant route by solving an asymmetric TSP.

Desrosiers, Laporte, Sauve, Soumis and Taillefer [12] develop an optimization algorithm for a multiple-depot full-truck-load problem in which the objective is to minimize the total distance traveled and there is a restriction on the length of the routes. The problem is transformed into an asymmetric TSP with a reduced set of subtour elimination constraints, prohibiting only subtours not including any depot, and distance constraints, prohibiting illegal subpaths of the Hamiltonian cycle. The transformed problem is solved by an algorithm of Laporte, Nobert and Taillefer [24], which is a direct extension of the algorithm of Carpaneto and Toth [7] for the asymmetric TSP.

4.2 The dynamic pickup and delivery problem

As in most combinatorial optimization problems, dynamic aspects of the pickup and delivery problem are not very well studied.

4.2.1 The dynamic single-vehicle pickup and delivery problem

Psaraftis [27] extends the dynamic programming algorithm described in Section 4.1.1 for the static immediate request dial-a-ride problem to the dynamic case. Indefinite deferment of customers, i.e., continuously reassigning service of a customer to the last position in the pickup and delivery sequence, is prevented with a special priority constraint.

The times at which requests for service are received define a natural order among the clients. In general, the position that a particular customer holds in the sequence of pickups will not be the same as his position in this natural order. The difference between these two positions defines a *pickup position shift*. A *delivery position shift* can be similarly defined as the difference between the position in the sequence of deliveries and the position in the natural order. A priority constraint bounding the two position shifts from above prevents indefinite deferment.

The dynamic problem is solved as a sequence of static problems. Each time a new request for service is received, a slightly modified instance of the static problem is solved to update the current route. Obviously, all clients that have already been picked up and delivered can be discarded and the new client has to be incorporated. The starting location of the vehicle and the origins of the clients that have been picked up but not yet delivered have to be set to the location of the vehicle at the time of the update.

4.2.2 The dynamic multiple-vehicle pickup and delivery problem

Psaraftis [32] develops an algorithm for the dynamic multiple-vehicle problem in which the vehicles are in fact ships. In this case, the capacity of the ports also has to be considered in order to avoid waiting times when loads are to be picked up or delivered. The algorithm is based on a rolling horizon principle. Let t_k be the 'current time', i.e., the time at the k th iteration of the procedure. At time t_k the algorithm only considers those known loads i whose earliest pickup time e_{i+} falls between t_k and $t_k + L$, where L , the length of the rolling horizon, is a user input. The algorithm then makes a tentative assignment of loads to eligible ships. However, only loads with e_{i+} within the front end of the interval $[t_k, t_k + L]$ are considered for permanent assignment. In this way the algorithm places less (but still some) emphasis on the less reliable information about the future. Iteration $k + 1$ will move the 'current time' to t_{k+1} , which is equal to the time a significant input update has to be made, or to the lowest e_{i+} of all yet unassigned loads, whichever of the two is the earliest.

The tentative assignment of loads to ships is calculated in two phases. First calculate for each load i and ship j the utility $u_{i,j}$ of assigning load i to ship j . This utility is a complicated function measuring the assignment's effect on (1) the delivery time of load i , (2) the delivery times of all other loads already assigned to ship j , (3) the efficiency of use of ship j , and (4) the system's port resources. In the second phase solve the following assignment problem:

$$\begin{aligned} & \text{minimize} && \sum_i \sum_j u_{i,j} x_{i,j} \\ & \text{subject to} && \sum_i x_{i,j} \leq K \quad \text{for all } j \\ & && \sum_j x_{i,j} \leq 1 \quad \text{for all } i \\ & && x_{i,j} \in \{0, 1\} \quad \text{for all } i, j. \end{aligned}$$

where K is a user specified integer denoting that no more than K loads are assigned to each ship.

4.2.3 The dynamic full-truck-load pickup and delivery problem

Powell [25, 26] considers a dynamic full-truck-load pickup and delivery problem in which transportation requests may be rejected. He develops an algorithm based on a network flow representation of the problem. In this model the planning area is divided into regions R and the time axis with time horizon P is divided into days $\{0, 1, 2, \dots, P\}$. The network has node set $(R \times \{0, 1, 2, \dots, P\}) \cup S \cup T$, where S represents a source and T represents a sink. The arc set consists of arcs representing loaded moves that are known at time $t = 0$, arcs representing empty moves, all arcs $((r, t), (r, t + 1))$, where $r \in R$ and $0 \leq t < P$, and all arcs $(S, (r, 0))$ and $((r, P), T)$, where $r \in R$. Arcs representing a loaded move have capacity 1, arcs $(S, (r, 0))$ have capacity equal to the number of vehicles available in region r on day 0. All other arcs have capacity ∞ . The profit of an arc representing a move is the net profit or cost of the corresponding move. The profit of all other arcs is 0. Note that a maximum profit flow in this network corresponds to an allocation of vehicles to transportation requests, but that it is not required to honor all transportation requests.

To anticipate future transportation requests, the network is extended with stochastic links. These stochastic links correspond to future uncertain trajectories of vehicles. They emanate from each node (r, t) and end in T . The profit of the k th stochastic link emanating from (r, t)

reflects the expected marginal value of another vehicle in region r on day t given that there are already k vehicles in region r on day t . These expected marginal values of an extra vehicle are based on historical data. Each stochastic link has capacity 1. A maximum profit flow in this extended network not only represents a deterministic allocation of vehicles to loads known at $t = 0$, but also assigns vehicles to regions in order to be able to serve future transportation requests at minimal cost. The use of a stochastic link from (r, t) to T indicates that at time t the vehicle will be available in region r to serve some request that is not yet known. Because the vehicles are only allocated to known loads, the network has to be reoptimized a few times each day.

5 Conclusion

In this paper, we have discussed various characteristics of pickup and delivery problems and have given an overview of the solution approaches that have been proposed. In the process, we have identified several important research topics.

Most real-life pickup and delivery problems that we are aware of are demand responsive. Currently, very little is known about on-line algorithms for dynamic pickup and delivery problems. Besides the obvious practical importance of such algorithms, it seems to be a fascinating and challenging research area as well.

Although the single-vehicle pickup and delivery problem is \mathcal{NP} -hard, it can be solved very efficiently with dynamic programming as long as the number of transportation requests is relatively small, which is the case in many practical situations. Therefore, the main problem in solving multiple-vehicle pickup and delivery problems is the assignment of transportation requests to vehicles.

Consequently, pickup and delivery problems, as well as many other routing and scheduling problems, seem to be well suited for solution approaches based on set partitioning with column generation. Although this approach has already been explored successfully, it is likely to remain an active research area for the next decade.

An important obstacle in the development of effective and efficient approximation algorithms is the lack of a reliable measure of closeness. In vehicle routing problems, customers that are geographically close to each other are likely to be served by the same vehicle. A concept similar to that of geographical closeness in vehicle routing problems does not exist for pickup and delivery problems. Although several alternatives have been proposed, none of them has turned out to be entirely satisfactory. The development of a useful closeness concept seems crucial for progress in approximation algorithms.

Many interesting questions arise in the context of interactive planning systems for pickup and delivery problems. Representing solutions, for instance, is nontrivial: an optimal solution to a single vehicle pickup and delivery problem may look very bad when drawn on a map, even without time windows and with infinite vehicle capacity.

References

- [1] M. BALL, B. GOLDEN, A. ASSAD AND L. BODIN (1983). Planning for truck fleet size in the presence of a common carrier option. *Decision Sciences* 14, 103-120.

- [2] J. BEARDWOOD, J.H. HALTON AND J.M. HAMMERSLEY (1959). The shortest path through many points. *Proc. Cambridge Philos. Soc.* 55, 299-327.
- [3] L. BODIN, B. GOLDEN, A. ASSAD AND M. BALL (1983). Routing and scheduling of vehicles and crews – The state of the art. *Computers and Operations Research* 10, 63-211.
- [4] L. BODIN AND T. SEXTON (1983). *The multi-vehicle subscriber dial-a-ride problem*. Management Science & Statistics Working Paper No. 83-009, University of Maryland at College Park.
- [5] L.J.J. VAN DER BRUGGEN, J.K. LENSTRA AND P.C. SCHUUR (1990). *A variable depth approach for the single-vehicle pickup and delivery problem with time windows*. COSOR Memorandum 90-48, Eindhoven University of Technology.
- [6] F.H. CULLEN, J.J. JARVIS AND H.D. RATLIFF (1981). Set partitioning based heuristics for interactive routing. *Networks* 11, 125-143.
- [7] G. CARPANETO AND P. TOTH (1980). Some new branching and bounding criteria for the asymmetric travelling salesman problem. *Management Science* 26, 736-743.
- [8] M. DESROCHERS, J.K. LENSTRA, M.W.P. SAVELSBERGH AND F. SOUMIS (1988). Vehicle routing with time windows: optimization and approximation. B.L. GOLDEN AND A.A. ASSAD (editors), *Vehicle Routing: Methods and Studies*, North-Holland, Amsterdam, 65-84.
- [9] J. DESROSIERS, Y. DUMAS AND F. SOUMIS (1984, revised in 1986). *The multiple vehicle many to many routing problem with time windows*. Ecole des Hautes Etudes Commerciales in Montreal, report G-84-13.
- [10] J. DESROSIERS, Y. DUMAS AND F. SOUMIS (1986). A dynamic programming solution of the large-scale single-vehicle dial-a-ride problem with time windows. *American Journal of Mathematical and Management Sciences* 6, 301-325.
- [11] J. DESROSIERS, Y. DUMAS, F. SOUMIS, S. TAILLEFER AND D. VILLENEUVE (1991). *An algorithm for mini-clustering in handicapped transport*. Ecole des Hautes Etudes Commerciales in Montreal, report G-91-02.
- [12] J. DESROSIERS, G. LAPORTE, M. SAUVE, F. SOUMIS AND S. TAILLEFER (1988). Vehicle routing with full loads. *Operations Research* 15, 219-226.
- [13] Y. DUMAS AND J. DESROSIERS (1986). *A shortest path problem for vehicle routing with pick-up, delivery and time windows*. Ecole des Hautes Etudes Commerciales in Montreal, report G-86-09(A).
- [14] Y. DUMAS, J. DESROSIERS AND F. SOUMIS (1989). *Large scale multi-vehicle dial-a-ride systems*. Ecole des Hautes Etudes Commerciales in Montreal, report G-89-30.
- [15] Y. DUMAS, J. DESROSIERS AND F. SOUMIS (1991) The pickup and delivery problem with time windows. *European Journal of Operational Research* 54, 7-22.
- [16] M.T. FIALA TIMLIN AND W.R. PULLEYBLANK (1990). *Precedence constrained routing and helicopter scheduling: Heuristic design*. Institute for Computer Research, University of Waterloo, Ontario, Canada, UW/ICR 90-02.

- [17] M.FISCHETTI AND P. TOTH (1989) An additive bounding procedure for combinatorial optimization problems. *Operations Research* 37, 319-328.
- [18] D.S. JOHNSON (1992) Private communication.
- [19] J.J. JAW, A.R. ODONI, H.N. PSARAFTIS AND N.H.M. WILSON (1982). *A heuristic algorithm for the multi-vehicle many-to-many advance-request dial-a-ride problem*. Working Paper MIT-UMTA-82-3, M.I.T. Cambridge, Massachusetts 02139.
- [20] J.J. JAW, A.R. ODONI, H.N. PSARAFTIS AND N.H.M. WILSON (1986). A heuristic algorithm for the multi vehicle advance-request dial-a-ride problem with time windows. *Transportation Research* 20B, 243-257.
- [21] B. KALANTARI, A.V. HILL AND S.R. ARORA (1985) An algorithm for the traveling salesman problem with pickup and delivery customers. *European Journal of Operational Research* 22, 377-386.
- [22] G.A.P. KINDERVATER AND M.W.P. SAVELSBERGH (1992). *Local search in physical distribution management*. Memorandum COSOR 92-30, Eindhoven University of Technology.
- [23] J.D.C. LITTLE, K.G. MURTY, D.W. SWEENEY AND C. KAREL (1963). An algorithm for the traveling salesman problem. *Operations Research* 11, 972-989.
- [24] G. LAPORTE, Y. NOBERT AND S. TAILLEFER (1987). A branch and bound algorithm for the asymmetrical distance constrained vehicle routing problem. *Math. Modelling* 9, 857-868.
- [25] W.B. POWELL (1987). An operational planning model for the dynamic vehicle allocation problem with uncertain demands. *Transportation Research* 21B, 217-232.
- [26] W.B. POWELL (1988). Maximizing profits for North American Van Lines' truckload division: a new framework for pricing and operations. *Interfaces* 18, 21-41.
- [27] H. PSARAFTIS (1980). A dynamic programming solution to the single vehicle many-to-many immediate request dial-a-ride problem. *Transportation Science* 14, 130-154.
- [28] H. PSARAFTIS (1983). An exact algorithm for the single vehicle many-to-many dial-a-ride problem with time windows. *Transportation Science* 17, 351-357.
- [29] H. PSARAFTIS (1983). k -Interchange procedures for local search in a precedence-constrained routing problem. *European Journal of Operational Research* 13, 391-402.
- [30] H. PSARAFTIS (1983). Analysis of an $O(N^2)$ heuristic for the single vehicle many-to-many euclidean dial-a-ride problem. *Transportation Research* 17B, 133-145.
- [31] H. PSARAFTIS (1986). Scheduling large-scale advance-request dial-a-ride systems. *American Journal of Mathematical and Management Sciences* 6, 327-367.
- [32] H.N. PSARAFTIS (1988). Dynamic vehicle routing problems, B.L. GOLDEN AND A.A. ASSAD (editors), *Vehicle Routing: Methods and Studies*, North-Holland, Amsterdam, 223-248.
- [33] M.W.P. SAVELSBERGH (1992). *Computer Aided Routing*, CWI Tract 75, Amsterdam .

- [34] M.W.P. SAVELSBERGH (1990). An efficient implementation of local search algorithms for constrained routing problems. *European Journal of Operational Research* 47, 75-85.
- [35] T. SEXTON AND L. BODIN (1985). Optimizing single vehicle many-to-many operations with desired delivery times: I Scheduling. *Transportation Science* 19, 378-410.
- [36] T. SEXTON AND L. BODIN (1985). Optimizing single vehicle many-to-many operations with desired delivery times: II Routing. *Transportation Science* 19, 411-435.
- [37] T.R. SEXTON AND Y-M. CHOI (1986). Pickup and delivery of partial loads with 'soft' time windows. *American Journal of Mathematical and Management Sciences* 6, 369-398.
- [38] M.M. SOLOMON AND J. DESROSIERS (1988). Time window constrained routing and scheduling problems. *Transportation Science* 22, 1-13.
- [39] R.A. SKITT AND R.R. LEVARY (1985). Vehicle routing via column generation. *European Journal of Operational Research* 21, 65-76.
- [40] D.M. STEIN (1978). An asymptotic probabilistic analysis of a routing problem, *Mathematics of Operations Research* 3, 89-101.
- [41] D.M. STEIN (1978). Scheduling dial-a-ride transportation systems. *Transportation Science* 12, 232-249.

List of COSOR-memoranda - 1992

Number	Month	Author	Title
92-01	January	F.W. Steutel	On the addition of log-convex functions and sequences
92-02	January	P. v.d. Laan	Selection constants for Uniform populations
92-03	February	E.E.M. v. Berkum H.N. Linssen D.A. Overdijk	Data reduction in statistical inference
92-04	February	H.J.C. Huijberts H. Nijmeijer	Strong dynamic input-output decoupling: from linearity to nonlinearity
92-05	March	S.J.L. v. Eijndhoven J.M. Soethoudt	Introduction to a behavioral approach of continuous-time systems
92-06	April	P.J. Zwietering E.H.L. Aarts J. Wessels	The minimal number of layers of a perceptron that sorts
92-07	April	F.P.A. Coolen	Maximum Imprecision Related to Intervals of Measures and Bayesian Inference with Conjugate Imprecise Prior Densities
92-08	May	I.J.B.F. Adan J. Wessels W.H.M. Zijm	A Note on "The effect of varying routing probability in two parallel queues with dynamic routing under a threshold-type scheduling"
92-09	May	I.J.B.F. Adan G.J.J.A.N. v. Houtum J. v.d. Wal	Upper and lower bounds for the waiting time in the symmetric shortest queue system
92-10	May	P. v.d. Laan	Subset Selection: Robustness and Imprecise Selection
92-11	May	R.J.M. Vaessens E.H.L. Aarts J.K. Lenstra	A Local Search Template (Extended Abstract)
92-12	May	F.P.A. Coolen	Elicitation of Expert Knowledge and Assessment of Im- precise Prior Densities for Lifetime Distributions
92-13	May	M.A. Peters A.A. Stoorvogel	Mixed H_2/H_∞ Control in a Stochastic Framework

Number	Month	Author	Title
92-14	June	P.J. Zwietering E.H.L. Aarts J. Wessels	The construction of minimal multi-layered perceptrons: a case study for sorting
92-15	June	P. van der Laan	Experiments: Design, Parametric and Nonparametric Analysis, and Selection
92-16	June	J.J.A.M. Brands F.W. Steutel R.J.G. Wilms	On the number of maxima in a discrete sample
92-17	June	S.J.L. v. Eijndhoven J.M. Soethoudt	Introduction to a behavioral approach of continuous-time systems part II
92-18	June	J.A. Hoogeveen H. Oosterhout S.L. van der Velde	New lower and upper bounds for scheduling around a small common due date
92-19	June	F.P.A. Coolen	On Bernoulli Experiments with Imprecise Prior Probabilities
92-20	June	J.A. Hoogeveen S.L. van de Velde	Minimizing Total Inventory Cost on a Single Machine in Just-in-Time Manufacturing
92-21	June	J.A. Hoogeveen S.L. van de Velde	Polynomial-time algorithms for single-machine bicriteria scheduling
92-22	June	P. van der Laan	The best variety or an almost best one? A comparison of subset selection procedures
92-23	June	T.J.A. Storcken P.H.M. Ruys	Extensions of choice behaviour
92-24	July	L.C.G.J.M. Habets	Characteristic Sets in Commutative Algebra: an overview
92-25	July	P.J. Zwietering E.H.L. Aarts J. Wessels	Exact Classification With Two-Layered Perceptrons
92-26	July	M.W.P. Savelsbergh	Preprocessing and Probing Techniques for Mixed Integer Programming Problems

Number	Month	Author	Title
92-27	July	I.J.B.F. Adan W.A. van de Waarsenburg J. Wessels	Analysing $E_k E_r c$ Queues
92-28	July	O.J. Boxma G.J. van Houtum	The compensation approach applied to a 2×2 switch
92-29	July	E.H.L. Aarts P.J.M. van Laarhoven J.K. Lenstra N.L.J. Ulder	Job Shop Scheduling by Local Search
92-30	August	G.A.P. Kindervater M.W.P. Savelsbergh	Local Search in Physical Distribution Management
92-31	August	M. Makowski M.W.P. Savelsbergh	MP-DIT Mathematical Program data Interchange Tool
92-32	August	J.A. Hoogeveen S.L. van de Velde B. Veltman	Complexity of scheduling multiprocessor tasks with prespecified processor allocations
92-33	August	O.J. Boxma J.A.C. Resing	Tandem queues with deterministic service times
92-34	September	J.H.J. Einmahl	A Bahadur-Kiefer theorem beyond the largest observation
92-35	September	F.P.A. Coolen	On non-informativeness in a classical Bayesian inference problem
92-36	September	M.A. Peters	A Mixed H_2/H_∞ Function for a Discrete Time System
92-37	September	I.J.B.F. Adan J. Wessels	Product forms as a solution base for queueing systems
92-38	September	L.C.G.J.M. Habets	A Reachability Test for Systems over Polynomial Rings using Gröbner Bases
92-39	September	G.J. van Houtum I.J.B.F. Adan J. Wessels W.H.M. Zijm	The compensation approach for three or more dimensional random walks

Number	Month	Author	Title
92-40	September	F.P.A. Coolen	Bounds for expected loss in Bayesian decision theory with imprecise prior probabilities
92-41	October	H.J.C. Huijberts H. Nijmeijer A.C. Ruiz	Nonlinear disturbance decoupling and linearization: a partial interpretation of integral feedback
92-42	October	A.A. Stoorvogel A. Saberi B.M. Chen	The discrete-time H_∞ control problem with measurement feedback
92-43	October	P. van der Laan	Statistical Quality Management
92-44	November	M. Sol M.W.P. Savelsbergh	The General Pickup and Delivery Problem