

**An Analysis of the Behavior of
Heuristics for the Vehicle Routing
Problem for a selection of problems
with Vehicle-related, Customer-related,
and Time-related Constraints**

Alex Van Breedam
Faculty of Applied Economics
University of Antwerp - RUCA
Middelheimlaan 1
B-2020 Antwerp (Belgium)
E-mail: alexvb@ruca.ua.ac.be

September 1994

Contents

Preface	7
1 Introduction and experimental plan	8
1.1 The Vehicle Routing Problem	8
1.2 Mathematical programming formulation	9
1.3 Experimental plan	10
2 Implementation of the initial heuristics	14
2.1 Sequential Nearest Neighbor heuristic	15
2.2 Parallel Nearest Neighbor heuristic	16
2.3 Sequential Savings heuristic	17
2.4 Parallel Savings heuristic	19
2.5 Generalized Savings heuristic	20
2.6 Sequential Insertion heuristic	20
2.7 Parallel Insertion heuristic	23
2.8 Parallel Assignment-based Insertion heuristic	24
2.9 Generalized Assignment heuristic	25
2.10 Two Phase heuristic	27
2.11 Sweep heuristic	28
3 Analysis of the effect of vehicle-related constraints	30
3.1 Parametric analyses for the test sets G1, G2 and G3	31
3.1.1 Sequential Nearest Neighbor heuristic	32
3.1.2 Parallel Nearest Neighbor heuristic	33
3.1.3 Sequential Savings heuristic	34
3.1.4 Parallel Savings heuristic	35
3.1.5 Generalized Savings heuristic	36
3.1.6 Sequential Insertion heuristic	37
3.1.7 Parallel Insertion heuristic	38
3.1.8 Parallel Assignment-based Insertion heuristic	40
3.1.9 Generalized Assignment heuristic	41
3.1.10 Two Phase heuristic	42

3.1.11	Conclusions of the parametric analysis	44
3.2	Heuristic analysis	44
3.2.1	Heuristic analysis for test set G1	45
3.2.2	Heuristic analysis for test set G2	47
3.2.3	Heuristic analysis for test set G3	48
4	Analysis of the effect of Customer-related constraints	53
4.1	Parametric analyses for the test set P1	54
4.1.1	Sequential Nearest Neighbor heuristic	54
4.1.2	Parallel Nearest Neighbor heuristic	55
4.1.3	Sequential Savings heuristic	56
4.1.4	Parallel Savings heuristic	57
4.1.5	Generalized Savings heuristic	58
4.1.6	Sequential Insertion heuristic	58
4.1.7	Parallel Insertion heuristic	59
4.1.8	Parallel Assignment-based Insertion heuristic	60
4.1.9	Generalized Assignment heuristic	61
4.1.10	Two Phase heuristic	61
4.1.11	Conclusions of the parametric analysis for test set P1	62
4.2	Heuristic analysis for test set P1	62
4.3	Parametric analysis for the test set P2	64
4.3.1	Sequential Nearest Neighbor heuristic	65
4.3.2	Parallel Nearest Neighbor heuristic	66
4.3.3	Sequential Savings heuristic	66
4.3.4	Parallel Savings heuristic	67
4.3.5	Generalized Savings heuristic	68
4.3.6	Sequential Insertion heuristic	68
4.3.7	Parallel Insertion heuristic	69
4.3.8	Parallel Assignment-based Insertion heuristic	70
4.3.9	Generalized Assignment heuristic	71
4.3.10	Two Phase heuristic	72
4.3.11	Conclusions of the parametric analysis for test set P2	73
4.4	Heuristic analysis for test set P2	73
5	Analysis of the effect of time-related constraints	76
5.1	Parametric analysis for test sets T1 and T2	77
5.1.1	Sequential Nearest Neighbor heuristic	77
5.1.2	Parallel Nearest Neighbor heuristic	79
5.1.3	Sequential Savings heuristic	80
5.1.4	Parallel Savings heuristic	81
5.1.5	Generalized Savings heuristic	82
5.1.6	Sequential Insertion heuristic	82
5.1.7	Parallel Insertion heuristic	84
5.1.8	Parallel Assignment-based Insertion heuristic	85

5.1.9	Generalized Assignment heuristic	86
5.1.10	Two Phase heuristic	87
5.1.11	Conclusions of the parametric analysis	88
5.2	Heuristic analysis	89
5.2.1	Heuristic analysis for test set T1	89
5.2.2	Heuristic analysis for test set T2	90
6	Implementations of the improvement heuristics	93
6.1	Types of move	94
6.2	Local Improvement heuristic	97
6.3	Simulated Annealing metaheuristic	98
6.4	Tabu Search metaheuristic	103
7	Analysis of the improvement heuristics	107
7.1	Reduced test set	108
7.2	Parametric analyses of the improvement heuristics	109
7.2.1	Local Improvement heuristic	109
7.2.2	Simulated Annealing metaheuristic	111
7.2.3	Tabu Search metaheuristic	113
7.3	Heuristic analysis of the improvement heuristics	115
7.3.1	String Cross	117
7.3.2	String Exchange	118
7.3.3	String Relocation	122
7.3.4	String Mix	123
7.4	Conclusions for the analyses of the improvement methods	125
8	Synthesis and guidelines for further research	128
A	The basic set of geographical problems	134
B	Descriptions of some procedures and tests used	138
B.1	Seed point generation methods	138
B.1.1	Circle Covering	139
B.1.2	Cone covering	139
B.2	Heuristic for the TSP	140
B.2.1	Insertion heuristic	140
B.2.2	3-opt improvement heuristic	141
B.3	Post-processor	141
B.4	Automatic Interaction Detection	141
B.5	Friedman test	142
C	Graphs for the improvement heuristics	144

List of Tables

1.1	The five values of meta-parameter A	12
3.1	Summary of the relative performances of the 11 heuristics for test set G1	47
3.2	Summary of the relative performances of the 11 heuristics for test set G2	49
3.3	Summary of the relative performances of the 11 heuristics for test set G3	51
4.1	Summary of the relative performances of the 11 heuristics for test set P1	64
4.2	Summary of the relative performances of the 11 heuristics for test set P2	75
5.1	Summary of the relative performances of the 11 heuristics for test set T1	90
5.2	Summary of the relative performances of the 11 heuristics for test set T2	92
7.1	The reduced test set with the 15 problems.	110
7.2	Heuristic analysis of LI, TS and SA with move type SC.	119
7.3	Heuristic analysis of LI, TS and SA with move type SE.	121
7.4	Heuristic analysis of LI, TS and SA with move type SR.	123
7.5	Heuristic analysis of LI, TS and SA with move type SM.	127
8.1	Summary of the parametric analysis of the initial heuristics for the seven test sets.	130
8.2	Summary of the parametric analysis of the improvement heuristics for the four types of move and the aggregated solutions.	132

List of Figures

6.1	Example of a String Cross.	95
6.2	Example of a String Exchange.	96
6.3	Example of a String Relocation.	96
7.1	Example of the path of the objective function value for the LI, TS and SA heuristics.	116
A.1	Geographical problems of the basic set. Across: spreading pat- terns; down: grouping patterns. Depot location: central.	135
A.2	Geographical problems of the basic set. Across: spreading pat- terns; down: grouping patterns. Depot location: inside.	136
A.3	Geographical problems of the basic set. Across: spreading pat- terns; down: grouping patterns. Depot location: outside.	137
B.1	Example of a tree produced by the AID-analyses.	142
C.1	String Cross for problem 1	145
C.2	String Cross for problem 2	145
C.3	String Cross for problem 3	145
C.4	String Cross for problem 4	146
C.5	String Cross for problem 5	146
C.6	String Cross for problem 6	146
C.7	String Cross for problem 7	147
C.8	String Cross for problem 8	147
C.9	String Cross for problem 9	147
C.10	String Cross for problem 10	148
C.11	String Cross for problem 11	148
C.12	String Cross for problem 12	148
C.13	String Cross for problem 13	149
C.14	String Cross for problem 14	149
C.15	String Cross for problem 15	149
C.16	String Exchange for problem 1	150
C.17	String Exchange for problem 2	150

C.18 String Exchange for problem 3	150
C.19 String Exchange for problem 4	151
C.20 String Exchange for problem 5	151
C.21 String Exchange for problem 6	151
C.22 String Exchange for problem 7	152
C.23 String Exchange for problem 8	152
C.24 String Exchange for problem 9	152
C.25 String Exchange for problem 10	153
C.26 String Exchange for problem 11	153
C.27 String Exchange for problem 12	153
C.28 String Exchange for problem 13	154
C.29 String Exchange for problem 14	154
C.30 String Exchange for problem 15	154
C.31 String Relocation for problem 8	155
C.32 String Relocation for problem 9	155
C.33 String Relocation for problem 10	155
C.34 String Relocation for problem 11	156
C.35 String Relocation for problem 12	156
C.36 String Relocation for problem 13	156
C.37 String Relocation for problem 14	157
C.38 String Relocation for problem 15	157
C.39 String Mix for problem 8	158
C.40 String Mix for problem 9	158
C.41 String Mix for problem 10	158
C.42 String Mix for problem 11	159
C.43 String Mix for problem 12	159
C.44 String Mix for problem 13	159
C.45 String Mix for problem 14	160
C.46 String Mix for problem 15	160

Preface

This book contains the main results of my Ph.D. thesis, defended successfully on the 28th of September 1994 at the Faculty of Applied Economics, University of Antwerp-RUCA. The original Dutch title is: *'Een analyse van het gedrag van rittenplanningsheuristieken voor een selectie van planningsproblemen met vrachtwagengebonden, klantgebonden en tijdgebonden randvoorwaarden'*. The main subject of the thesis consists of the analysis of the behavior of a set of heuristics for the Vehicle Routing Problem. Eleven initial heuristics and three improvement heuristics were surveyed on specifically designed problems. The behavior of the heuristics was examined on a parametric level and a heuristic level.

The Ph.D. thesis is the emanation of five years of intensive research, 60000 lines of computer code and thousands of computer runs. In order to realize this, the support of some people was indispensable.

First, I wish to thank professor H. Müller-Malek for his supervision. Besides, the recommendations of the professors F. Broeckx, G. Thiers, G. Janssens were an important aid. The comments of the professors R. Dillmann, M. Labbé and D. Van Oudheusden were very useful.

Furthermore, the following people were very helpful: G. Asselberghs, W. Bleumers, F. Bogaerts, I. Buitink, C. Coeck, A. De Beuckelaer, W. Deceuninck, P. De Pelsmacker, D. De Maeyer, A. De Schepper, D. Newman, J. Piot, I. Roosen, W. Tielemans, J. Van Breedam, L. Verheggen, J. Verelst, P. Willemé and W. Winkelmans.

I thank my wife Hilde for the cover design and most illustrations.

Last, but not least, I wish to thank my parents and particularly my wife Hilde and my elder daughter Stephanie for their patience and support. Moreover, I am grateful to my second daughter Elise for having postponed her arrival to the most appropriate time. I dedicate this work to them.

Alex Van Breedam

Chapter 1

Introduction and experimental plan

1.1 The Vehicle Routing Problem

The standard Vehicle Routing Problem (VRP) can be defined as one of finding a set of routes for a fleet of vehicles which have to service a number of stops n from a depot 0. It is assumed that every vehicle has the same capacity Q and that the number of vehicles is unlimited. The vehicles depart and arrive at the depot. The demand quantity q_i at each stop i is known in advance and is deterministic. No single demand quantity exceeds vehicle capacity.

The VRP was first formulated by Dantzig and Ramser (1959). Besides vehicle capacity, other side-constraints can be included in the VRP, such as maximum travel distance and/or travel time, time-windows at the stops and mixed pick-up and delivery. The VRP with or without side-constraints is an \mathcal{NP} -hard problem (Lenstra and Rinnooy Kan (1981)).

A number of assumptions were made for the research presented here. The VRP is an Euclidean network problem. The distance between two nodes is the Euclidean distance (measured in km.). The same figure is used for the travel time, i.e. a vehicle drives at a constant speed of one km. per minute. All vehicles are homogeneous with respect to their capacity. The speed of a vehicle is constant and is 1 kilometer per minute.

All planning problems have 100 customers and a single depot. The number of vehicles K available is unrestricted. For solving problems of this size at this point in time, only heuristics are applicable. Exact algorithms have been used to solve problems of 30 to 50 stops.

The heuristics considered belong to two categories: initial and improvement

heuristics. Initial heuristics generate a feasible solution for the VRP, given the data on customers, depot and side-constraints.

The improvement heuristics considered are assumed to enhance a feasible solution, generated by means of an initial heuristic, through moving customers between routes.

1.2 Mathematical programming formulation

For elucidation purposes, the mathematical programming formulation of the VRP is proposed. Given is an Euclidean graph $G = (N, E)$ with a set N of vertices and a set E of arcs. The set N is defined as $N = 0, 1, \dots, n$ where $\{0\}$ indicates the depot. All VRPs considered have $n = 100$ customers.

A three-index VRP formulation is used. It is based on Bodin et al. (1983). A good overview of VRP formulations is presented in Laporte and Nobert (1987).

The objective is to minimize the total travel time, given by

$$\min \sum_{i=0}^n \sum_{j=0}^n \sum_{k=0}^K t_{ij} x_{ijk} \quad (1.1)$$

subject to

$$\sum_{i=0}^n \sum_{k=1}^K x_{ijk} = 1 \quad j = 1, \dots, n \quad (1.2)$$

$$\sum_{i=0}^n x_{ipk} - \sum_{j=0}^n x_{jpk} = 0 \quad k = 1, \dots, K, \quad p = 0, \dots, n \quad (1.3)$$

$$\sum_{i=0}^n q_i \left(\sum_{j=0}^n x_{ijk} \right) \leq Q_k \quad k = 1, \dots, K \quad (1.4)$$

$$\sum_{j=1}^n x_{0jk} \leq 1 \quad k = 1, \dots, K \quad (1.5)$$

$$\sum_{i \in S} \sum_{j \in S} x_{ijk} \leq |S| - 1 \quad S \subseteq N \setminus \{0\}, 2 \leq |S| \leq n - 2, k = 1, \dots, K \quad (1.6)$$

$$x_{ijk} \in \{0, 1\} \quad i, j = 0, \dots, n \quad k = 1, \dots, K \quad (1.7)$$

The travel time from node i to node j with vehicle k , t_{ijk} , is assumed to be symmetrical, i.e. $t_{ijk} = t_{jik}$. The triangle inequality is satisfied:

$$t_{ij} + t_{jp} \geq t_{ip} \quad i, j, p = 0, \dots, n$$

Equations 1.2 ensure that each stop is serviced by exactly one vehicle. A vehicle must leave a stop after having it served. This route continuity constraint is provided by equations 1.3. Constraints 1.4 are the capacity constraints. It is assumed that there is no demand associated with the depot, i.e. $q_0 = 0$. Inequalities 1.5 guarantee that each vehicle completes one route at a time. Subtour-elimination constraints are represented by constraints 1.6, and imply that the number of arcs linking the nodes of each subset of nodes S must be less than the number of nodes in that subset. Finally, conditions 1.7 guarantee an integer solution.

1.3 Experimental plan

The experimental plan for the research consists of formulating the objectives of the research, the treatments, the experimental units and the analyses required to meet the objectives.

The main *objective* of the research is to gain insight into the behavior of a number of heuristics for various VRPs. The planning problems considered are grouped into several test sets. These test sets are built by adding different types of side-constraints to a basic set of geographic problems. The nature of the vehicle-, the customer- and the time-related side-constraints contained in the test sets frequently occur in real-life situations.

Two levels of *experimental treatments* are distinguished: a heuristic and a parameter level. The behavior of both is analyzed under various planning conditions provided by the above-mentioned test sets. The implementations of the initial and the improvement heuristics are presented in chapters 2 and 6, respectively. All heuristics are programmed in the C++ language.

The *experimental units* are the planning problems. As mentioned previously, all test problems were obtained by adding side-constraints to a basic set of geographic problems. The basic problem set is developed, based on three criteria: the location of the depot, the grouping of customers and the spreading of customers.

Three different depot locations are considered: *central*, *inside* and *outside*. For the grouping of customers, five patterns are distinguished: *singleton*, *clusters*, *50% clusters*, *cones* and *50% cones*. Four patterns are taken into account for the spreading of customers: *uniform*, *50% central*, *concentric* and *compressed*. Exhaustive combination of these patterns gives rise to a basic set of 60 geographic problems.

The conception of the patterns is based on deterministic rules, as can be seen on

the graphical representations of the problems in figures A.1 to A.3 on pages 135 to 137. Stochastic influences are avoided to the greatest possible extent in order to preserve the internal validity of the experiments.

The actual test problems for the initial heuristics are obtained by adding various side-constraints to the 60 problems of the basic set. The construction and the analysis of the test sets with vehicle-, customer- and time-related side-constraints are discussed in chapters 3, 4 and 5, respectively. For the analysis of the behavior of the improvement heuristics, a reduced set of test problems is used. This is justified by the large run times and by the generic parameters of the improvement heuristics, which are a threat to the external validity of the results. The reduced test set is presented in chapter 7.

The *analyses* required to meet the research objectives, are to be performed at two levels: a parametric level and a heuristic level. Common to both analyses is the use of the total travel time of a feasible solution as the dependent variable.

The *parametric analysis* evaluates the behavior of the parameters of each heuristic subjected to a VRP. Automatic Interaction Detection (AID) is a vital tool for realizing the parametric analysis. The purpose of AID is to search a structure in the relation between variables. The relationships found by AID can be represented by a tree structure.

The input of AID-analysis contains all solutions to a planning problem obtained by combining all values of the parameters. AID requires a continuous dependent variable and nominally-scaled independent variables. Therefore, all parameters of the heuristics are transformed in nominally-scaled parameters. Some parameters have a nominal scaled by nature. An example is a parameter which represents the initialization criterion of a sequential route-building heuristic. The two values for the parameter represent the initialization of a route with the stop either farthest from the depot or closest to the depot.

The transformation procedure is harder for parameters representing continuous weights of criteria. Without loss of generality, we decided to transform all criteria in order to have continuous weights summing at 1.

Assume, for example, the savings criterion, originally stated as

$$\begin{aligned} \max_{i,j} [t_{0i} + t_{0j} - \alpha t_{ij}] \\ \alpha \geq 0 \end{aligned}$$

At first, this expression is transformed into

$$\begin{aligned} \max_{i,j} [\alpha_1(t_{0i} + t_{0j}) - \alpha_2 t_{ij}] \\ \alpha_1 \geq 0, \alpha_2 \geq 0, \alpha_1 + \alpha_2 = 1 \end{aligned}$$

The combinations of the continuous weights α_1 and α_2 can be represented by a single meta-parameter. As an example, the five combinations of values for α_1

A	α_1	α_2
1	1	0
2	0.75	0.25
3	0.50	0.50
4	0.25	0.75
5	0	1

Table 1.1: The five values of meta-parameter A .

and α_2 are categorized in a meta-parameter A .

We decided to consider always the same five values for each continuous weight: 0, 0.25, 0.50, 0.75 and 1. Due to the fact that the sum of the weights α_1 and α_2 always sum to 1, the parameter A can take five values with all possible combinations of both weights, as can be seen from table 1.1. Similarly, a meta-parameter representing three continuous weights of a criterion needs fifteen values to represent all combinations.

This approach limits the number of combinations of continuous weights. A drawback to this procedure is that the best solution for a heuristic can be obtained by the combination of values between the restricted set of five values we propose for a continuous weight.

The dependent variable of each AID analysis, the total travel time, is of a continuous nature. With AID one can get an idea of the sets of values for a parameter which yield significantly better solutions for any given problem. A limitation related with the application of AID is that only one replication is available for each combination of parameter values. Consequently, it is not possible to obtain a single significantly better value for each parameter. For most parameters, only a set of significantly better values can be derived. The reader is referred to appendix B.4 for additional information on AID and its specific use in the framework of the parametric analysis.

AID is used for the parametric analysis of the initial as well as for the improvement heuristics. The raw results of the AID-analyses are not reproduced because of their abundance. Nevertheless, the interpretation and description of the results are reported in detail in the chapters 3, 4 and 5 for the initial heuristics, and in chapter 7 for the improvement heuristics.

The *heuristic analysis* is different for both types of heuristics. A common feature is that only the best solution of each heuristic is used for the experimental units considered.

The Friedman test is used for analyzing the relative behavior of the initial heuristics on a set of problems. The fact that the total travel time of different

problems of a set is not comparable, imposes the use of ranks. Heuristics are ranked in order of increasing total travel time. The test uses the rank totals of the heuristics in order to determine significant differences among heuristics on a set of problems. The use of the Friedman test for comparing heuristics has been suggested by Golden and Stewart (1985). The reader is referred to appendix B.5 for additional information on the Friedman test. Chapters 3, 4 and 5 describe the results of the heuristic analysis of the initial heuristics. The huge amount of results does not permit a reproduction of the raw results in this book.

The heuristic analysis is quite different for the improvement heuristics. Two out of the three improvement heuristics are global optimisation heuristics. These methods are characterized by the absence of a deterministic stop criterion. A comparison of final solutions of the improvement heuristics only would cause the local optimisation heuristic to be prejudiced a priori. Hence, the solution of the improvement heuristics is traced at predefined points in time. The raw results of this analysis are reported in appendix C. Chapter 7 of this book contains the heuristic analysis for the improvement heuristics.

All computations required for the research have been run on an 80486DX processor at 33Mhz.

Chapter 2

Implementation of the initial heuristics

Eleven initial heuristics are presented in the survey. They belong to two groups of heuristics: route-building methods and two-phase methods.

Route-building methods construct routes iteratively by adding unrouted stops. Two types of heuristics can be distinguished within this group: sequential and parallel route-building heuristics. Sequential route-building heuristics construct one route at a time, while parallel implementations build all routes simultaneously.

The sequential route-building heuristics included are: the Sequential Nearest Neighbor, the Sequential Savings and the Sequential Insertion heuristics.

The Parallel Nearest Neighbor, the Parallel Savings, the Parallel Generalized Savings, the Parallel Insertion and the Parallel Assignment-based Insertion heuristics are the parallel route-building heuristics.

Two-phase methods produce an initial solution to the VRP in two phases. We will consider two categories of two-phase algorithms. The first category includes cluster-first route-second methods, in which stops are first clustered and then routed with a Traveling Salesman Problem (TSP) being solved within each cluster. The Generalized Assignment and the Sweep heuristic belong to this category.

The second category of two-phase methods contains heuristics which can be considered as a superposition of a sequential and a parallel route-building algorithm. The so-called Two-phase heuristic, is part of this category.

Other classification schemes of initial heuristics can be found in Christofides (1985) and Bodin et al. (1983).

All heuristics are equipped with a *3-opt* within-route improvement procedure, which is invoked each time a change is made to a route. The reader is referred to appendix B.2.2 for additional information on the *3-opt* branch-exchange procedure for the TSP.

Some heuristics cannot guarantee the routing of all stops. If some stops remain unrouted after the execution of a heuristic, a post-processor is used for assigning the unrouted stops to an existing route or, if not possible, to a new route.

For each heuristic, the stepwise procedure and its parameters are described. Each parameter is denoted by a single character given in brackets.

2.1 Sequential Nearest Neighbor heuristic

The Sequential Nearest Neighbor heuristic (SN) builds one route at a time by adding to the current route those unrouted stops, which satisfy the nearest neighbor criterion.

The major drawback of the heuristic is the naivety of its criterion. Routes are built to full capacity utilization. This causes the resulting routes to be poorly separated. Like all route-building heuristics, the SN heuristic is myopic in so far that it cannot look more than one iteration step ahead. Consequently, the added stops are irrevocable.

Implementations of a similar nearest neighbor heuristic have been proposed by Tyagi (1968), Baker and Schaffer (1986), Solomon (1987) and Balakrishnan (1993). The last three publications present a time-oriented nearest neighbor heuristic.

Procedure

- Step 1:* Initialize a new current route with the initialization criterion.
- Step 2:* Select the unrouted stop satisfying the nearest neighbor criterion with respect to an endstop of the current route. If all stops are routed, then stop.
- Step 3:* If the stop selected in step 2 can be added to the current route, then add it and go to step 2, else go to step 4.
- Step 4:* In case of multiple initialization, go to step 1. In case of single initialization, try to insert the unrouted stop in an existing route. If this is possible the route considered becomes the new current route and go to step 2. If the unrouted stop cannot be added to a route, then start a new current route with that stop and go to step 2.

Parameters

Initialization criterion (I)

1. The unrouted stop farthest from the depot.
2. The unrouted stop closest to the depot.
3. Merging the two unrouted nearest neighbors.

Frequency of initialization (P)

1. Multiple initialization: each new route is initialized with the initialization criterion.
2. Single initialization: only the first route is initialized with the initialization criterion.

Places to add stops (A)

1. Unrouted stops are added at the end of the current route.
2. Unrouted stops are added at the begin or at the end of the current route.

Nearest neighbor criterion (C)

For the VRP without time-windows, the nearest neighbor criterion is used.

$$\min_{i,j \in N \setminus \{0\}} t_{ij} \quad (2.1)$$

t_{ij} : travel time between stop i and j .

For the VRP with time-windows, a time-oriented nearest neighbor criterion is used. If stop i is the last stop of the current route, the unrouted stop j is added at the end of the current route if the following expression is satisfied

$$\min_{i,j \in N \setminus \{0\}} [\alpha_1 t_{ij} + \alpha_2 (b_j - (b_i + s_i)) + \alpha_3 (l_{2j} - (b_i + s_i + t_{ij}))] \quad (2.2)$$

$$\alpha_1, \alpha_2, \alpha_3 \geq 0, \alpha_1 + \alpha_2 + \alpha_3 = 1$$

b_i : time to begin service at stop i .

s_i : service time at stop i .

l_{2j} : closing time of stop j . This is the latest time of the second time-window. For the experiments, it is assumed that each stop has two time-windows at most (see chapter 5). The combinations of weights α_1, α_2 , and α_3 are represented by the fifteen values of parameter C (see table 1.1).

2.2 Parallel Nearest Neighbor heuristic

The Parallel Nearest neighbor heuristic (PN) creates routes simultaneously by iteratively adding the nearest unrouted stop to one of the routes. Routes are

initialized by merging the two unrouted nearest neighbors or by means of seed stops.

The main drawbacks of this heuristic are its naive criterion and its inability to fill vehicles to a high capacity utilization.

Procedure

Step 1: Perform the parallel initialization of the routes.

Step 2: Select the unrouted stop nearest to one of both endpoints of a route and of which the addition to the route is feasible. If no stop can be selected, then stop. If stops remain unrouted, invoke the post-processor (cfr. appendix B.3) in order to assign these stops to an existing route or if not possible to a new route.

Step 3: Add the selected stop to the route and go to step 2.

Parameters

Parallel initialization procedure (S)

1. Routes are initialized by merging the two unrouted nearest neighbors in a single route.
2. Seeds are generated by means of the two variants of the circle covering method (see appendix B.1.1). Seeds coincide with the location of stops.
3. Seeds are generated by means of the cone covering method with five different load fractions (see appendix B.1.2). Seeds coincide with the location of stops.

Nearest neighbor criterion

The nearest neighbor criterion 2.1 is used. The use of this criterion does not require any parameter.

No time-oriented nearest neighbor criterion can be used for the PN heuristic.

2.3 Sequential Savings heuristic

The Sequential Savings heuristic (SS) builds one route at a time by adding unrouted stops, which satisfy the savings criterion, to the current route.

Important advantages of the heuristic are primarily the good quality of the savings criterion and its ability to fill up routes to a high utilization rate. The consequence of the latter feature is that routes tend to overlap. Like all route-building methods, also this heuristic is myopic.

The idea of the savings criterion for merging two stops in a route is due to Clarke and Wright (1964). Some authors (Gaskell (1967), Yellow (1970) and

Paessens (1988)) extended the savings criterion with additional terms and parameters. Implementations of a sequential savings heuristic have been proposed by among others Gaskell (1967), Yellow (1970), Webb (1972) and Ziegler et al. (1988).

Procedure

- Step 1:* Initialize the current route with the initialization criterion.
Step 2: Select the unrouted stop satisfying the savings criterion with respect to an endstop of the current route. If all stops are routed, then stop.
Step 3: If the stop selected in step 2 can be added to the current route, then add it and go to step 2, else go to step 4.
Step 4: In case of multiple initialization, go to step 1. In case of single initialization, try to insert the unrouted stop in an existing route. If this is possible, then the route considered becomes the new current route and go to step 2. If the unrouted stop cannot be added to a route, then start a new current route with the stop considered and go to step 2.

Parameters

Initialization criterion (I)

1. The unrouted stop farthest from the depot.
2. The unrouted stop closest to the depot.
3. Merging the two unrouted stops with the highest savings value.

Frequency of initialization (P)

1. Multiple initialization: each new route is initialized with the initialization criterion.
2. Single initialization: only the first route is initialized with the initialization criterion.

Places to add stops (A)

1. Unrouted stops are added at the end of a route.
2. Unrouted stops are added at the begin or at the end of a route.

Savings criterion (C)

Stops i and j are merged in one route if they satisfy the following expression

$$\max_{i,j \in N \setminus \{0\}} [\alpha_1(t_{0i} + t_{0j}) - \alpha_2 t_{ij} + \alpha_3 | t_{0i} - t_{0j} |] \quad (2.3)$$

$$\alpha_1, \alpha_2, \alpha_3 \geq 0, \alpha_1 + \alpha_2 + \alpha_3 = 1$$

At most, one of both stops i or j is an end point of a route.

The combinations of the weights α_1, α_2 , and α_3 are represented by the fifteen values of parameter C .

2.4 Parallel Savings heuristic

The Parallel Savings heuristic (PS) creates routes simultaneously by iteratively adding the unrouted stop which satisfies the savings criterion together with one of both endstops of one of the routes.

The main advantages of the heuristic are the good overall quality of the savings criterion and its ability to keep routes well-separated in spite of a maximal capacity utilization.

The savings heuristic of Clarke and Wright (1964) was the first parallel savings heuristic to be developed. A lot of authors focused on storing, sorting and updating the long savings list proper to the parallel savings implementation (see Golden (1977), Nelson et al. (1988), Bodin (1983)).

Van Landeghem (1988) extended the savings criterion with a time-oriented part in order to take account of time-windows. He used seeds to initialize routes. Some features to relax to some extent the irrevocability of the added stops are proposed by Knowles (1967), Tillman and Cochran (1968), Holmes and Parker (1976), McDonald (1972) and Buxey (1979).

Procedure

Step 1: Perform the parallel initialization of the routes.

Step 2: Select the unrouted stop which yields the greatest savings value with respect to one of the endpoints of a route and of which the addition is feasible. If no stop can be selected in step 2, then stop. If stops remain unrouted, invoke the post-processor (cfr. appendix B.3) to assign these stops to an existing route or, if not possible, to a new route.

Step 3: Add the stop selected to the route and go to step 2.

Parameters

Parallel initialization procedure (S)

1. Routes are initialized by merging the two unrouted stops which yield the highest savings value.
2. Seeds are generated by means of the two variants of the circle covering method (see appendix B.1.1). Seeds coincide with the location of stops.
3. Seeds are generated by means of the cone covering method with five different load fractions (see appendix B.1.2). Seeds coincide with the location of stops.

Savings criterion (C)

The savings criterion 2.3 is used. The combinations of the weights α_1, α_2 , and α_3 are represented by the fifteen values of parameter C .

2.5 Generalized Savings heuristic

The Generalized Savings heuristic (GS) can be considered as an extension of the Parallel Savings heuristic. Not only stops, but entire routes can be merged to form a new route if this yields a savings in travel time. This heuristic is very time-consuming, because a TSP has to be solved in order to evaluate each merger.

The heuristic proposed here is based on the ideas propounded by Altinkemer and Gavish (1991), and Desrochers and Verhoog (1989).

Procedure

Step 1: Merge all stops two by two, using the generalized savings criterion.

Step 2: Select the two routes with the highest value for the generalized savings criterion. If no merger is feasible, then stop.

Step 3: Merge the routes selected in step 2. Go to step 2.

Parameters**Generalized Savings criterion (P)**

The routes K and L are merged if they satisfy the following criterion

$$\begin{aligned} \max_{K,L} [\alpha_1(T_K + T_L) - \alpha_2 T_{KL}] \\ \alpha_1, \alpha_2 \geq 0, \alpha_1 + \alpha_2 = 1 \end{aligned} \quad (2.4)$$

T_K, T_L, T_{KL} : travel time of route K, L and combined route KL respectively.

The combinations of the weights α_1 and α_2 are represented by the five values of parameter P .

2.6 Sequential Insertion heuristic

The Sequential Insertion heuristic (SI) constructs routes one by one. The current route is initialized by the initialization criterion. Stops are selected by a selection criterion to be inserted in the current route at the place determined by the insertion criterion.

This heuristic allows the insertion of an unrouted stop at any place in the current route, and not just at one of both ends as was the case with the SN and SS heuristics.

As a result of its sequential nature, the heuristic tends to build routes to full utilization as a primary objective at the expense of minimizing the travel time. Also this heuristic is myopic.

The first implementation of a sequential insertion has been conceived by Mole and Jameson (1976). Baker and Schaffer (1986), and Solomon (1987) proposed different insertion and selection criteria for time-window problems.

Procedure

- Step 1:* Initialize a new current route with the initialization criterion.
- Step 2:* For every unrouted stop, determine the best feasible insertion place in the current route using the insertion criterion. If no feasible insertion is possible and all stops are not routed, then go to step 1. If all stops are routed, then stop.
- Step 3:* Select the stop with the best value for the selection criterion. Insert the stop selected in the current route at the place determined by the insertion criterion in step 2. Go to step 2.

Parameters

Initialization criterion (I)

1. The unrouted stop farthest from the depot.
2. The unrouted stop closest to the depot.

Insertion and selection criterion (M, L, E)

The first combination of insertion and selection criterion is given by

$$\begin{aligned} & \min_{u,i \in N \setminus \{0\}} [\mu_1(t_{iu} + t_{i+1u}) - \mu_2 t_{ii+1}] & (2.5) \\ & \max_{u,i \in N \setminus \{0\}} [\lambda_1 t_{0u} - \lambda_2 (\mu_1(t_{iu} + t_{i+1u}) - \mu_2 t_{ii+1})] \\ & \mu_1, \mu_2 \geq 0, \mu_1 + \mu_2 = 1 \\ & \lambda_1, \lambda_2 \geq 0, \lambda_1 + \lambda_2 = 1 \end{aligned}$$

This combination causes unrouted stop u to be inserted between the successive stops i and $i + 1$ of the route.

The second combination uses the same insertion criterion, but selects the unrouted stop u which gives the minimal travel time of the current route.

$$\begin{aligned}
& \min_{u,i \in N \setminus \{0\}} [\mu_1(t_{iu} + t_{i+1u}) - \mu_2 t_{ii+1}] & (2.6) \\
& \min_{u,i \in N \setminus \{0\}} T \\
& \mu_1, \mu_2 \geq 0, \mu_1 + \mu_2 = 1
\end{aligned}$$

The combinations of weights μ_1 and μ_2 and of weights λ_1 and λ_2 are represented by the five values of parameters M and the six values of parameter L .

For time-window problems, three combinations of insertion and selection criteria are considered.

The first combination of time-oriented insertion and selection criterion takes account of the delay at the beginning of the service time at stop $i + 1$ caused by the insertion of stop u .

$$\begin{aligned}
& \min_{u,i \in N \setminus \{0\}} [\eta_1(\mu_1(t_{iu} + t_{i+1u}) - \mu_2 t_{ii+1}) + \eta_2(b'_{i+1} - b_{i+1})] & (2.7) \\
& \max_{u,i \in N \setminus \{0\}} [\lambda_1 t_{0u} - \lambda_2(\eta_1(\mu_1(t_{iu} + t_{i+1u}) - \mu_2 t_{ii+1}) + \eta_2(b'_{i+1} - b_{i+1}))] \\
& \mu_1, \mu_2 \geq 0, \mu_1 + \mu_2 = 1 \\
& \lambda_1, \lambda_2 \geq 0, \lambda_1 + \lambda_2 = 1 \\
& \eta_1, \eta_2 \geq 0, \eta_1 + \eta_2 = 1
\end{aligned}$$

Here b_{i+1} and b'_{i+1} represent the beginning of service time at stop $i + 1$ before and after the insertion of stop u in the route, respectively.

The second combination is obtained by changing the selection criterion in order to select the stop whose insertion yields the minimal travel time T for the route.

$$\begin{aligned}
& \min_{u,i \in N \setminus \{0\}} [\eta_1(\mu_1(t_{iu} + t_{i+1u}) - \mu_2 t_{ii+1}) + \eta_2(b'_{i+1} - b_{i+1})] & (2.8) \\
& \min_{u,i \in N \setminus \{0\}} T \\
& \mu_1, \mu_2 \geq 0, \mu_1 + \mu_2 = 1 \\
& \lambda_1, \lambda_2 \geq 0, \lambda_1 + \lambda_2 = 1
\end{aligned}$$

The third combination of insertion and selection criterion considers the urgency of servicing the stop u to be inserted. The urgency is expressed as the difference of the latest closing time of stop u , l_{2u} , and its begin of service time, b_u .

$$\begin{aligned}
& \min_{u,i \in N \setminus \{0\}} [\eta_1(\lambda_1(t_{iu} + t_{i+1u}) - \lambda_2 t_{ii+1}) + \eta_2(b'_{i+1} - b_{i+1}) + \eta_3(l_{2u} - b_u)] \\
& \min_{u,i \in N \setminus \{0\}} [\eta_1(\lambda_1(t_{iu} + t_{i+1u}) - \lambda_2 t_{ii+1}) + \eta_2(b'_{i+1} - b_{i+1}) + \eta_3(l_{2u} - b_u)] \\
& \lambda_1, \lambda_2 \geq 0, \lambda_1 + \lambda_2 = 1 & (2.9) \\
& \eta_1, \eta_2, \eta_3 \geq 0, \eta_1 + \eta_2 + \eta_3 = 1
\end{aligned}$$

The combinations of the weights μ_1 and μ_2 , the weights λ_1 and λ_2 , and the weights η_1 , η_2 and η_3 of combinations 2.7, 2.8 and 2.9 are represented by combining the six values of parameter M , the six values of parameter L and the fifteen values of parameter E , respectively.

2.7 Parallel Insertion heuristic

The Parallel Insertion heuristic (PI) develops several routes simultaneously. The routes are initialized by means of seeds. For each unrouted stop, the best insertion place is determined by means of an insertion criterion. The selection criterion selects the unrouted stop to be inserted in one of the routes at its best place.

The traditional disadvantages of the PI heuristic are its myopic character and the fact that the utilization of the vehicles is not a primary concern. The routes obtained tend to be well-separated.

Potvin and Rousseau (1993) present a two-phase parallel insertion heuristic, of which the seeds are generated by performing a sequential insertion heuristic with the same insertion and selection criteria as for the parallel insertion heuristic. This heuristic can also be considered to be a two-phase heuristic because it is a superposition of a sequential and a parallel route-building heuristic. The generation of seeds with a sequential route-building heuristic is a valuable alternative to the circle and the cone covering in case of hard time-windows. The circle- and cone covering methods for seed generation only take capacity constraints into account.

Procedure

- Step 1:* Perform the parallel initialization of the routes.
- Step 2:* For every unrouted stop, determine the best feasible insertion place using the insertion criterion. All places of all seed routes are to be evaluated. If no feasible insertion places can be found for any unrouted stop, then go to step 4.
- Step 3:* Select the stop with the best feasible value for the selection criterion. Insert the stop selected at the place in the route determined by the insertion criterion in step 2. Go to step 2.
- Step 4:* If stops remain unrouted, invoke the post-processor (cfr. appendix B.3) to assign these stops to an existing route or, if not possible, to a new route.

Parameters

Parallel initialization procedure (S)

1. Seeds are generated by means of the two variants of the circle covering method (see appendix B.1.1). Seeds coincide with the location of stops.
2. Seeds are generated by means of the cone covering method with five different load fractions (see appendix B.1.2). Seeds coincide with the location of stops.

Insertion and selection criterion (M, L, E)

The same combinations of selection and insertion criteria as for the SI heuristic are used. Criteria combinations 2.5 and 2.6 are used for problems without time-windows. The remaining combinations 2.7, 2.8 and 2.9 are exclusively used for time-windows problems.

2.8 Parallel Assignment-based Insertion heuristic

The Parallel Assignment-based Insertion heuristic (PAI) is comparable to the PI heuristic, but the insertion and selection criteria are different. In the PAI heuristic, unrouted stops are selected by solving a generalized assignment problem (GAP) and inserted into a seed route by solving a TSP. The order of selection and insertion is inverted, as compared to the PI heuristic.

The advantages and disadvantages of this heuristic are identical to those of the PI heuristic.

The conception of this heuristic is due to Savelsbergh (1990b).

Procedure

- Step 1:* Generate seed stops. Build seed routes by linking the stops to the depot.
- Step 2:* Compute the cost of a feasible assignment of each unrouted stop to each seed route. If no feasible assignment is possible, then stop and go to step 5.
- Step 3:* Select the unrouted stop associated with the best assignment.
- Step 4:* Solve a TSP (see appendix B.2) for inserting the stop selected in step 3 into the route associated with the seed selected in step 3. If the TSP is unsolvable, then discard the stop closest to depot out of the route associated with the seed. Repeat this until a feasible solution is obtained for the TSP. Go to step 2
- Step 5:* If stops remain unrouted, invoke the post-processor (cfr. appendix B.3) to assign these stops to an existing route or, if not possible, to a new route.

Parameters**Parallel initialization procedure (S)**

1. Seeds are generated by means of the two variants of the circle covering method (see appendix B.1.1). Seeds coincide with the location of stops.
2. Seeds are generated by means of the cone covering method with five different load fractions (see appendix B.1.2). Seeds coincide with the location of stops.

Assignment procedure (E, R)

1. The unrouted stop with the largest difference between its best and its second best assignment to different seed points is selected. This stop is assigned to the seed point corresponding with the lowest assignment cost. Assigning all stops to seeds in this way is denoted as the assignment based on the minimal regret function. This procedure for solving the GAP has been proposed by Martello and Toth (1981).

The cost of assigning a stop $i \in N \setminus \{0\}$ to a seed point $k \in N \setminus \{0\}$ is given by

$$\begin{aligned} & \epsilon_1(t_{0i} - t_{0k}) + \epsilon_2 t_{ik} & (2.10) \\ & \epsilon_1, \epsilon_2 \geq 0, \epsilon_1 + \epsilon_2 = 1 \end{aligned}$$

The combinations of the weights ϵ_1 and ϵ_2 are represented by the five values of parameter E .

2. The unrouted stop $i \in N \setminus \{0\}$ is assigned to seed point $k \in N \setminus \{0\}$ if this assignment yields the lowest assignment cost of all possible assignments. The assignment cost for this assignment procedure is extended with a term which takes account of the magnitude of the demand associated with a stop. The rationale is that stops with a high demand are more difficult to assign when most other stops have already been assigned. Hence, priority is given to the assignment of stops with a large demand. In order to take account of this, the assignment cost of stop i to seed point k is augmented by the ratio of the average vehicle capacity \bar{Q} to the demand of stop i .

$$\begin{aligned} & \rho_1(\epsilon_1(t_{0i} - t_{0k}) + \epsilon_2 t_{ik}) + \rho_2 \left(\frac{\bar{Q}}{q_i} \right) & (2.11) \\ & \epsilon_1, \epsilon_2 \geq 0, \epsilon_1 + \epsilon_2 = 1 \\ & \rho_1, \rho_2 \geq 0, \rho_1 + \rho_2 = 1 \end{aligned}$$

The combinations of the weights ϵ_1 and ϵ_2 and those of weights ρ_1 and ρ_2 are represented by the five values of parameter E and the five values of parameter R , respectively.

In the same way, the hardness of time-windows can be taken into account in the expression above. This is not relevant to this research because all time-windows are homogeneous.

2.9 Generalized Assignment heuristic

The Generalized Assignment heuristic (GA) is a two-phase heuristic. During the first phase, the clustering phase, a GAP is solved in order to assign the stops to previously generated seed points. The second phase, the routing phase, consists of building a route sequence with all stops assigned to each seed point. Therefore, a TSP must be solved within each cluster of assigned stops to a

seed point. The implementation proposed solves both the GAP and the TSP heuristically.

A major drawback of this heuristic is inherent in its phased nature. In case of hard time-windows, for example, there is no guarantee that the TSPs are solvable with all stops assigned to a seed point. Only capacity constraints are taken into account for the generation of seed points and for solving the GAP during the first phase.

An advantage is that the routes produced are well-separated.

The generalized assignment heuristic was first proposed by Fisher and Jaikumar (1981). Nygard et al. (1988) adapted the generalized assignment heuristic to the deadline VRP. The effect of human interaction in order to improve the solutions generated by the heuristic is considered by Baker (1992). The use of the heuristic for soft time-windows has been evaluated by Koskodis et al. (1992).

Procedure

Step 1: Generate seed points.

Step 2: Assign the unrouted stops to the seed points by solving the GAP.

Step 3: Improve the initial assignment through exchanging or relocating stops between clusters.

Step 4: Solve a TSP (see appendix B.2) with all stops assigned to a seed point. Do this for each seed point. If a TSP cannot be solved with all stops, the stop closest to the depot is discarded and left unrouted. Repeat this until a feasible solution is obtained for the TSP.

Step 5: If stops remain unrouted, invoke the post-processor (cfr. appendix B.3) to assign these stops to an existing or, if not possible, to a new route.

Parameters

Parallel initialization procedure (S)

1. Seeds are generated by means of the two variants of the circle covering method (see appendix B.1.1). Seeds coincide with the location of stops.
2. Seeds are generated by means of the cone covering method with five different load fractions (see appendix B.1.2). Seeds do not need to coincide with the location of stops.

Assignment procedure (E, R)

The same two assignment procedures as for the PAI heuristic (section 2.8) are used, i.e. the assignment based on the minimal regret function and the direct assignment.

Seed nature (P)

1. Seeds must coincide with the locations of stops.
2. Seeds do not necessarily coincide with the location of stops. This is only possible with the cone covering method.

2.10 Two Phase heuristic

The Two Phase heuristic (TP) solves the VRP by successively performing a sequential and a parallel route-building heuristic. The initialization stops resulting from the sequential route-building phase are used as seed points for the parallel route-building heuristic.

The major advantage of this heuristic is the guarantee for a sufficient number of seed points, even in case of hard side-constraints. Moreover, the advantages of both sequential and parallel route-building methods are exploited. During the sequential phase, routes are built to full capacity utilization. This yields a minimal number of seeds, which are used to construct well-separated routes during the second parallel phase.

The first implementation of a two phase algorithm was proposed by Christofides et al. (1979). The parallel Insertion heuristic of Potvin and Rousseau (1993) mentioned in section 2.7 can also be considered as a two phase heuristic in the way it is presented here.

Procedure

- Step 1:* Initialize a new current route with the initialization criterion.
- Step 2:* Select the unrouted stop, which feasibly satisfies the sequential selection criterion with respect to the current route. If this is not possible and all stops are not routed, then go to step 1. If all stops are routed go to step 4.
- Step 3:* Add the stop selected in step 2 to the current route by solving a TSP (see appendix B.2). Go to step 2.
- Step 4:* Make seed routes by linking all initialization stops with the depot. All other stops are assigned the status unrouted.
- Step 5:* Select the unrouted stop which satisfies the parallel selection criterion. If no stop can be selected, then go to step 7.
- Step 6:* Add the stop selected in step 5 to the route selected by solving a TSP (see appendix B.2).
- Step 7:* If stops remain unrouted, invoke the post-processor (cfr. appendix ref-post) to assign these stops to an existing or, if not possible, to a new route.

Parameters

Initialization criterion (I)

1. The unrouted stop farthest from the depot.
2. The unrouted stop closest to the depot.

Sequential selection criterion (G)

During the sequential route-building phase, stop j is added to the current route of initialization stop $i \in N \setminus \{0\}$ if the following expression is satisfied

$$\begin{aligned} \min_{j \in N \setminus \{0\}} [\gamma_1 t_{0j} + \gamma_2 t_{ij}] \\ \gamma_1, \gamma_2 \geq 0, \gamma_1 + \gamma_2 = 1 \end{aligned} \quad (2.12)$$

The combinations of the weights γ_1 and γ_2 are represented by the five values of parameter G .

Parallel selection criterion (L)

For each stop $j \in N \setminus \{0\}$ the assignment cost to initialization stop $i \in N \setminus \{0\}$ is computed by means of the expression

$$\begin{aligned} \lambda_1(t_{0i} - t_{0k}) + \lambda_2 t_{ik} \\ \lambda_1, \lambda_2 \geq 0, \lambda_1 + \lambda_2 = 1 \end{aligned} \quad (2.13)$$

The unrouted stop j with the largest difference between its best and second best feasible assignment to two different initialization stops is inserted into the route of initialization stop i , which yields the minimal assignment cost for stop j . The combinations of the weights λ_1 and λ_2 are represented by the five values of parameter L .

2.11 Sweep heuristic

The Sweep heuristic (SW) is a two-phase heuristic of which the two phases are not completely separated. The clustering and routing phases are interwoven. Unrouted stops are added to the current route on the basis of their polar angle with respect to the depot. Physically this corresponds to a counterclockwise sweep movement with the depot as central point, starting from and ending in a reference line. Every stop is selected to form the reference point.

Routes formed by the Sweep heuristic tend to be well separated. No stops remain unrouted in the case of an unlimited number of vehicles. A drawback is that the heuristic is highly dependent on the location of the depot for the sweep movement. The sweep heuristic is also very time consuming. The use of each stop to form the reference point for the sweep movement can be inefficient, particularly in the case of a homogeneous vehicle fleet, because a lot of reference points can give equal solutions.

This heuristic was first implemented by Gillet and Miller (1974). In Wren and Holliday (1972) the use of polar angles for ordering stops is also mentioned.

Procedure

- Step 1:* Compute the polar angle of each stop. Rank the stops in ascending order of their polar angle. The first stop of this list is the reference stop.
- Step 2:* The reference stop determines the reference line with the depot. Make a second list of polar angles with the stops rearranged with respect to the reference stop.
- Step 3:* If the second list of polar angles is not empty, then select the next stop and add it to the current route and go to step 4, else the sweep-arm is back at the reference line. In the latter case, all stops are routed and a feasible solution is obtained. If this solution is better than all previous solutions, it is stored temporarily. If the first list of polar angles is not empty, then select the next stop of the first list as the reference stop and go to step 2, else go to step 5.
- Step 4:* Solve a TSP (see appendix B.2) with all stops of the current route. If the TSP is solvable, go to step 3. If not, then terminate the current route without adding the last stop selected. Evaluate improvements by relocating or exchanging stops between the two routes that were last built.
A new current route is started with the last stop selected as the first stop. Go to step 3.
- Step 5:* The final best solution is the solution saved.

Parameters

The implementation of the Sweep heuristic contains no parameters. Consequently, the heuristic is not included in the parametric analyses.

Chapter 3

Analysis of the effect of vehicle-related constraints

This chapter is devoted to the analysis of the influence of the vehicle capacity on the behavior of the eleven initial heuristics. For this purpose, three test sets have been conceived, which were named $G1$, $G2$ and $G3$. These three test sets were obtained by adding side-constraints to the geographical basic set of 60 problems (cfr. appendix A). First a demand of $q_i=10$ units was assigned to each stop $i \in N \setminus \{0\}$ of each problem. For test set $G1$, a homogeneous vehicle capacity of $Q=100$ units is associated with each vehicle. This implies the substitution of constraint 1.4 in section 1.2 by:

$$10 \left(\sum_{i=1}^n \sum_{j=0}^n x_{ijk} \right) \leq 100 \quad k = 1, \dots, K$$

For test set $G2$ the homogeneous vehicle capacity amounts to $Q=50$ units. Constraint 1.4 in section 1.2 becomes:

$$10 \left(\sum_{i=1}^n \sum_{j=0}^n x_{ijk} \right) \leq 50 \quad k = 1, \dots, K$$

A capacity of $Q=200$ units is allocated to each vehicle for the problems of test set $G3$. Constraint 1.4 in section 1.2 is substituted by:

$$10 \left(\sum_{i=1}^n \sum_{j=0}^n x_{ijk} \right) \leq 200 \quad k = 1, \dots, K$$

As mentioned previously, the number of vehicles is unlimited. The minimal number of routes is 10 for test set $G1$, 20 for test set $G2$ and 5 for

test set $G3$.

Below, the parametric analyses are presented, followed by the results of the three test sets $G1$, $G2$ and $G3$ per heuristic.

The second part of this chapter compares the relative performances of each heuristic in the heuristic analysis. The results of the heuristic analyses are discussed separately for the three test sets.

3.1 Parametric analyses for the test sets G1, G2 and G3

The parametric analysis encompasses an AID analysis for each heuristic applied to each problem of each test set. As mentioned in the introductory chapter, the AID uses all solutions of a problem, obtained by combining the values of all parameters of a heuristic exhaustively. It is clear that we are far more interested in the interpretation of the AID results on an aggregate level rather than on the individual problem level.

Next, the AID results are considered for each of the ten heuristics. The Sweep heuristic contains no parameters and is as such not involved in the parametric analyses. The AID results for a heuristic are discussed for each parameter separately (see appendix B.4).

A general remark can be made on the consequences of the use of problems with a deterministic structure. Some problems are pathologic for some heuristics. These pathological problems are a threat to the internal validity of the experiments.

Problems with an important pathological effect are those which disturb the initialization criterion of the sequential route-building heuristics. If the initialization criterion consists of the stop farthest from or closest to the depot, then problems with a large number of stops on the same highest or lowest travel time from the depot can disturb the expected behavior of the initialization criterion. In cases where several stops satisfy the initialization criterion simultaneously, the stop with the lowest rank number in the file of stops is selected. The bias is more substantial if the number of routes is limited. Consequently, the distortion tends to be higher for the sequential heuristics of test set $G3$ and lower for those of test set $G1$.

The problems of the test sets which are pathological for the sequential initialization criterion are those with a *concentric* or a *50 % central* spreading and a *central* depot (see appendix A).

These problems also affect the starting position of the sweep procedure of the cone covering method (cfr. appendix B.1.2). If the largest difference between

the polar angles of two successive stops occurs between more than one pair of stops, the first evaluated pair is selected. The bias caused is in fact negligible.

3.1.1 Sequential Nearest Neighbor heuristic

The results of the AID analyses for the SN heuristic are somewhat biased due to the large number of problems which are pathological for the nearest neighbor criterion. Many test set problems have a large number of equal nearest neighbor pairs. In such case, the pair evaluated first is selected. However, the sequential nature of the SN heuristic keeps the bias most often limited.

The number of solutions per problem obtained by combining all values of the three parameters of the SN heuristic amounts to 18. Applying AID to this set of replications leads to the following findings.

Initialization criterion (I). As mentioned previously, the initialization of the current route with the stop at highest or lowest travel time from depot can be perturbed by some pathological problems.

Initialization by merging the two nearest neighbors in one route is strongly biased in the case of the pathological problems with a great number of nearest neighbors pairs. A large number of problems of each test set are pathological in the previously defined sense (see appendix A).

A significant influence of the initialization criterion is only meaningful if the frequency of initialization (P) (cfr. infra) is the multiple initialization. In that case, every new current route is initialized with the initialization criterion. However, the number of problems of the three test sets for which this interaction is observed is insufficient to make inferences.

Significant differences between the three values of I tend to increase with an increasing number of routes. This can be explained by the increasing number of initializations required. The number of problems where significant differences can be observed, is too limited to base conclusions on, even for test set $G2$.

Frequency of initialization (P). Significant differences between solutions obtained with a single and with a multiple initialization occur only for a very limited number of problems of the three test sets. These problems have no clear common characteristics.

Places to add stops (A). For nearly all problems of the three test sets, no significant difference is observed between the addition of stops at the end of a route and the addition at both ends of a route.

These findings reveal that no parameter of the SN heuristic has a systematic effect on the solution value of the problems of the three test sets. Moreover, the changes in capacity constraints do not induce substantial shifts in the be-

havior of the parameters.

3.1.2 Parallel Nearest Neighbor heuristic

This heuristic is highly perturbed by the pathological problems containing a great number of nearest neighbor pairs. The parallel nature of the heuristic makes it more vulnerable to biased results than a sequential heuristic. In the case of problems with more nearest neighbors, the first pair that is evaluated is selected. The three test sets contain a substantial number of pathologic problems of this kind (see appendix A). Consequently, the internal validity of the AID results is not guaranteed. Moreover, the only parameter of the PN heuristic gives rise to only 8 solutions per problem.

Simultaneous initialization procedure (S). Routes can be initialized by means of the two variants of the circle covering or the five variants of the cone covering in the case of seed generation or by successively merging the two nearest neighbors.

The significant effect of S on the total travel time is very limited in the case of solutions with a few number of large routes (test set $G\beta$). This is mainly due to the small number of routes, and thus the low number of seeds required. Changing the positions of seeds does not significantly alter the compositions of the routes in those situations.

For test sets $G1$ and $G2$, a vast majority of problems are observed for which the differences between the values of S yield significantly different results. For these problems, the circle covering with the ascending radius selection procedure (cfr. appendix B.1.1) results in an overestimation of the number of seeds and gives significantly worse solutions. The initialization procedure through merging the nearest neighbor also gives bad results, probably due to the effect of the pathological problems. The cone covering and the circle covering with the descending radius selection procedure give significantly better results than all other initialization procedures. As far as the cone covering is concerned, no other significant differences were observed between the five different load fractions.

If only the value of S in the best solutions is considered, a dominance of the cone covering with load fraction 0.75 is observed.

From these results it is possible to conclude that the significant effect of the single parameter of the PN heuristic is affected by changes in capacity constraints and by the problem characteristics as well as by the pathological problems.

3.1.3 Sequential Savings heuristic

The exhaustive combination of the four parameters of the SS heuristic gives rise to 270 solutions per problem. AID analyses performed on each problem of the three test sets reveal the following findings.

Initialization criterion (I). As mentioned previously, the expected behavior of the initialization criterion is perturbed for a number of problems. The initialization procedure through merging the stops with the greatest savings value can be perturbed for problems with more than one pair of stops that yields the greatest savings. This has also been observed by Golden (1977).

The results of the three test sets reveal that the differences between the three values of I are significant for most of the problems. Initialization with the stop closest to the depot tends to yield significantly worse results than initialization with the stop farthest from the depot and the initialization through merging the two unrouted greatest savings stops. This occurs mainly for problems of the *clusters* and the *cones* patterns of test sets $G1$ en $G3$. The problems of both patterns have well-separated groups of stops. For test set $G2$ this tendency is observed for other problems too. The insignificant difference between initialization through merging the greatest savings stops and initialization with the stop farthest from depot is evident. The stops which yield the greatest savings value are most often very far from the depot and very close to each other.

Frequency of initialization (P). The significant effect of the initialization criterion I is meaningful if the multiple initialization procedure is used. This interaction is frequently observed. Multiple initialization tends to favor the construction of well-separated routes in the event of problems with a rather homogeneous spread of stops.

In the case of single initialization a new current route is started in the neighborhood of the route just terminated. This does not benefit the formation of separated routes, but avoids the need for additional routes to service the remaining stops spread out between the constructed routes.

Places to add stops (A). The AID results do not permit to draw meaningful conclusions concerning the addition of stops at the end or at both ends of the current route.

Savings criterion (C). The effect of the parameter representing the combinations of the weights of the savings criterion 2.3 is significant for all problems. The significantly better values of C are those which preserve the inequality $\alpha_1 \leq \alpha_2$ and minimize or even neglect the weight α_3 . This implies that the savings value is primarily defined in terms of the proximity of the stops to be merged and secondary in terms of remoteness from the depot. This is observed for the three test sets. Slight shifts among the significantly better values can be observed

only on the individual problem level.

The weights for the savings criterion proposed by Gaskell (1967), Yellow (1970), Eilon et al. (1971), Webb (1972), Mole (1979) and Paessens (1988) correspond with the significantly better values resulting from the AID analyses.

From these results, we can deduce that only the parameter representing the combinations of the weights of the savings criterion has a consistent significant effect and ditto significantly better values. The effect of the three other parameters is most often insignificant and can hardly be related to the problem characteristics or shifts in the capacity constraints.

3.1.4 Parallel Savings heuristic

Combining the values of the two parameters of the PS heuristic gives rise to 120 replications for each problem of the three test sets.

Simultaneous initialization procedure (S). Routes can be initialized by means of the two variants of the circle covering or the five variants of the cone covering in the case of seed generation or by successively merging the two stops yielding the highest savings value.

Some problems have, due to geographical symmetry, several pairs of stops with an equal savings value. This can cause a slight bias especially when the traditional simultaneous initialization procedure without seeds is used. As mentioned previously, the starting position of the cone covering can also be biased for some problems.

The AID analyses reveal that the circle covering is usually significantly worse due to an overestimation of the seeds required. For three problem types, however, the circle covering is not significantly worse than the cone covering. The first type contains some problems with a rather homogeneous spread of stops and the depot located among the stops. For these problems, solutions with a larger number of routes than the theoretically minimum number are not worse a priori. This can be observed mainly for the test sets *G1* and *G3*.

The second type of problems are problems of which the stops are compressed within a narrow area, like those of the *concentric* pattern. The number of circles generated with the circle covering for these problems is not overestimated. This is particularly true for test set *G3*.

The third type of problems are problems with well-separated groups of stops (*clusters* and *cones* patterns). The circle covering has especially been designed by Savelsbergh (1990b) for this kind of problems.

As far as the traditional method without seeds is concerned, one can observe that it yields good results for problems with a homogeneous spread of stops of test set *G1* and for almost all the problems of *G2*. For test set *G3* the solutions obtained with this method are significantly worse. This could be explained by

the fact that for test set $G3$ the number of routes are insufficiently reduced by merging stops iteratively.

The cone covering belongs to the significantly better values of S for almost all problems. A general remark is that the differences between the distinct load fractions fade away for the *concentric* and *50% central* patterns with a depot *central* or *inside*. This is because the stops are compressed within a narrow area. For most problems with a rather homogeneous spread of stops, the cone covering with the lowest load fraction, (0.05), is significantly worse. This causes a positioning of seeds too close to the depot, which does not guarantee good separation of the resulting routes. The cone covering with the highest load fraction (0.95) gives significantly worse results for the majority of problems of test set $G2$. So, in the case of routes with a small number of stops, the seeds should not be too far from the depot.

Savings criterion (C). The effect of the parameter representing the combinations of the weights of savings criterion 2.3 is significant for all problems. The significantly better values of C are comparable to those of the SS heuristic. They contain the combinations of the weights, which stress the importance of the proximity of the stops to merge in spite of the remoteness of both stops from the depot. This can be expressed as $\alpha_1 \leq \alpha_2$ and minimizing or even neglecting the weight α_3 . Again, this is observed for the three test sets. Further refinements within these significantly better values can only be observed on the individual problem level.

The weights for the savings criterion proposed by Clarke and Wright (1964), Yellow (1970), Eilon et al. (1971), Webb (1972), McDonald (1972) and Paesens (1988) can be found among the significantly better values resulting from the AID analyses.

Both parameters of the PS heuristic have a significant effect on the final solution. The problem characteristics and the side-constraints only affect the significantly better values of the simultaneous initialization procedure. The significantly better values of the combinations of the savings criterion weights are consistent.

3.1.5 Generalized Savings heuristic

The single parameter P of the GS heuristic contains only five values. This number of replications is insufficient to make a reliable AID analysis.

The GS heuristic can be affected by pathological problems only during its first step where stops are merged two by two. Problems with a large number of pair of stops with an equal savings value can cause a perturbation of the expected behavior of the heuristic. The order of evaluation is decisive in these cases.

Generalized savings criterion (P). Parameter P represents the combinations of the weights α_1 and α_2 of the generalized savings criterion 2.4.

For all problems of test set $G2$, the effect of P is insignificant. The solutions for the problems of $G2$ contain routes with a small number of stops. Consequently, only a very few number of merging cycles are required to obtain a solution. The average number of stops is higher for the problems of test set $G1$. Hence, for a small number of those problems, the effect of P is significant.

For the problems of test set $G3$, a substantial larger number of merging cycles is required to construct routes with a high number of stops. As a result, the effect of P becomes significant for the majority of problems. The significantly better values of P favor the merging with greater stress on the minimization of the travel time of the resulting route instead of on the maximization of the travel time of each of the routes to be merged.

By considering only the value of P giving the best solution, it can be observed that the combination $\alpha_1 = \alpha_2$ is clearly dominant. This is the original generalized savings criterion conceived by Altinkemer and Gavish (1991).

The number of replications is too limited to make important conclusions on the behavior of the parameter of the GS heuristic.

3.1.6 Sequential Insertion heuristic

Combining the values of the three parameters of the SI heuristic results in 60 replications per problem for each of the three test sets. AID analyses performed on these solutions yield the following findings.

Initialization criterion (I). As mentioned previously, the initialization criterion is biased for a number of problems.

Differences between the initialization of the current route with the stop closest to or farthest from depot are significant for only a minority of problems in the three test sets. No meaningful deductions can be made with these results.

Insertion and selection criterion (M, L) . Parameters M and L represent the combination of weights of the combination 2.5 and 2.6 of insertion and selection criterion. The results of the AID analyses are consistent for the three test sets.

For the insertion criterion, which is the same for both the combinations proposed, significantly worse results are obtained with the combination of weights $\mu_1 = 0$, $\mu_2 = 1$. An insertion criterion with these weights determines the insertion place for all unrouted stops as the largest link between two nodes of the current route.

The combination of the weights $\mu_1 = 0.25$, $\mu_2 = 0.75$ is also significantly worse than the three other combinations for most problems of the *concentric* and the

50 % central pattern. Both patterns have long links inherent in the spread of their stops. Consequently, for these problems the relation $\mu_1 \geq \mu_2$ gives significantly better solutions.

In general, significantly better solutions are obtained by emphasizing the importance of the minimization of the increase in travel time at least as much as the maximization of the travel time between the two successive stops i and $i + 1$ of the current route.

By considering only the combinations of the weights in the best solution, one can observe that the two combinations $\mu_1 = \mu_2 = 0.50$ and $\mu_1 = 0.75, \mu_2 = 0.25$ appear more frequently.

For the selection criterion, two combinations can be distinguished. As far as the selection criterion of combination 2.5 is concerned, the combination $\lambda_1 = 1, \lambda_2 = 0$ gives significantly worse results. This is because stops are selected for insertion in the current route on the basis of maximal travel time to the depot. For a vast majority of the problems, the combination $\lambda_1 = 0.75, \lambda_2 = 0.25$ is also significantly worse for the same reasons. So, the significantly better combinations of the weights select the unrouted stop which causes the minimal increase in travel time due to the insertion rather than the stop at the largest travel time from depot, i.e. $\lambda_1 \leq \lambda_2$. These weight combinations favor the formation of well-separated routes. This corresponds to the findings of Mole and Jameson (1976).

The selection criterion of combination 2.6 is not significantly worse for problems where stops are grouped. For these problems, the selection of the stop nearest to the current route is the stop whose insertion gives the minimal route time of the current route.

These results demonstrate the consistency of the significant effect and the significantly better values of the parameters representing the insertion and selection criteria. The effect of the initialization criterion is most often insignificant and cannot be related to the problem characteristics and/or side-constraints.

3.1.7 Parallel Insertion heuristic

The Parallel Insertion heuristic uses three parameters resulting in 210 replications for each problem of test sets *G1*, *G2* and *G3*.

Parallel initialization procedure (S). The two variants of the circle covering and the five variants of the cone covering can be used for generating seeds. As mentioned previously, the starting position of the cone covering can also be biased for some problems. The results of the AID analyses correspond to those of the PS heuristic (see section 3.1.4).

This implies that the cone covering is significantly better than the circle covering, except for a number of problems belonging to three types: problems with a homogeneous spread of stops and with a depot among the stops, problems

with separated groups of stops, and problems with stops compressed within a narrow area.

For test set *G2* it can also be observed that the solutions obtained with the cone covering with the highest load fraction (1.00) are significantly worse than those with other load fractions.

Another correspondence with the findings of the PS heuristic is the fading of the differences between the load fractions of the cone covering method for the problems of the *concentric* pattern with a depot *central* or *inside*.

The reader is referred to the parametric analysis of the PS heuristic (see section 3.1.4) for further explanations on these observations.

The best solution is most frequently generated by the cone covering with a load fraction ranging from 0.25 to 0.75.

Insertion and selection criterion (M, L). Parameters M and L represent the weights combinations of the combinations 2.5 and 2.6 of the insertion and selection criterion. The AID results are comparable to those of the SI heuristic. The results for the three test sets are very similar too.

As far as the insertion criterion is concerned, significantly worse solutions are obtained with the combination of weights $\mu_1 = 0, \mu_2 = 1$. An insertion criterion with these weights determines the insertion place for all unrouted stops as the largest link between two nodes among all parallel routes. The combinations of the weights $\mu_1 = 0.25, \mu_2 = 0.75$ are significantly worse than the other three combinations for most problems of the *concentric* and the *50 % central* pattern, as well as for the *compressed* pattern with a decentralized depot. The above-mentioned patterns have long links inherent in the spread of their stops. Consequently, for the problems of these patterns the relation $\mu_1 \geq \mu_2$ yields significantly better results.

In general, good combinations of weights for the insertion criterion determine an insertion place by attaching greater importance to a minimization of the increase in travel time due to the insertion than to the maximal remoteness of the two successive stops i and $i + 1$ between which the unrouted stop u will be inserted.

With respect to selection criterion 2.5, the combination $\lambda_1 = 1, \lambda_2 = 0$ gives significantly worse results. This is because stops are selected for insertion in a route on the basis of maximal travel time to the depot. The selection of the unrouted stop by the significantly better combinations of weights is based slightly more on the minimal increase in travel time due to the insertion than on the maximal remoteness of the unrouted stop from the depot.

The specific effect of the selection criterion of combination 2.6 used with the parallel heuristic is that routes are built proportionally. Stops whose insertion in a route yields the minimal travel time of that route, are selected. Consequently, stops in the proximity of a route are favored for selection. So, this selection criterion gives good results for problems with grouped stops, like those of the *concentric* and the *clusters* pattern.

Further refinements between the significantly better combinations of weights of insertion and selection criterion can only be observed on the individual problem level.

All parameters of the PI heuristic have a significant effect on the solution value. The significantly better values for the parameters of the insertion and selection criteria are consistent. Those for the simultaneous initialization procedure are more dependent on the problem characteristics and/or the side-constraints.

3.1.8 Parallel Assignment-based Insertion heuristic

The three parameters of the PAI heuristic give rise to 70 replications for each problem of the three test sets.

Parallel initialization procedure (S). The two variants of the circle covering and the five variants of the cone covering can be used for generating seeds. As mentioned previously, the starting position of the cone covering can also be biased for some problems. The results of the AID analyses are similar to those of the PS and PI heuristic.

The cone covering is significantly better than the circle covering, except for a number of problems essentially belonging to three types: problems with a homogeneous spread of stops and with a depot among the stops, problems with separated groups of stops, and problems with stops compressed within a narrow area.

For test set $G2$ it can also be observed that the solutions obtained with the cone covering with the highest load fraction (1.00) are significantly worse than the other load fractions.

Another similarity with the findings of the PS and PI heuristic is the fading of the differences between the load fractions of the cone covering method for the problems of the *concentric* pattern with a depot *central* or *inside*.

Explanations for all these findings can be found in the above section of the PS heuristic (see section 3.1.4).

The best solution is most frequently obtained with the cone covering with a load fraction of 0.50 and 0.75.

Assignment procedure (E, R). The first part of the assignment procedure consists of the determination of significantly better combinations of the weights of the assignment cost 2.10 of stop i to seed k .

The AID results reveal that the combination $\epsilon_1 = 1, \epsilon_2 = 0$ yields significantly worse solutions than the other four combinations. The combination considered assigns stops only on the basis of a comparable travel time of respectively stop and seed to the depot, without taking account of the proximity of stop and seed. For some problems the combination $\epsilon_1 = 0.75, \epsilon_2 = 0.25$ is also significantly worse. This combination is not significantly worse if stops and seeds are

at an approximately equal travel time from the depot. This happens for some problems of test set *G2* and *G1* with a *concentric* or a *50% central* pattern with a *central* depot. For most problems of *G3* the combination $\epsilon_1 = 0.75, \epsilon_2 = 0.25$ is not significantly worse, because on average the travel time between a stop and a seed is high, due to the limited number of seeds required for the problems of test set *G3*. Consequently, the magnitude of the travel time between a stop and a seed, weighted by ϵ_2 , exceeds largely the magnitude of the term associated with weight ϵ_1 .

In general, the assignment of stop *i* to seed *k* gives significantly better solutions if the proximity of stop and seed is at least as important as the difference in their remoteness from the depot.

Two approaches are possible for assigning stops to the seed routes. The first approach uses the minimal regret function, the second the direct assignment. For the three test sets the same assignment cost is used for both approaches. The AID results are not sufficiently clear to make reliable deductions. However, on the basis of the results one can deduce that the assignment based on the minimal regret function is preferred to the direct assignment if the geographical structure of the problem to be solved does not force the assignment of a stop to a dedicated seed.

3.1.9 Generalized Assignment heuristic

The GA heuristic is driven by four parameters. For the three test sets, the combination of the values of the four parameters gives rise to 120 different solutions for each problem.

Parallel initialization procedure (S). The two variants of the circle covering and the five variants of the cone covering can be used for generating seeds. As mentioned previously, the starting position of the cone covering can also be biased for some problems. The results of the AID analyses are similar to those of the PS, PI and PAI heuristic, described above.

The cone covering is significantly better than the circle covering, except for a number of problems, essentially belonging to three types: problems with a homogeneous spread of stops and with a depot among the stops, problems with separated groups of stops, and problems with stops compressed within a narrow area.

For test set *G2* it can also be observed that the solutions obtained with the cone covering with the lowest (0.05) and the highest load fraction (1.00) are significantly worse than the other load fractions.

In accordance with the findings of the PS, PI and PAI heuristic, the fading of the differences between the load fractions of the cone covering method is also observed for the problems of the *concentric* pattern with a depot *central* or *inside*.

Explanations for all these findings can be found in section 3.1.4 of the PS heuristic.

The best solution is most frequently obtained with the cone covering with a load fraction of 0.50 and 0.75. A load fraction of 0.75 has been proposed by Fisher and Jaikumar (1981), who conceived the cone covering and the GA heuristic.

Assignment procedure (E, R). As in the case of the PAI heuristic, two approaches are available for solving the GAP in order to assign the stops to the seed points: the assignment based on the minimal regret function and the direct assignment. Both approaches use the same assignment cost 2.10. The AID results are similar to those of the PAI heuristic (see section 3.1.8). This implies that the significantly worse solutions are obtained with the combination of the weights $\epsilon_1 = 1, \epsilon_2 = 0$ of the assignment costs. The combination $\epsilon_1 = 0.75, \epsilon_2 = 0.25$ is significantly worse for the same types of problems with specific geographical structures encountered with the PAI heuristic (cfr. section 3.1.8). So, in general, it can be stated that the proximity of stop and seed is fundamental in computing the assignment cost for the three test sets.

As far as the choice between the two approaches for solving the GAP is concerned, no meaningful deductions can be drawn from the results of the AID analyses. The number of problems of the three test sets for which the two approaches yield significantly different results is very small in absolute terms and even in comparison with the number observed for the PAI heuristic.

Seed nature (P). Only the cone covering permits the generation of seeds which do not necessarily coincide with stops. Differences between solutions obtained with seed points and with seed customers are seldom significant. The characteristics of the problems of the three test sets for which the difference is significant do not permit us to make meaningful deductions on a level higher than that of the individual problem.

Consistency of the significant effect and the significantly better values is only observed for the parameter representing the combinations of the weights of the assignment cost. The significant effect of the simultaneous initialization procedure can be related to the problem characteristics and/or the side-constraints. The effect of the assignment procedure is most often insignificant.

3.1.10 Two Phase heuristic

The combination of the values of the three parameters of the TP heuristic gives rise to 50 different solutions for each problem of the three test sets.

Initialization criterion (I). The expected behavior of the initialization criterion is biased by the above-mentioned pathological problems.

A new current route can be started by the stop closest to or farthest from the depot during the first sequential phase of the heuristic.

The AID results for the three test sets reveal that initialization with the stop farthest from the depot yields significantly better results, especially for problems with a homogeneous spread of stops (*uniform* and *compressed* patterns).

The use of the farthest stop as initialization stop in the second phase can be considered to be to the use of the cone covering method with a high load fraction.

Potvin and Rousseau (1993) also used the farthest stop as initialization stop in the first phase of their two phase/parallel insertion heuristic.

Sequential selection criterion (G). Unrouted stops are added to the current route during the first phase of the heuristic if they satisfy the sequential selection criterion 2.12. The combinations of the weights γ_1 and γ_2 of this criterion are represented by G .

The AID results show that the significantly better solutions are obtained with the combinations for which $\gamma_1 \leq \gamma_2$. This implies that stops are added to the current route on the basis of a criterion in which the proximity of the unrouted stop j to the initialization stop i is at least as important as that of stop j to the depot.

The weights proposed by Christofides et al. (1979) for the selection criterion confirm our findings.

Parallel selection criterion (L). The parallel selection criterion 2.13 can be considered as a cost for inserting the unrouted stop j into the route of the initialization stop i . The actual insertion is made by means of a minimal regret function.

The results reveal that the combination of weights $\lambda_1 = 1, \lambda_2 = 0$ gives significantly worse solutions. This combination preferably inserts a stop which is at approximately the same travel time from the depot as an initialization stop, without considering the proximity between the stop and the initialization stop.

The combination $\lambda_1 = 0.75, \lambda_2 = 0.25$ is not significantly worse if stops and seeds are at approximately equal travel time from the depot. This happens for some problems of test sets $G2$ and $G1$ with a *concentric* or a *50% central* pattern, both with a *central* depot. For most problems of test set $G2$, the combination $\lambda_1 = 0.75, \lambda_2 = 0.25$ is significantly worse because on average the travel time between a stop and a seed is limited, due to the high number of seeds required. Consequently, the magnitude of the travel time between a stop and a seed, weighted by λ_2 , is relatively low compared with the magnitude of the term associated with weight λ_1 .

In general, the insertion of stop j into the route of initialization stop i gives significantly better solutions if the proximity of stop and initialization stop is at least as important as the difference in remoteness of both stops from the depot.

These findings are also confirmed by the weights for this criterion proposed

by Christofides et al. (1979).

The significant effect and the significantly better values of the parameters representing the sequential and parallel selection criteria are consistent. The significance of the effect of the initialization criterion is not consistent and the significantly better values are partially dependent on the problem characteristics and/or side-constraints.

3.1.11 Conclusions of the parametric analysis

Based on the parametric analysis, two groups of parameters of the initial heuristics can be distinguished.

The first group contains the parameters representing the combinations of weights of the selection criteria of the initial heuristics. Selection criteria are the savings criteria of the savings heuristics (SS and PS), the insertion and selection criteria of the insertion heuristics (SI and PI), the assignment costs of the PAI and GA heuristics, and the two selection criteria of the TP heuristic. All these parameters have a significant effect on the total travel time. Moreover, the significantly better values for these parameters are consistent throughout the three test sets. Minor shifts in the significantly better values are mainly induced by different problem characteristics and/or side-constraints.

The second group of parameters contains the remaining parameters. The parameters of the non-sequential heuristics are significant for far more problems than those of the sequential heuristics. This applies particularly to parameters related with the initialization of routes. This can be explained by the fact that the primary criterion of sequential heuristics is to fill up the routes to maximal capacity utilization. The effect of the sequential initialization criterion is not significant for all problems. In addition, the significantly better values occurring can hardly be related to specific problem characteristics and/or side-constraints. The effect of the simultaneous initialization procedure for non-sequential heuristics is significant for almost all problems. Moreover, the significantly better values for these parameters can be related to the problem characteristics and/or side-constraints.

3.2 Heuristic analysis

The heuristic analysis evaluates the relative performances of the eleven initial heuristics. The statistical analysis required to perform this analysis is the Friedman test (see appendix B.5). Only the total travel time of the best solution of

a heuristic obtained for a problem is considered. Consequently, eleven solutions are compared for each problem.

The statistical analysis is performed on an aggregate level. This implies that the problems are grouped, based on the three geographical criteria: spreading of stops, grouping of stops and location of the depot (see section 1.3).

3.2.1 Heuristic analysis for test set G1

All problems. The GA heuristic gives significantly better results than all other heuristics for the problems of *G1*. This dominance is undoubtedly related to the ability of the GA heuristic to keep the routes well-separated for problems with only capacity constraints.

The solutions of the PN heuristic are significantly worse than all other heuristics. The pathological problems are probably not strange to this. Further refinements of these findings for the problem groups corresponding to the geographical criteria yield some interesting observations. They are discussed below.

Depot location. The GA heuristic gives the better solutions for the three depot locations. Additionally, it has been observed that the more decentralized the depot, the better the performance of the sequential heuristics relative to that of the non-sequential heuristics. This is partially due to the depot-dependent cone covering method used by the non-sequential heuristics for generating seeds. The cone covering method turns out to perform rather badly with a depot *outside*. The depot-independent circle covering procedure offers no good alternative for this depot pattern because the number of seeds is usually overestimated. The SW heuristic also generates bad results for the depot *outside* due to its depot-dependency; the sweep mechanism of the SW heuristic is comparable to that used for creating cones of stops in the cone covering method.

For a *central* depot and a depot *inside* most non-sequential heuristics give good results compared to the sequential heuristics. Non-sequential heuristics are more appropriate for keeping the routes separated in these situations.

Grouping patterns. A general remark about the different grouping patterns is that their results are far more difficult to interpret in comparison to those of the spreading pattern. This is due to the fact that the impact of the grouping patterns on the total geographical structure is less important than that of the spreading patterns (see appendix A).

The results for the *clusters* and the *cones* patterns are worth mentioning. Both patterns contain well-separated groups of ten stops. Ideally, each such group of stops can be serviced by exactly one vehicle.

As far as the problems of the *clusters* pattern is concerned, the majority of heuristics is able to find the best solutions. Only the performance of the SW heuristic is significantly worse, which is mainly caused by its bad performance for

the problems of the *clusters* pattern with a decentralized depot. As mentioned in the parametric analysis, heuristics with seed points use the depot-independent circle covering method for generating seeds for most problems of this pattern. The results for the problems of the *cones* pattern are not so clear because the separation of the groups of stops is not so hard in comparison to those of the *clusters* pattern.

Nevertheless, better results for both patterns are obtained with the SI heuristic, while the solutions of the SW, GS and PN heuristics are significantly worse.

Spreading patterns. For problems with a rather homogeneous spread of stops, it can be observed that the GA heuristic gives the better solutions. For the *uniform* pattern, the PS, PI and PAI heuristics are not significantly worse. For the *compressed* pattern, the PS and SI heuristics are not significantly worse than the GA heuristic. These results confirm the tendency that the ability of non-sequential heuristics to keep the routes separated as much as possible gives very good results for problems with a rather homogeneous spread of stops.

Sequential heuristics tend to produce better solutions for the *50% central* pattern, due to their ability to minimize the links between central stops and peripheral stops. The solutions of the SI and SS heuristics are very good, but those of the GA heuristic are not significantly worse.

The fact that the stops of the *concentric* pattern are compressed within a narrow area, results in limited differences between the solutions of the heuristics. Moreover, some problems of this pattern are pathological for some heuristics. Nevertheless, the GA heuristic gives significantly better solutions than the PN, SN, GS, PAI and SW heuristic.

Table 3.1 presents an overview of the relative performances of the eleven initial heuristics for test set *G1*. The geographical pattern categories distinguished in the table may be considered as the most important ones. The category of the homogeneous spread aggregates the *uniform* and the *compressed* results, except the results of the problems with a *clusters* grouping.

The results of the *clusters* pattern, together with those of the *cones* pattern are represented by the *clusters* category. The category containing the problems with central and peripheral stops contains the results of the *50% central* and the *concentric* patterns.

Comparison of the computing times of the eleven heuristics reveal that the sequential heuristics are less time consuming than their parallel alternatives. The computing times are mostly far below the 10 seconds. The run time of the GA heuristic is the shortest in every respect.

Exceptionally high is the time required by the SW, TP and GS heuristics, due to the large numbers of TSPs to be solved in the three heuristics. We do not exclude that more efficient and time-saving implementations are possible. However, for this research we are not interested in absolute computing times. The

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	+	--	--	+	--
PN	--	--	--	--	-	--
SS	--	+	+	-+	++	++
PS	+	+	-	+	++	-
GS	--	--	--	--	+	-
SI	--	++	+	-+	++	++
PI	+	++	-	-+	+	+
PAI	+	-	--	-	-	-
GA	++	-	++	++	+	+
TP	--	++	-+	+	++	-
SW	+	--	--	--	--	--

Table 3.1: Summary of the relative performances of the 11 heuristics for test set G1 ($q = 10, Q = 100$). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

times are used only for comparative purposes.

3.2.2 Heuristic analysis for test set G2

All problems. The SI heuristic gives significantly better results than any other heuristic, except the SS heuristic for the problems of $G2$. The dominance of both sequential heuristics can be explained by the limited maximal number of five stops in a route. Consequently, the routes are sufficiently short for the sequential heuristics not to make them overlap too much with their objective of filling routes to full capacity. Buxey (1979) also noticed that the quality of sequential heuristics increases if the routes are less overlapping.

The bad performance of the non-sequential heuristics is probably related to the large number of seeds required for the problems of set $G2$. This makes it more difficult to assign the stops to seeds or seed routes.

The solutions of the GS heuristic are significantly worse than those of all other heuristics. The GS heuristic cannot make the appropriate mergers of routes in order to obtain five stops per route.

Depot location. The solutions of the SI heuristic improve as the depot gets more decentralized. For the depot *outside*, the SI heuristics performs significantly better than all other heuristics.

If the depot lies *inside*, the solutions of the SS, PS and TP heuristics are not significantly worse than those of the SI heuristic.

The solutions of the SS, PS and SI are only significantly better than those of the GS, SW, PN and SN heuristics for the *central* depot pattern.

Grouping patterns. As mentioned previously, the results of the grouping patterns are difficult to interpret. However, the results confirm the dominance of the SI heuristic for almost all grouping patterns. The SS and the TP heuristics also give good results for some of these patterns. The *clusters* and *cones* patterns are dominated by the sequential heuristics SI and SS.

The dominance of the sequential heuristics for problems with grouped stops has also been observed by Potvin and Rousseau (1993).

Spreading patterns. The SI and SS heuristic yield better solutions for most spreading patterns. Both heuristics are significantly better than all other heuristics for the *compressed* and the *50 % central* pattern. The sequential heuristics already proved in test set *G1* to be appropriate for the latter pattern due to their ability to minimize links between central and peripheral stops.

The *concentric* pattern does not induce a large number of significant differences between the heuristics, which is consonant with the results of test set *G1*. This is caused by the narrow area within which the stops are compressed. Nevertheless, the sequential heuristics give better solutions for this pattern too.

Better solutions for the *uniform* pattern are obtained with the SI heuristic. The solutions of the PS and TP heuristics are not significantly worse. So, even for the problems with a homogeneous spread of stops, sequential heuristics, in particular the SI heuristic, performs better than almost all non-sequential heuristics. As mentioned previously, this is mainly due to the limited number of stops per route for the problems of testset *G2*.

Again, it can be noticed that the savings criterion is appropriate for problems with a homogeneous spread of stops.

Table 3.2 summarizes the relative performances of the eleven initial heuristics for the principal geographical pattern categories of test set *G2*.

The average computing time is considerably shorter than it was for test set *G1*. Again, the time required by the SW, GS and TP heuristic are disproportionately longer compared to that of the other heuristics. The PI heuristic also requires a longer computing time than most other heuristics, excluding the SW, GS and TP because a large number of insertion places has to be evaluated at each iteration.

The GA, SN, SS and SI heuristics remain very fast.

3.2.3 Heuristic analysis for test set G3

All problems. The results of the comparison of the eleven heuristics for the

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	-+	-+	--	-+	+
PN	--	--	--	--	--	--
SS	+	+	++	+	-+	++
PS	++	-+	-+	+	--	-+
GS	--	--	--	--	--	--
SI	++	+	++	++	++	++
PI	-	-	-	--	--	-
PAI	+	--	--	--	--	--
GA	+	++	+	+	--	-+
TP	-+	+	+	+	++	+
SW	-	-	-+	--	--	--

Table 3.2: Summary of the relative performances of the 11 heuristics for test set G2 ($q = 10, Q = 50$). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

problems of test set $G3$ reveal that the GA and PI heuristics give significantly better results than all other heuristics. The large number of stops per route, maximally twenty, favors the non-sequential heuristics, in particular GA and PI. Sequential heuristics cannot prevent the long routes from overlapping.

Depot location. The depot location does not considerably affect the dominance of the non-sequential heuristics.

The better solutions for a *central* depot are obtained with the GA, PI, SS and SW heuristics. The good performance of the SW heuristic for problems with long routes has also been observed by Solomon (1987) and Paessens (1988). This heuristic is specially designed for keeping the routes well-separated.

For both other depot patterns, a large number of non-sequential heuristics with the GA and PI heuristics on top, give good solutions. The SI is the only sequential heuristic which yields no significantly worse solutions than the non-sequential heuristics for the depot *inside* and *outside* patterns.

Grouping patterns. As mentioned previously, the results for the distinct grouping patterns are more difficult to interpret in comparison with those of the spreading patterns.

Nevertheless, it has to be noticed that the sequential heuristics and more precisely the SS and SI heuristics, give significantly better solution for the problems of the *clusters* pattern, than all other heuristics, the PI heuristic excluded. Sequential heuristics are appropriate for making routes with stops of more than one cluster. For the *cones* pattern, the groups of stops are less separated. Although

the GA heuristic gives better solutions for the problems of this pattern, it is not significantly better than the SI, SS and SW heuristic.

The GA and PI heuristics give good solutions for the others patterns.

Spreading patterns. The results for the four spreading patterns are more or less comparable. The GA and PI heuristic give better solutions for all patterns.

For the *uniform* pattern, the SW heuristic is not significantly worse because of its ability to keep the long routes well-separated.

The SI and PS heuristics are not significantly worse than the GA and the PI heuristic for the *compressed* pattern. The SI heuristics proved to give good solutions for this pattern in all three test sets.

The problems of the *50 % central* pattern are normally better solved with a sequential heuristic due to its ability to minimize the number of links between central and peripheral stops. In the case of test set *G3*, however, this ability becomes irrelevant because of the limited number of routes. Consequently, the better solutions for this pattern are obtained with the GA heuristic, but the solutions of the PI and SS heuristic are not significantly worse.

As far as the *concentric* pattern is concerned, the results are in correspondence with those of the previous two test sets. This implies that no segregation can be made between sequential and non-sequential heuristics, and that the number of significant differences between heuristics is rather limited due to the compression of stops within a narrow area.

Table 3.3 summarizes the relative performances of the eleven initial heuristics for the principal geographical pattern categories of test set *G3*.

With respect to the computing times, we can observe that they are substantially longer than those of test sets *G1* and *G2*. The increase is more than linear. The time required for the GS, TP and SW heuristics is much longer than for the other heuristics. For the SW heuristic, Gillet and Miller (1974) noticed that the computing time grows quadratically with the number of stops in a route, the total number of stops remaining equal.

Again, the GA heuristic requires the least computing time.

To conclude the heuristic analyses for test sets *G1*, *G2* and *G3* the most important observations are gathered.

The relative behavior of the eleven heuristics is affected by the geographic structure of the problem as well as by changes in the vehicle-related side-constraints. Some clear tendencies can be observed.

A major shift can be observed if the average number of stops per route changes. For a maximum of five stops per route the sequential heuristics dominate. If the number of stops per route increases, the non-sequential heuristics tend to

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	-	--	--	-	--
PN	--	--	--	--	-+	-
SS	--	++	++	--	+	+
PS	-+	--	-	+	-+	+
GS	--	--	-	--	-	--
SI	-	+	+	-+	-+	+
PI	++	+	++	++	+	++
PAI	--	-	--	-+	--	+-
GA	++	-+	++	++	-+	++
TP	-+	--	-	++	--	--
SW	++	-	++	++	-	--

Table 3.3: Summary of the relative performances of the 11 heuristics for test set G3 ($q = 10, Q = 200$). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

generate better solutions due to their ability to keep the routes separated. This is also the reason why non-sequential heuristics give better results for problems with a rather homogeneous spread of stops.

The more the depot is decentralized, the better the sequential heuristics perform in comparison to the non-sequential heuristics. With the sequential heuristics, good solutions are also obtained for the 50 % *central* pattern due to their ability to limit the number of links between central and peripheral stops. In addition, good solutions for problems with clustered stops are mostly obtained with sequential heuristics.

By considering each heuristic separately, the following observations can be made.

The SN heuristic offers satisfactory results for the 50 % *central* pattern and for problem with clustered stops, where each cluster can be served by exactly one vehicle. However, the effect of the pathological problems must be taken into account while considering the performances of the SN and PN heuristic. The latter heuristic is characterised by overall bad solutions.

In general, the solutions obtained with the GS heuristic are also poor. The GS heuristic is very bad for problems with a solution containing very few, and particularly an odd number of stops per route ($G2$).

The SS heuristic is a typical sequential heuristic, in so far that its results are excellent for problems with a sufficiently small number of stops per route, for problems with a decentralized depot, for problems with clustered stops and/or for problems with a mix of central and peripheral stops.

The SI heuristic performs at least as good as the SS heuristic for the above-mentioned problems. The results of both heuristics are poor for problems with a large number of stops per route (test set *G3*). For these kinds of problems the GA and PI heuristics stand out. The GA heuristic also outperforms all heuristics for the problems of set *G1*. Additionally, this heuristic is very attractive due to its extremely short computing time.

The PS, PAI and TP are moderate, in so far that there are no particular problems for which these heuristics are excellent or really bad. However, the PS heuristic produces good solutions for problems with a homogeneous spread of stops.

Finally, the SW heuristic gives excellent results for problems with a large number of stops per route. Due to the depot-dependent nature of the heuristic, good results are obtained for problems for which the sweeping results in meaningful routes.

Chapter 4

Analysis of the effect of Customer-related constraints

This chapter is dedicated to the analysis of the effect of two types of customer-related side-constraints on the behavior of the eleven initial heuristics and their parameters.

The first customer-related constraint is that of the mixed pick-up and delivery. In order to analyse the effect of mixed pick-ups and deliveries, a new test set was created. This test set, called *PI* by convention, was constructed by using the 60 problems of test set *G1* with $q=10$ and $Q=100$. Half the stops of each problem of *G1* were transformed to pick-up points. This implies that each problem contains 50 pick-up and 50 delivery points. Each stop $i \in N \setminus \{0\}$ has a demand of $q_i = 10$. A systematic rule was used for determining the pick-up and delivery points among the stops. The rule alternates the pick-up and delivery points for all 60 problems of the test set: Stop 1 in the file of stops is a pick-up point, stop 2 a delivery point, stop 3 a pick-up point, and so on...

All deliveries do not need to be done before pick-ups in a route; pick-ups can be done between deliveries. Some applications require that all deliveries must be performed before all the pick-ups in a route (see Thangiah et al. (1994), Goetschalckx and Jacobs-Blecha (1986)).

Combining pick-ups and deliveries in a single route imposes a route-sequence on the route. If a vehicle of $Q = 100$ is filled to full capacity, then it can service at most 10 pick-up and 10 delivery points in a single route. In that case the first stop had to be a delivery point, while the last stop is obligatory a pick-up point. Moreover, there are only a limited number of ways to combine pick-up and delivery points in a feasible route-sequence.

Preferably, the results of the analyses for test set *P1* must be compared with these of *G3*. Both test sets have solutions for their problems containing at least five routes. Sometimes, comparisons with the results of test set *G1* are appropriate too.

The second customer-related constraint is the heterogeneous demand of the customers. For analyzing the effects of this constraint, a test set called *P2* was conceived. This test set is also based on test set *G1*. For each problem of *G1*, the demand of $q=10$ units for the 100 stops was replaced by four different demands $q=4, 8, 12$ and 16 units. Consequently, four demand-groups of 25 stops each resulted. The total demand for all stops remains 1000 units, which implies that at least 10 vehicles are required. Therefore, comparing the results of the analyses of test set *P2* with these of test set *G1* is appropriate.

The systematic rule used for assigning the heterogeneous demands to the stops allocates sequentially the four different demands of 4, 8, 12 and 16 to every set of four successive stops in the file of stops.

4.1 Parametric analyses for the test set P1

In the following, the results of the AID analysis for the problems of test set *P1* are presented for ten out of the eleven initial heuristics. The SW heuristic contains no parameters and is as such not involved in the parametric analyses.

As in the previous test sets, a general remark concerns the pathological problems which are a threat for the internal validity of the experiments. Again, problems with an important pathological effect are those which disturb the initialization criterion of the sequential route-building heuristics. Referring to the problems of the test sets, the pathological ones are those with a *concentric* or a *50 % central* spreading and a *central* depot (see appendix A).

These problems also affect the starting position of the sweep procedure of the cone covering method (cfr. appendix B.1.2). If the largest difference between the polar angles of two successive stops occurs between more than one pair of stops, the first evaluated pair is selected. Nevertheless, the bias caused is almost negligible.

The heuristic-specific pathological problems for test sets P_1 and P_2 are the same as those for the previous three test sets because the pathological nature is only due to the geographical characteristics of the problems.

4.1.1 Sequential Nearest Neighbor heuristic

The results of the AID analyses for the SN heuristic are somewhat biased due to the large number of problems which are pathological for the nearest neighbor

criterion. As mentioned previously, the sequential nature of the SN heuristic keeps this bias limited.

For each problem 18 replications result from combining all parameter values. The AID analyses give the following findings for the three parameters.

Initialization criterion (I). The results observed correspond entirely with those observed for test set $G3$. Hence, the number of problems where significant differences between the three values of I can be observed is too limited to base conclusions on.

If only the value of I in the best solution is considered, without taking account of any significance, a preference for the initialization with the stop farthest from the depot and the initialization through merging the two nearest neighbors is observed.

Frequency of initialization (P). The significant effect of P is even less than for test set $G3$ and $G1$. This could be caused by the route sequence imposed by the mixed pick-up and delivery. The route sequence is more determined by the order of pick-ups and deliveries than by the minimal travel time. Consequently, the multiple initialization cannot guarantee well-separated routes anymore, as it was the case for the previous three test sets.

If only the value of P in the best solution is considered, a dominance of the multiple initialization procedure is observed.

Places to add stops (A). The behavior of this parameter is different from that observed for $G1$ and $G3$. For a number of problems, the addition of stops at both ends of a route is significantly better than only at the end of a route, due to the mixed pick-ups and deliveries. This is the case, particularly for problems with a *clusters* pattern. For this pattern, two insertion places give more certainty to have the stops of a cluster as much as possible in the same route.

For almost all problems, the addition of stops at both ends of a route yields the best solution.

The addition of stops at both ends of a route partly relaxes the route-sequence imposed by mixed pick-ups and deliveries.

These results show that mixed pick-ups and deliveries do have an effect on the behavior of the parameters of the SN heuristic.

4.1.2 Parallel Nearest Neighbor heuristic

The numerous pathological problems containing a great number of nearest neighbor pairs together with the limited number of 8 replications result in unreliable AID analyses. This can be considered as a threat for the internal validity of the experiments.

Simultaneous initialization procedure (S). The significant effect of S on the total travel time is very limited for the problems of test set *P1*. This has also been observed for the problems of test set *G3*. A reason for this phenomenon can be sought in the small number of routes which require only a few seeds. Additionally, there is the effect of the insufficient number of replications and the high number of pathological problems. These findings give no evidence on an effect of mixed pick-ups and deliveries on the single parameter of the PN heuristic.

4.1.3 Sequential Savings heuristic

The exhaustive combination of the four parameters of the SS heuristic give rise to 270 solutions per problem of test set *P1*.

Initialization criterion (I). As observed for test set *G3*, the effect of parameter I is significant for a substantial number of problems. For these problems, the initialization with the stop closest to the depot yields significantly worse results than the initialization with the stop farthest from the depot and the initialization through merging the two unrouted greatest savings stops. The absence of significant difference between the initialization with the farthest stop from the depot and the initialization through merging has been explained for the parametric analysis of the SS heuristic of the previous test sets (see section 3.1.3).

By considering only the value of I in the best solution, it is observed that the initialization with the stop closest to the depot almost never gives the best solution.

Frequency of initialization (P). The significant effect of the initialization criterion is meaningful if the multiple initialization procedure is used. The results observed do not confirm this interaction and are even more diffuse in comparison to those of test set *G3*. The fading of the significant character of the difference between the single and the multiple initialization is probably induced by the mixed pick-ups and deliveries. The route-sequence imposed by the pick-ups and deliveries prevails on the minimization of the travel time and the shape of the routes. Consequently, the multiple initialization procedure cannot guarantee better separated routes than the single initialization in these cases.

Places to add stops (A). The effect of mixed pick-ups and deliveries is perceptible for this parameter. For a number of problems, the addition of stops at both ends of the current route gives significantly better solutions than only at the end of that route. Particularly, this is the case for problems with a *clusters* pattern. For this pattern, two insertion places offer more certainty for having the stops of a cluster in the same route.

By considering only the value of A contained in the best solution, it can be observed that the addition of stops at two ends of a route is clearly dominant. The use of two insertion places instead of one partly relaxes the binding nature of the route sequence imposed.

Savings criterion (C). The results observed for the parameter representing the combinations of the weights of the savings criterion 2.3 are similar with these of the three previous test sets.

The significantly better values of C preserve the inequality $\alpha_1 \leq \alpha_2$ and minimize or even neglect the weight α_3 . This implies that the savings value is primarily defined in terms of proximity of the stops to be merged and secondary in terms of remoteness from the depot.

Only on the individual problem level slight shifts can be observed among the significantly better values.

The findings for the SS heuristic demonstrate that there is an effect of mixed pick-ups and deliveries on the behavior of some of its parameters, particularly on that of the parameter representing the number of places to add stops to a route.

4.1.4 Parallel Savings heuristic

The PS heuristic with its three parameters gives rise to 120 replications per problem of test set $P1$.

Simultaneous initialization procedure (S). The behavior of parameter S is not visibly affected by mixed pick-ups and deliveries. The results corresponds entirely with these of test set $G3$.

The circle covering is significantly worse for all but three specific types of problems, described in section 3.1.4.

For most problems, the initialization through merging the two stops with the greatest savings gives significantly worse solutions. Probably, this initialization procedure reduces insufficiently the number of routes for the problems considered.

The values of S representing the cone covering can be found in the significantly better values of S for all problems. Only the cone covering with the lowest load fraction (0.05) gives significantly worse results than those with the other load fractions. The lowest load fraction tends to position the seeds too close to the depot. This has a negative effect on the separation of routes.

Savings criterion (C). The significance structure of the parameter representing the combinations of the weights of the savings criterion 2.3 corresponds to that of all previous test sets.

The effect of C is significant for all problems. The significantly better values contain the combinations of the weights, which stress the importance of the proximity of the stops to merge more than the remoteness of both stops from the depot. In terms of weights, this can be expressed as $\alpha_1 \leq \alpha_2$ and minimizing or even neglecting the weight α_3 .

These results demonstrate that the effect of mixed pick-ups and deliveries is almost invisible on the parameters of the PS heuristic.

4.1.5 Generalized Savings heuristic

The 5 replications emerging from the single parameter of the GS heuristic make the results of the AID analyses unreliable.

Generalized Savings criterion (P). The effect of parameter P , representing the combinations of the weights α_1 and α_2 of the generalized savings criterion 2.4, is only significant for about one third of the test set problems. This confirms the thesis formulated for the previous test sets, saying that the significance of P increases with an increasing number of stops per route, due to the larger number of mergers required.

The significantly better values of P realize the merging by stressing the minimization of the travel time of the resulting route instead of the maximization of the sum of the travel times of the routes to be merged.

If only the value of P is considered for the best solution, it can be observed that the combination $\alpha_1 = \alpha_2$ is dominant.

These findings permit us to conclude that there is no clear effect of mixed pick-ups and deliveries on the behavior of the parameter of the GS heuristic.

4.1.6 Sequential Insertion heuristic

Combining the values of the three parameters of the SI heuristic results in 60 replications per problem of test set *P1*.

Initialization criterion (I). Differences between the initialization of the current route with the stop closest to or farthest from the depot are significant for only a minority of problems of this test set. This has already been observed for test set *G3* and is probably a consequence of the reduced number of routes, and hence, initializations required.

Insertion and selection criterion (M, L). Parameters M and L represent the combination of weights of the combinations 2.5 and 2.6 of insertion

and selection criterion. The results of the AID analyses of this test set confirm entirely those of all previous test sets (see section 3.1.6). In general, the significantly better solutions are obtained by emphasizing the minimization of the increase in travel time at least as much as the travel time between the two successive stops i and $i + 1$ of the current route between which the unrouted stop u could be inserted ($\mu_1 \geq \mu_2$).

The significantly better combinations of the weights λ_1 and λ_2 of selection criterion 2.5 favor the selection of the unrouted stop which causes the minimal increase in travel time after the insertion rather than the stop at the largest travel time from the depot, i.e. $\lambda_1 \leq \lambda_2$.

Selection criterion 2.6 is not significantly worse for problems where stops are grouped. For these problems, the selection of the stop nearest to the current route is the stop whose insertion gives the minimal route time.

Minor shifts among these significantly better values of insertion and selection criterion can only be observed on the individual problem level.

These findings suggest that there is hardly an effect of mixed pick-ups and deliveries on the behavior of the parameters of the SI heuristic.

4.1.7 Parallel Insertion heuristic

Combining the three values of the parameters of the PI heuristic results in 210 replications for each problem of test set *P1*.

Parallel initialization procedure (S). The AID analyses reveal that the behavior of the parameter representing the parallel initialization procedure is similar to that of test set *G3* and to that of the PS heuristic for this test set. This implies that the circle covering is not significantly worse for the three specific types of problems, described in section 3.1.7. As far as the cone covering is concerned, the lowest load fraction (0.05) is also significantly worse than the other load fractions.

Insertion and selection criterion (M, L). Parameters M and L represent the combination of weights of the combinations 2.5 and 2.6 of insertion and selection criterion. The AID results observed are comparable to those of the PI heuristic in the previous test sets (see section 3.1.7).

The better combinations of weights for the insertion criterion determine an insertion place by emphasizing the minimization of the increase in travel time more than the travel time between the two successive stops i and $i + 1$ of the route between which the unrouted stop u could be inserted ($\mu_1 \geq \mu_2$).

The significantly better combinations of the weights of selection criterion 2.5 result in the selection of the unrouted stop which is slightly more based on the

minimal increase in travel time due to its insertion than on its remoteness from the depot.

The specific effect of the selection criterion of combination 2.6 used with the parallel heuristic is that routes are built proportionally. The stop whose insertion in a route yields the minimal travel time of that route, is selected. Consequently, stops in the proximity of a route are favored for selection. Selection criterion 2.6 performs well for problems with grouped stops, like problems with a *concentric* pattern, where stops are compressed on a narrow area.

Further refinements between the significantly better combinations of weights of insertion and selection criterion can only be observed on the individual problem level.

Based on these findings we can conclude that the behavior of the parameters of the PI heuristic is not visibly affected by the mixed pick-ups and deliveries.

4.1.8 Parallel Assignment-based Insertion heuristic

The three parameters of the PAI heuristic give rise to 70 replications for each problem of test sets *P1*.

Parallel initialization procedure (S). The behavior of the parameter representing the parallel initialization procedure is entirely similar to that of all previous heuristics of this test set and to that of the PAI of test set *G3* (see section 3.1.8). The cone covering with all but the lowest load fraction (0.05) is significantly better than the circle covering on all but three types of problems, described in section 3.1.4.

Assignment procedure (E, R). The results of the AID analyses for both parameters *E* and *R* reveal an almost perfect parallelism with the corresponding results for test set *G3*.

With respect to parameter *E*, which represents the combinations of the weights of the assignment cost 2.10 of stop *i* to seed *k*, it can be observed that the significantly worse combinations are those which only take account of the remoteness of stop and seed from the depot, without considering the proximity of stop and seed ($\eta_1 = 1, \eta_2 = 0$). For a number of problems, the proximity is at least as important as the remoteness of stop and seed from the depot, i.e. $\eta_1 \leq \eta_2$ (see section 3.1.8).

As far as the assignment procedure is concerned, the results of the AID analyses are insufficiently clear to support a choice on an aggregate level between an assignment based on the minimal regret function and a direct assignment. However, on the basis of the results one might deduce that the assignment based on the minimal regret function is preferred to the direct assignment if the geographical structure of the problem to be solved does not force the assignment

of a stop to a dedicated seed.

The conclusion for the PAI heuristic is that the behavior of its parameters does not seem to be really affected by mixed pick-ups and deliveries.

4.1.9 Generalized Assignment heuristic

Combining the four parameters of the GA heuristic gives rise to 120 different solutions for each problem of test set *P1*.

Parallel initialization procedure (S). The results of the AID analyses reveal that the behavior of parameter *S* is conform with that of the GA heuristic for test set *G3* and with that of all heuristics requiring seed generation for this test set. This means that the cone covering with all but the lowest load fraction is significantly better than the circle covering on all but the three types of problems, which were described in section 3.1.4.

Assignment procedure (E, R). As for test set *G3*, it can be stated that the proximity of stop and seed is fundamental in computing the assignment cost. For a number of problems, the proximity is at least as important as the remoteness of stop and seed from the depot.

As far as the choice between the two approaches for solving the GAP is concerned, no meaningful deductions can be drawn from the results of the AID analyses. The number of problems for which the the two approaches yield significantly different solutions, is very small.

Seed nature (P). Only the cone covering can be used for the generation of seeds which do not necessarily coincide with stops. Significant differences between the use of seed points and seed customers are only observed for only a small number of problems. The characteristics of these problems do not permit us to make meaningful deductions on a level higher than that of the individual problem.

These results reveal no clear effect of combined pick-ups and deliveries on the behavior of the parameters of the GA heuristic.

4.1.10 Two Phase heuristic

For each problem, 50 different solutions are obtained by combining the values of the three parameters of the TP heuristic.

Initialization criterion (I). The behavior of parameter *I* is similar to that observed for test set *G3* (see section 3.1.10). This implies that the initialization

with the stop farthest from the depot yields significantly better results, especially for problems with a homogeneous spread of stops.

Sequential selection criterion (G). Parameter G , representing the combinations of the weights γ_1 and γ_2 of the sequential selection criterion 2.12, shows a comparable behavior to that of all previous test sets. This implies that stops are added to the current route on the basis of a criterion in which the proximity of the unrouted stop j to the initialization stop i is at least as important as the proximity of stop j to the depot ($\gamma_1 \leq \gamma_2$).

Parallel selection criterion (L). The consistency in behavior with respect to all previous test sets is also observed for parameter L , representing the combinations of the weights λ_1 and λ_2 of the parallel selection criterion 2.13. In general, the insertion of stop j into the route of initialization stop i gives significantly better solutions if the proximity of stop j to initialization stop i is at least as important as the difference in remoteness of both stops from the depot.

These results indicate the almost negligible effect of mixed pick-ups and deliveries on the behavior of the parameters of the TP heuristic.

4.1.11 Conclusions of the parametric analysis for test set P1

Referring to the two groups of parameters distinguished for the previous three test sets (see section 3.1.11), the findings of the parametric analysis can be summarized as follows.

The first group of parameters is hardly affected by mixed pick-ups and deliveries. The significant effect of all these parameters as well as their significantly better values remain unaffected in comparison to all previous test sets.

The behavior of almost all parameters of the second group remain more or less invariant in comparison to those of test set $G3$. An exception is the parameter representing the number of places for adding stops, belonging to the SN and the SS heuristic.

4.2 Heuristic analysis for test set P1

All problems. The results of the comparison for the problems of test set $P1$ reveal that there is no real dominant heuristic for all problems. The GA

heuristics gives the better solutions, but these are only significantly better than the solutions of the GS, PAI, SW, SN and PN heuristics.

The fading of significant differences between heuristics is due to the route sequence imposed by the combined pick-ups and deliveries. The route sequence partly decreases the liberty of each heuristic to build its routes. Webb (1972) shares the same opinion by stating that the differences between the solution of heuristics tend to fade away if the side-constraints become harder because less space for variation remains.

Depot location. In correspondence with observations of previous test sets it has been observed that the more decentralized the position of the depot, the better the performance of the sequential heuristics relative to that of the non-sequential heuristics. For the *central* and *inside* pattern, approximately the same results as those for the test with all problems are obtained (cfr. supra). The TP heuristic gives the better solutions for the depot *outside*, but is not significantly better than the SI, SS and GA heuristics for this pattern.

Grouping patterns. As mentioned previously, the results for the distinct grouping patterns are more difficult to interpret in comparison with those of the spreading patterns.

As for the entire test set, no dominant heuristic can be observed for the grouping patterns. Several heuristics generate good solutions for most of these problems.

For the *clusters* pattern, no significant differences are observed at all.

Spreading patterns. The absence of dominant heuristics is confirmed by the results of the tests for the four spreading patterns. The good performances of the non-sequential heuristics for problems with a homogeneous spread of stops are only confirmed by the results of the *uniform* pattern. For this pattern, the significantly better solutions are obtained with the GA, PI, SW heuristic as well as with the PS heuristic.

As far as the *compressed* pattern is concerned, the sequential heuristics give no significantly worse solutions than the non-sequential heuristics.

For the *50% central* pattern, no significant differences are observed between the heuristics. The ability of sequential heuristic to minimize the links between central and peripheral stops becomes irrelevant due to the reduced number of routes.

For the *concentric* pattern, no clear difference can be made between sequential and non-sequential heuristics. The better solutions are obtained with the SS heuristics. Only the solutions of the SI and PI heuristics are not significantly worse.

Table 4.1 gives a summary of the relative performances of the eleven initial heuristics for the principal geographical pattern categories in the case of mixed

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	-+	--	--	-	--
PN	--	-+	--	--	--	--
SS	-	+	+	-+	+	+
PS	+	+	-+	-+	-+	-
GS	-	+	-+	-+	+	-
SI	-	-+	-+	+	-+	+
PI	++	-+	-	-+	-+	-+
PAI	-	-+	--	--	+	-+
GA	++	+	+	+	-+	-+
TP	+	-+	-	++	+	-+
SW	++	-+	--	-+	--	--

Table 4.1: Summary of the relative performances of the 11 heuristics for test set P1 (mixed pick-ups and deliveries). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

pick-ups and deliveries. This table reflects the absence of really dominant heuristics very well.

With respect to the computing times, we can observe that their magnitude is comparable to these of test set *G3*, but slightly shorter. The long time required is due to the large number of stops per route. As expected, the CPU-times of the GS, TP and SW heuristics remain much longer than that of the other heuristics.

In summary, mixed pick-ups and deliveries seems to affect more the relative behavior of the heuristics than the behavior of the parameters.

The most important conclusion is that the route sequence imposed by the mixed pick-ups and deliveries tends to fade the differences between heuristics. Exceptions are the SN and PN heuristics which remain significantly worse than the other heuristics, to some extent due to the large number of pathological problems for these heuristics.

Another important finding is the absence of one or more dominant heuristics. The latitude of heuristics in building the routes is to a large extent suppressed by the route sequence imposed.

4.3 Parametric analysis for the test set P2

In the following, the results of the AID analyses of the problems of test set *P2* are described for the ten initial heuristics separately. The SW heuristic is

excluded of these analyses due to the absence of parameters.

Again, the same pathological problems as detected in the previous test sets remain a threat for the internal validity of the experiments. It concerns mainly the problems which disturb the initialization criterion of the sequential route-building heuristics. Referring to the problems of the basic test set, the pathological ones are those with a *concentric* or a *50 % central* spreading pattern, both with a *central* depot (see appendix A).

These problems also affect the starting position of the sweep procedure of the cone covering method (cfr. appendix B.1.2). If the largest difference between the polar angles of two successive stops occurs between more than one pair of stops, the first evaluated pair is selected. Nevertheless, the bias caused is almost negligible.

The heuristic-specific pathological problems for this test sets are the same as those for the previous test sets.

4.3.1 Sequential Nearest Neighbor heuristic

Combining the values of the parameters of the SN heuristic gives a total of 18 solutions for per problem.

Initialization criterion (I). The results observed are similar to those observed for test set *G1*. Hence, the number of problems where significant differences between the three values of *I* can be observed is too limited to draw any conclusion.

If only the value of *I* is considered in the best solution, without taking account of any significance, the dominance of the initialization with the stop farthest from the depot is observed.

Frequency of initialization (P). The results for parameter *P* are more enunciated than they were for test set *G1*.

The multiple initialization is significantly better than the single initialization for more than half of the problems, especially for the problems with a homogeneous spread of stops.

If only the value of *P* in the best solution is considered, a clear dominance of the multiple initialization is observed.

The multiple initialization guarantees better separated routes. Each new route is started by means of the initialization criterion and hence, the interaction with the initialization criterion *I* is more meaningful in this situation. The single initialization, on the contrary, starts a new route approximately at the place where the previous route has been terminated.

Places to add stops (A). No significant difference can be observed between

the addition of stops at the end or at both ends of the current route. This corresponds with the findings for parameter A for test set $G1$.

The heterogeneous demand causes minor shifts in the significance structure of the parameters in comparison with the homogeneous demand. However, a clearer behavior of the parameters of the SN heuristic is observed throughout the test set.

4.3.2 Parallel Nearest Neighbor heuristic

The AID analyses for the PN heuristic are relatively unreliable due to the numerous pathological problems with a great number of nearest neighbor pairs in combination with the limited number of 8 replications per problem.

Simultaneous initialization procedure (S). The results of the AID analyses for this parameter are comparable to those of test set $G1$ (see section 3.1.2).

For the problems of the *clusters* pattern, no significant differences between the initialization procedures are observed.

In general, the circle covering with selection of circles with ascending radius as well as the initialization procedure through merging the nearest neighbors are significantly worse than the cone covering, respectively due to an overestimation of the seeds and to the pathological problems.

The circle covering with selection of circle with descending radius is not significantly worse for some problems with a depot among the stops and for stops compressed within a narrow area. In these cases, the number of seeds does not tend to be overestimated.

No visible effect is induced by the heterogeneous demand on the behavior of the parameter the PN heuristic.

4.3.3 Sequential Savings heuristic

The exhaustive combination of the four parameters of the SS heuristic gives rise to 270 solutions per problem for test set $P2$.

Initialization criterion (I). For most problems, the initialization with the closest stop to the depot yields significantly worse solutions than the initialization with the farthest stop from the depot and the initialization through merging the pair of stops giving the greatest savings.

The difference between the initialization with the farthest stop from the depot and the initialization through merging the greatest savings stops is seldom significant. This is due to the fact that the highest savings values are most often

obtained by merging two stops close to each other, but far from the depot.

This reinforces the tendency already observed for the problems of test set *G1*.

Frequency of initialization (P). The analyses reveal that the multiple initialization is significantly better than the single initialization for every five out of six problems. The multiple initialization combined with the two significantly better values for the initialization criterion *I* (cfr. supra) provide a guarantee for well-separated routes.

These results are also more emphasized than these of test set *G1*.

Places to add stops (A). The difference between the addition of stops at one end or at both ends of the current route is for most problems insignificant. This confirms the findings of test set *G1*.

Savings criterion (C). The results for the parameter representing the combinations of the weights of the savings criterion 2.3 are similar with those of all previous test sets.

The significantly better values of *C* preserve the inequality $\alpha_1 \leq \alpha_2$ and minimize or even neglect the weight α_3 . Hence, the savings value is primarily defined in terms of proximity of the stops to be merged and secondary in terms of remoteness of the stops from the depot.

Only on the individual problem level slight shifts can be observed among the significantly better values.

The effect of the heterogeneous demand causes a clearer behavior of the parameters of the SS heuristic compared with that of test set *G1*.

4.3.4 Parallel Savings heuristic

The PS heuristic with its three parameters give rise to 120 replications per problem of test set *P2*.

Simultaneous initialization procedure (S). The results for the parameter representing the parallel route initialization procedure show conformity with the results of test set *G1*.

The significantly better results are obtained with the cone covering. For problems with a homogeneous spread of stops, the cone covering with the lowest load fraction (0.05) gives significantly worse solutions than the other load fractions because the seeds are positioned too close to the depot. This hinders the formation of well-separated routes.

For the same problems the traditional initialization method without seeds gives no significantly worse solutions. Moreover, the best solutions for these problems are most often obtained with the initialization procedure without seeds.

The circle covering gives no significantly worse solutions for the problems of the same three specific types defined for test set *G1* (see section 3.1.4).

Savings criterion (C). The significance structure of the parameter representing the combinations of the weights of the savings criterion 2.3 is comparable to that of all previous test sets. The effect of *C* is significant for all problems. The significantly better values of *C* contain the combinations of the weights, which stress the importance of the proximity of the stops to merge more than the remoteness of both stops from the depot. In terms of weights, this can be expressed as $\alpha_1 \leq \alpha_2$ and minimizing or even neglecting the weight α_3 .

These results demonstrate that the heterogeneous demand has no clearly visible effect on the behavior of the parameters of the PS heuristic.

4.3.5 Generalized Savings heuristic

The AID analyses for the GS heuristic are unreliable due to the restricted number of 5 replications per problem for test set *P2*.

Generalized Savings criterion (P). The effect of parameter *P*, representing the combinations of the weights α_1 and α_2 of the generalized savings criterion 2.4, is insignificant for all problems. Besides the moderate number of mergers required, the limited number of replications can be cited as a possible reason for this phenomenon.

If only the value of *P* in the best solution is considered, it can be observed that the combination $\alpha_1 = \alpha_2$ is absolutely dominant.

These findings give evidence that the heterogeneous demand has no clear effect on the behavior of the parameter of the GS heuristic.

4.3.6 Sequential Insertion heuristic

Combining the values of the three parameters of the SI heuristic results in 60 replications per problem of test set *P2*.

Initialization criterion (I). The difference between the initialization of the current route with the stop closest to or farthest from the depot is significant for only a minority of problems. These problems mainly have a homogeneous spread of stops. For these problems the initialization with the stop farthest from the depot gives significantly better solutions.

If only the value of *I* in the best solution is considered, a dominance of the initialization with the stop farthest from the depot can be observed.

The same findings for this parameter has been made for test set *G1* (see section 3.1.6).

Insertion and selection criterion (M, L). Parameters *M* and *L* represent the combination of weights of the combinations 2.5 and 2.6 of insertion and selection criterion.

The results of the AID analyses of this test set correspond to a great extent with those in all previous test sets (see section 3.1.6).

In general, the significantly worse solutions are obtained with an insertion criterion which determines an insertion place as the place between the two most distant successive stops of the current route, without taking account of the minimization of the increase in travel time due to the insertion.

As far as selection criterion 2.5 is concerned, the significantly better combinations of its weights λ_1 and λ_2 favor the selection of the unrouted stop which causes the minimal increase in travel time after insertion rather than the stop at the largest travel time from the depot.

Selection criterion 2.6 is not significantly worse for problems where stops are grouped. For these problems, the selection of the stop nearest to the current route is the stop whose insertion results in the minimal route time.

Minor shifts among these significantly better values of the weight combinations of insertion and selection criteria can only be observed on the individual problem level.

These results do not prove a clear effect of the heterogeneous demand on the behavior of the parameters of the SI heuristic.

4.3.7 Parallel Insertion heuristic

Combining the three parameters of the PI heuristic results in 210 replications for each problem of test set *P2*.

Parallel initialization procedure (S). The AID analyses reveal that the behavior of the parallel route initialization procedure is similar to that of test set *G1*.

The cone covering yields significantly better solutions than the circle covering for all but three specific types of problems, described in section 3.1.7.

For the problems with a homogeneous spread of stops, the cone covering with the lowest load fraction (0.05) most often gives significantly worse solutions because stops are positioned too close to the depot.

Insertion and selection criterion (M, L) . Parameters *M* and *L* represent the combinations of weights of the combinations 2.5 and 2.6 of the insertion and selection criterion. The AID results are comparable to those of the PI heuristic

in the previous test sets (see section 3.1.7).

The significantly worse solutions are obtained with an insertion criterion which determines an insertion place as the place between the most distant successive stops of a route, without taking account of the increase in travel time due to the insertion.

Concerning selection criteria 2.5, the significantly better combinations of its weights λ_1 and λ_2 favor the selection of the unrouted stop which primarily causes the minimal increase in travel time and secondary lies at the maximal travel time from the depot.

Selection criterion 2.6 is not significantly worse for problems where stops are grouped. Used with the PI heuristic, this criterion provides a proportional growth of the parallel routes.

Minor shifts among these significantly better values of insertion and selection criteria can only be observed on the individual problem level.

Based on these findings we can conclude that the behavior of the parameters of the PI heuristic is hardly affected by the heterogeneous demand in comparison to the homogeneous demand.

4.3.8 Parallel Assignment-based Insertion heuristic

The three parameters of the PAI heuristic give rise to 210 replications for each problem of test sets $P2$. The higher number of replications for a problem of this test set in comparison to that for a problem of the previous test sets is due to parameter R which allows to take account of a heterogeneous demand pattern in the case of the direct assignment procedure.

Parallel initialization procedure (S). The behavior of the parallel initialization procedure corresponds entirely with that of PI heuristic of this test set and with that of the PAI heuristics of the previous test sets (see section 3.1.8).

The cone covering is significantly better than the circle covering on all but the three previously mentioned types of problems.

The cone covering with the lowest load fraction (0.05) is for most problems significantly worse than the other load fractions because the stops are positioned too close to the depot. This hinders the formation of well-separated of the routes.

Assignment procedure (E, R). As far as parameter E , representing the combinations of the weights ϵ_1 and ϵ_2 of the assignment cost 2.10 of stop i to seed k , is concerned, it is observed that the significantly better values for most problems are obtained with the combinations $\epsilon_1 \leq \epsilon_2$. This implies that the proximity of stop and seed is at least as important as the remoteness of stop

and seed from the depot. These findings correspond with those observed in all previous test sets.

With respect to the assignment procedure, it is observed that the direct assignment procedure with assignment cost 2.11 gives significantly worse results if only the magnitude of the customer demand is taken into account ($\rho_1 = 0$). In this case, stops are not assigned based on the increase in travel time caused by the insertion. The results obtained demonstrate that the heterogeneity of the demand is insufficient to be taken into account for realizing a better assignment of stops to seeds. The direct assignment only based on the customer demand is never significantly better than the direct assignment only with assignment cost 2.10 or than the assignment based on the minimal regret function.

Again, on the basis of the results one might deduce that the assignment based on the minimal regret function is preferred to the direct assignment if the geographical structure of the problem to be solved does not force the assignment of a stop to a dedicated seed.

The results show that the heterogeneous demand is not sufficiently heterogeneous to affect substantially the behavior of the parameters of the assignment procedure of the PAI heuristic.

4.3.9 Generalized Assignment heuristic

The GA heuristic is also equipped with a demand-sensitive assignment procedure, as the PAI heuristic. Consequently, 360 solutions are obtained for each problem of test set *P2* by combining the four parameters of the GA heuristic.

Parallel initialization procedure (S). The results of the AID analyses reveal that the cone covering is significantly better than the circle covering, except for some problems with well-separated groups of stops. The circle covering was specially designed for such problems.

Again, the cone covering with the lowest load fraction (0.05) is most often significantly worse than the other load fractions.

The explicit dominance of the cone covering could be explained by the fact that for the cone covering seeds must not necessarily coincide with the locations of the stops.

Assignment procedure (E, R). As far as the assignment cost 2.10 is concerned, the significantly better combinations of the weights ϵ_1 and ϵ_2 are conform with those of the previous test sets. This means that the proximity of stop and seed is fundamental in computing the assignment cost. For a number of problems (see section 3.1.9), the proximity is at least as important as the remoteness of stop and seed from the depot ($\eta_1 \leq \eta_2$).

As for the PAI heuristic, it can be observed that the direct assignment of

stops to seeds only based on their demand gives significantly worse results ($\rho_1 = 0$ in assignment cost 2.11). Nevertheless, taking the demand into account beside measures of proximity in the assignment cost for the direct assignment does not give significantly worse results than the assignment based on the minimal regret function or the direct assignment with assignment cost 2.10 for most problems.

Again, the results do not permit us to make meaningful deductions for the assignment procedure to be used. However, one cannot reject the thesis that the assignment based on the minimal regret function is preferred to the direct assignment if the geographical structure of the problem to be solved does not force the assignment of a stop to a dedicated seed.

Seed nature (P). Only the cone covering allows the generation of seeds which do not necessarily coincide with stops. Differences obtained with seed points or seed customers are seldom significant. The results do not permit us to make meaningful deductions on a level higher than that of the individual problem.

In comparison with the results of test set *G1* one might conclude that the effect of the heterogeneous demand on the behavior of the parameter of the GA heuristic is limited.

4.3.10 Two Phase heuristic

For each problem, 50 different solutions are obtained by combining the values of the three parameters.

Initialization criterion (I). For about half the problems, the difference between the initialization with the stop farthest from the depot and the stop closest to the depot gives significantly different solutions. The initialization with the stop farthest from the depot gives the significantly better solutions for these problems. However, these problems do not seem to have clear common characteristics.

By considering only the value of I in the best solution, the dominance of the initialization with the stop farthest from the depot can be observed.

These results tend to be more explicit than those of test set *G1*.

Sequential selection criterion (G). Parameter G , representing the combinations of the weights γ_1 and γ_2 of the sequential selection criterion 2.12, shows a comparable behavior to that of all previous test sets. This implies that stops are added to the current route based on the criterion in which the proximity of the unrouted stop j to the initialization stop i is at least as important as the proximity of stop j to the depot ($\gamma_1 \leq \gamma_2$).

Parallel selection criterion (L). The consistency of behavior with respect to all previous test sets is also observed for parameter L , representing the combinations of the weights λ_1 and λ_2 of the parallel selection criterion 2.13.

In general, the insertion of stop j into the route of initialization stop i gives significantly better solutions if the proximity of stop and initialization stop is at least as important as the difference in remoteness of both from the depot.

As far as the TP heuristic is concerned, these results give evidence on the consistency of behavior of its parameters and particularly those representing the weights of both selection criteria.

4.3.11 Conclusions of the parametric analysis for test set P2

The results of the parametric analysis for test set $P2$ can be summarized according to the two groups of parameters distinguished previously (see section 3.1.11).

The first group of parameters with consistent behavior is hardly affected by the heterogeneous demand. The significant effect of all these parameters as well as their significantly better values remain for the greater part invariant in comparison to all previous test sets.

The parameters of the second group are subjected to some shifts in comparison to their behavior for test set $G1$.

The significant effect of the sequential initialization criterion for the SN, SS, SI and TP heuristics becomes more explicit for this test set. The initialization with the stop farthest from the depot gives significantly better solutions for the majority of problems. For the SN and SS heuristics, the multiple initialization procedure dominates significantly the single initialization.

The experiments with a demand-sensitive assignment procedure for the PAI and GA heuristics showed that such an assignment procedure is probably more appropriate for problems with a higher degree of demand heterogeneity.

All together we may state that the behavior of the parameters of the second group is not severely affected by a heterogeneous demand in comparison to a homogeneous one.

4.4 Heuristic analysis for test set P2

All problems. The results of the comparison of the eleven initial heuristics for the problems of test set $P2$ reveal that there is no real dominant heuristic for all problems. The SS gives the better solutions, but the solutions of the PS, SI,

GA and TP heuristic are not significantly worse.

Both, the PN and SN heuristic give significantly worse results than almost all other heuristics, to some extent due to effect of the pathological problems.

By comparing these results with those of test set *G1* one can see that the dominance of the GA heuristic has faded away. A reason for this can be sought in the heterogeneous demand which hinders the formation of well separated groups of ten stops. Consequently, the heuristics requiring the generation of seeds or even the SW heuristic tend to perform less good for this test set.

In general, the latitude of building the routes is slightly decreased by the heterogeneous demand constraint.

Depot location. The tendency observed for the three depot patterns is even more explicit compared to that of test set *G1*.

The non-sequential heuristic give the better results for the *central* depot. The more the depot moves away from the central position, the better the performances of the sequential heuristics become. For the depot *outside* the SI and SS heuristics give significantly better solutions than all other heuristics.

Grouping patterns. As mentioned previously, the results for the distinct grouping patterns are more difficult to interpret in comparison to those of the spreading patterns.

As for the entire test set, no dominant heuristic can be observed for the grouping patterns. Several heuristics give good solutions for the various patterns.

Spreading patterns. The two patterns with a homogeneous spread of stops differ somewhat. For the *uniform* pattern, the PS as well as the GA heuristic give significantly better results than all other heuristics. These heuristics have the ability to keep the routes well-separated. These results confirm once more the good performance of the savings criterion for this pattern.

For the *compressed* pattern, the better solutions are obtained with the PS and TP heuristics. The solutions of the SS, SI and GA heuristics are not significantly worse. These results show that the differences for this pattern are not polarized around sequential and non-sequential heuristics.

The SS heuristic gives significantly better solutions than these of all other heuristics for the *concentric* pattern.

The sequential heuristics, in particular the SI and SS heuristic give significantly better solutions for the problems of the *50% central* pattern than all other heuristics, with exception of the PS heuristic. The good performances of the sequential heuristics for this pattern are explained by their ability to minimize the links between central and peripheral stops.

Table 4.2 gives a summary of the relative performances of the eleven initial heuristics for the principal geographical pattern categories of test set *P2*.

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	--	--	-	-+	-+
PN	--	--	--	--	--	--
SS	-	+	++	++	+	++
PS	++	+	-	++	-+	-+
GS	--	--	--	-+	--	--
SI	-	-+	-+	++	++	++
PI	+	+	-+	-+	-+	-
PAI	--	-+	-	--	+	-
GA	+	+	++	+	++	-+
TP	++	-+	-+	++	-	-+
SW	+	--	+	-+	--	--

Table 4.2: Summary of the relative performances of the 11 heuristics for test set P2 (heterogeneous demand). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

With respect to the computing times of the heuristics for test set *P2*, we can observe that they are comparable to these of test set *G1*. As expected, the computing times of the GS, TP and SW heuristics surpass largely that of the other heuristics. The CPU-time of the GA heuristic remains the shortest.

To conclude the heuristic analysis of test set *P2*, it can be stated that the heterogeneous demand partly suppresses the dominance of the GA heuristic, compared to test set *G1*. No single significantly better heuristic can be distinguished anymore.

Nevertheless, some tendencies observed for *G1* are preserved, like the good performances of the non-sequential heuristics for problems with a homogeneous spread of stops or the good performances of the sequential heuristics for a decentralized depot or a 50% central pattern.

Chapter 5

Analysis of the effect of time-related constraints

This chapter is dedicated to the analysis of the effect of time-windows on the behavior of the initial heuristics and their parameters. In order to realize this analysis, two test sets with time-windows were generated.

The first test set, called *T1* by convention, is used to test the effect of homogeneous time-windows. Test set *T1* is built by defining a time-window of 60 minutes for all customers of all problems of test set *G1* ($q = 10, Q = 100$). The opening and closing time of the time-window of each customer are equal. If e_i denotes the opening time and l_i the closing time of a stop $i \in N \setminus \{0\}$, then the time-window for stop i is defined by $[e_i, l_i] = [0, 60]$. No binding time-window is associated with the depot.

Adding the time-window constraint to the mathematical programming formulation of section 1.2 requires the definition of some additional variables. The end of service time (departure time) at stop i , d_i , is defined as the sum of the time to begin service, b_i and the actual service time s_i at stop i : $d_i = b_i + s_i$. For both test sets, the actual service time for a stop i is considered to be inexistent, $s_i = 0$ and thus is $d_i = b_i$. The time-window constraints can be formulated as:

$$\begin{aligned} x_{ijk} = 1 &\rightarrow d_i + t_{ijk} \leq d_j && i, j = 0, \dots, n \\ &0 \leq d_i \leq 60 && i = 1, \dots, n \end{aligned}$$

In the second test set, *T2*, two time-windows of 30 minutes each are associated with each customer of test set *G1* ($q = 10, Q = 100$). The two time-windows are separated by a closing time of 20 minutes. The opening and closing times of both time-windows are equal for all stops. For each stop $i \in N \setminus \{0\}$, two time-windows $[e_{1i}, l_{1i}] = [0, 30]$ and $[e_{2i}, l_{2i}] =$

[50, 80] are defined. The formulation is:

$$\begin{aligned} x_{ijk} = 1 &\rightarrow d_i + t_{ijk} \leq d_j & i, j = 0, \dots, n \\ 0 &\leq d_i \leq 80 & i = 1, \dots, n \end{aligned}$$

The departure time for the vehicles is adapted to the opening time of the first stop in the route. Consequently, waiting time can only occur for the problems of test set *T2*.

A general remark concerns the binding effect of time-windows. Homogeneous time-windows do not induce the same homogeneous effect as a homogeneous demand, for example. The binding character of a time-window depends heavily on the location of the customer. Consequently, time-windows do not have the same effect on all problems of the test sets *T1* and *T2*. Therefore, the values defining the time-windows of both sets were determined in such a way that for the majority of problems the constraints imposed by the time-windows would dominate those imposed by the vehicle capacity.

5.1 Parametric analysis for test sets T1 and T2

The results of the parametric analysis by means of the AID technique are discussed simultaneously for both test sets.

As in all previous test sets, the pathological problems can be considered as a threat for the internal validity of some results. The pathological problems are the same as in all previous test sets.

Another remark applies exclusively to the heuristics requiring the generation of seeds. It can happen that, due to very hard time-windows, the number of seeds is insufficient for allocating all stops to a route. The remaining unrouted stops are handled by the post-processor (see appendix B.3), which tries to insert each unrouted stop in an existing route or, if not possible, in a new route.

It is obvious that solutions requiring a large number of stops to be rerouted are partly biased. Consequently, it cannot be excluded that the results of the parametric analyses for the PN, PS, PI, PAI and GA heuristics are perturbed for the problems with hard time-windows.

5.1.1 Sequential Nearest Neighbor heuristic

For time-window problems, a time-oriented nearest neighbor criterion can be used with the SN heuristic. Therefore, a parameter is added to the three other

parameters of the SN heuristic. As a result the number of solutions rises to 144 per problem of test sets $T1$ and $T2$. The time-oriented nearest neighbor can only be used when stops are added at the end of the current route.

Initialization criterion (I). The effect of the initialization criterion is significant for almost all problems of both test sets. The significantly worse solutions are obtained by initializing the current route with the stop closest to the depot. Exceptions are some problems of the *concentric* and the *50 % central* pattern, inherent in the particular spread of their stops.

The significantly better solutions are obtained primarily by initializing the current route with the stop farthest from the depot and secondary through merging the two nearest neighbors. Both initialization procedures prevent that stops at high travel time from the depot have to be serviced by a separate route afterwards. This can happen when the current route is initialized with the stop closest to the depot. The vehicle of the current route could then arrive too late for servicing the stops at high travel times from the depot.

The implementations of a time-oriented SN heuristic by Solomon (1987) and Baker and Schaffer (1986) for heterogeneous time-windows used the initialization with the stop closest to the depot. They also used the criterion of nearest neighbor for initializing purposes.

Frequency of initialization (P). The effect of parameter P is significant for more than half of the problems. The multiple initialization procedure yields significantly better solutions for these problems. Moreover, for most of these problems the meaningful interaction with the significant effect of the initialization criterion can be observed.

For some problems, the single initialization is significantly better or no significant effect of P can be observed at all. This could be explained by the observation that the use of the multiple initialization tends to leave some stops between routes unrouted because a new current route is not necessarily started where the previous one was terminated. These stops have to be serviced in a separate route afterwards, which deteriorates the quality of the solution. For instance, the problems of the *50% clusters* pattern are vulnerable to this phenomenon.

If only the value of P in the best solution is considered, a dominance of the multiple initialization is observed.

Places to add stops (A). No significant differences are observed between the solutions obtained with the addition of stops at the end or at both ends of the current route. As a secondary effect, this can be interpreted as an indication of the fact that the simple and the time-oriented nearest neighbor criterion do not differ significantly because the time-oriented criterion can only be used when stops are added at the end of the current route.

Time-oriented nearest neighbor criterion (C). Parameter C represents the combinations of the weights α_1 , α_2 and α_3 of criterion 2.2. It is observed that the simple nearest neighbor is almost never significantly worse than its time-oriented alternative.

The significantly better combinations of the time-oriented nearest neighbor minimize the urgency of servicing, weighted by α_3 , and compute the proximity of two stops in terms of travel time, weighted by α_1 , and begin of service time, weighted by α_2 . The results indicate the trivial nature of the time-oriented nearest neighbor for problems without waiting times.

The results show that there is a change in the behavior of the parameters of the SN heuristic induced by time-window constraints. The waiting time gives rise to a less trivial use of the time-oriented nearest neighbor criterion.

5.1.2 Parallel Nearest Neighbor heuristic

For the PN heuristic, no time-oriented extension is provided. The AID analyses remain unreliable due to the numerous pathological problems containing a great number of nearest neighbor pairs and the limited number of 8 solutions per problem of test sets $T1$ and $T2$.

Simultaneous initialization procedure (S). The results reveal that the initialization through merging the two nearest neighbors gives significantly worse results for almost all problems.

The circle covering with selection of circles in descending order of their radius is not significantly worse than the cone covering for problems with binding time-windows. This can be explained by the fact that the overestimated number of seeds obtained by the circle covering is accidentally more appropriate than the insufficient number of seeds generated by the cone covering.

For the problems of the *clusters* and some of the *cones* pattern, with very loose time-windows, the circle covering is not significantly worse than the cone covering. The same observation has been made for test set $G1$ (see section 3.1.2). Nevertheless, none of both seed generating method is appropriate for problems with hard time-windows. The circle nor the cone covering, takes account of time-windows in order to determine the number of seeds required.

The behavior of the parameter of the PN heuristic is affected by the time-windows. However, no clear differences are observed between the behavior of the parameters among both time-window test sets. This is an indication of the fact that waiting time has no visible effect on the parameter of the PN heuristic.

5.1.3 Sequential Savings heuristic

The exhaustive combination of the four parameters of the SS heuristic gives rise to 270 solutions per problem of each of both test sets.

Initialization criterion (I). For most problems the initialization with the stop closest to the depot yields significantly worse solutions than the initialization with the stop farthest from the depot and the initialization through merging the greatest savings stops. The difference between the initialization of the stop farthest from the depot and the initialization through merging the greatest savings stops is seldom significant. This is due to the fact that the greatest savings are most often obtained by merging two stops close to each other but far from the depot.

For some problems of the *concentric* and the *50 % central* patterns either no significant differences between the values of I are observed or the initialization with stop closest to the depot is significantly better. This is explained by the typical spread of stops for these two patterns.

Frequency of initialization (P). The tendency observed for parameter P is similar but less explicit to that of the same parameter of the SN heuristic in this test set (see section 5.1.1). The multiple initialization is significantly better than the single initialization for half of the problems. Especially for problems with a homogeneous spread of stops, the multiple initialization gives significantly better solutions, most often in combination with a significant effect of the initialization criterion. The multiple initialization combined with the initialization with the stop farthest from the depot or with the initialization through merging the two greatest savings stops, favor the formation of separated routes.

For some problems, either no significant effect is observed or the single initialization is significantly better. This could be explained by the observation that the use of the multiple initialization tends to leave some stops between routes unrouted because a new current route is not necessarily started where the previous one was terminated. These stops have to be serviced with a separate route afterwards, which deteriorates the quality of the solution. Particularly the problems of the *50% clusters* pattern can be subjected to these situations. The same observation has been made for the above-described SN heuristic.

Places to add stops (A). The difference between the addition of stops at one end or at both ends of the current route is significant for a number of problems. Nevertheless, the results do not allow to clearly relate the significantly better values to the problem characteristics on an aggregate level.

However, for a number of problems a tendency is observed where the addition of stops at both ends of the current route is significantly better if the initialization with the stop closest to the depot is significantly worse. This can be explained by the fact that starting a new route with the stops closest to the

depot makes it difficult to insert some more stops before the initialization stop.

Savings criterion (C). The results for the parameter representing the combinations of the weights of the savings criterion 2.3 is similar to these of all previous test sets.

The significantly better values of C preserve the inequality $\alpha_1 \leq \alpha_2$ and minimize or even neglect the weight α_3 . Consequently, the better savings values are primarily defined in terms of proximity of the stops to be merged and secondary in terms of remoteness from the depot.

Only on the individual problem level, slight shifts can be observed among the significantly better values.

The time-windows have some effect on the significance structure of all the parameters of the SN heuristic, except on the savings criterion. The waiting time does not seem to have a particular effect on the behavior of the parameters.

5.1.4 Parallel Savings heuristic

The PS heuristic with its three parameters gives rise to 120 replications per problem of both test sets.

Simultaneous initialization procedure (S). For problems with hard time-windows, the results of the AID analyses show that the initialization procedure without seeds, through merging the greatest savings stops, gives no significantly worse solutions. The main advantage of this initialization procedure is the fact that it takes account of all side-constraints.

The circle covering is accidentally not significantly worse than the cone covering for the above-mentioned problems with hard time-windows. This has also been observed for the PN heuristic in this test set.

For the problems of the *clusters* and some of the *cones* pattern, with very loose time-windows, the circle covering is not significantly worse than the cone covering. The same observation has been made for test set *G1* (see section 3.1.4). Although neither seed generating procedures is appropriate for handling time-windows, the cone covering is almost never significantly worse than the procedure through merging the greatest savings stops. The cone covering with the lowest load fraction (0.05) gives significantly worse solutions than these with the other load fractions. The positioning of seeds too close to the depot can make it difficult to service stops far from the depot with other than dedicated routes.

Savings criterion (C). The significance structure of the parameter representing the combinations of the weights of the savings criterion 2.3 corresponds to that of all previous test sets.

The effect of C is significant for all problems. The significantly better values of C contain the combinations of the weights which stress more the importance of the proximity of the stops to merge than the remoteness of both stops from the depot. In terms of weights this can be expressed as $\alpha_1 \leq \alpha_2$ and minimizing or even neglecting weight α_3 .

These results demonstrate that there is only an effect of the time-windows on the behavior of the initialization procedure of the PS heuristic. No effect of the waiting time is observed.

5.1.5 Generalized Savings heuristic

The AID analyses for the GS heuristic remain unreliable due to the restricted number of 5 replications per problem of test sets *T1* and *T2*.

Generalized Savings criterion (P). The effect of parameter P , representing the combinations of weights α_1 and α_2 of the generalized savings criterion 2.4, is insignificant for the solution of all problems. The limited number of replications can be considered as one of the possible reason for this phenomenon.

If only the value of P in the best solution is considered, it can be observed that the combination $\alpha_1 = \alpha_2$ is most often dominant. The combination $\alpha_1 = 1, \alpha_2 = 0$ never generates the best solution because this merges the routes only on the magnitude of their travel time.

The effect of the time-windows on the parameter of the GS heuristic is difficult to quantify due to the limited reliability of the AID analysis.

5.1.6 Sequential Insertion heuristic

The SI heuristic comes with three different combinations of insertion and selection criterion. These combinations can be represented by three parameters. In addition to the parameter representing the initialization criterion, a total of four parameters results. Combining the values of the four parameters results in 450 replications per problem of test set *T1* and *T2*.

Initialization criterion (I). The difference between the initialization of the current route with the stop closest to or farthest from the depot is significant for a large number of problems.

A clear tendency can be observed indicating that the significantly better solutions for the *50% central* and the *concentric* patterns with a decentralized depot are obtained with the initialization with the stop closest to the depot.

This can be explained by the position of the depot among the stops for these patterns.

The unclear significance structure for the initialization criterion has also been observed for all previous test sets.

Insertion and selection criterion (M, L, E). Parameters M , L and E represent the combination of weights of the combinations 2.7, 2.8 and 2.9 of the insertion and selection criterion.

The results of the AID analyses reveal that, for most problems, criterion combination 2.9 gives significantly better solutions. For this criterion, selection and insertion criterion are equal. The criterion determines insertion places and selects stops for insertion in the current route based on the minimal increase in travel time, the minimal increase in begin of service time at stop $i + 1$, and the urgency of servicing stop u respectively weighted by η_1 , η_2 and η_3 . The significantly better combinations of these three weights are those which minimize or even neglect weight η_3 . This means that the urgency of servicing stop u , whose insertion is considered, is of less importance than the increase in travel time and the increase of begin of service time at stop $i + 1$. Stop $i + 1$ is the successor of the inserted stop u in the current route.

Due to the aggregation of effects in criterion 2.9 it is difficult to extract the significantly better combinations of weights λ_1 and λ_2 for the term representing the increase in travel time. Nevertheless, it can be observed that the combination $\lambda_1 = 0, \lambda_2 = 1$ is significantly worse because this combination determines the insertion place for all stops as the largest link between two successive nodes of the current route.

For the small number of problems where combination 2.9 is not significantly better, the combination 2.8 gives significantly better solutions. For this combination the stop, whose insertion yields the minimal route time of the current route, is selected. The insertion criterion is the same as for combination 2.9, with exclusion of the term representing the urgency of servicing stop u .

A tendency can be observed, which indicates that the less binding the time-windows are, the more important the term representing the increase in travel time becomes with respect to the term representing the change in the begin of service time at stop $i + 1$.

Combination 2.7 is significantly worse for most problems, except for the problems for which combination 2.8 is also retained.

Although the SI heuristic uses specific time-oriented combinations of selection and insertion criterion, the behavior of its parameters remains to a high degree comparable to that of the previous test sets. The effect of waiting time does not induce a visible shift in the behavior of the parameters of the SI heuristic. Nevertheless the combinations of insertion and selection criterion allow a waiting time-sensible route-building.

5.1.7 Parallel Insertion heuristic

The PI heuristic uses the same time-oriented combinations of insertion and selection criterion as the SI heuristic (see section 5.1.6). Three parameters are required for their representation. In addition to the parallel initialization parameter, this gives a total of four parameters. Combining these four parameters of the PI heuristic results in 1575 replications for each problem of test sets $T1$ and $T2$.

Parallel initialization procedure (S). The AID analyses reveal that the cone covering method is significantly better than the circle covering. Exceptions to this statement are made for the problems with hard time-windows. In these cases the overestimated number of seeds of the circle covering is not significantly worse than the seeds generated with the cone covering.

For the problems of the *clusters* and some of the *cones* pattern with very loose time-windows, the circle covering is not significantly worse than the cone covering. The same observation has been made for test set $G1$ (see section 3.1.7).

Another tendency indicates that the cone covering with the highest (1.00) or lowest (0.05) load fractions tends to give significantly worse results than that with the other load fractions, especially in the case of problems with binding time-windows. This can be explained by the fact that the seeds generated in these cases are too close or too far from the depot, which can give rise to a number of unrouted stops too far from the seeds.

Insertion and selection criterion (M, L, E). Parameters M , L and E represent the same weights combinations of the combinations 2.7, 2.8 and 2.9 of the insertion and selection criterion as for the SI heuristic (see section 5.1.6).

The results of the AID analyses for test sets $T1$ and $T2$ are a little more fuzzy than for the SI heuristic. This is probably due to the remaining unrouted stops which have to be assigned to a route by the post-processor afterwards.

For five out of six problems, combination 2.9 gives significantly better solutions. As for the SI heuristic, the significantly better combinations of the weights η_1 , η_2 and η_3 are those which minimize or even neglect weight η_3 of the term representing the urgency of servicing the stop u considered for insertion. For the term weighted by η_1 , representing the increase in travel time, the combination $\lambda_1 = 0, \lambda_2 = 1$ is significantly worse. In latter case the insertion place for all stops is given by the largest link between two successive nodes in one of the existing routes.

For the small number of problems for which combination 2.9 is not better, combination 2.8 gives significantly better solutions. For this combination the stop whose insertion yields the minimal route time among all routes, is selected. In the case of the PI heuristic, this selection criterion performs a proportional building of all routes.

Again, the less binding the time-windows, the more important the term repre-

sending the increase in travel time becomes with respect to the term representing the change in the begin of service time.

Rarely, combination 2.7 is not significantly worse. However, this can be considered as an indication of the inferior importance of the remoteness of a stop from the depot when selecting it for insertion.

These findings show that particularly the behavior of the initialization procedure of the PI heuristic is affected by the time-windows. The results for the dedicated time-oriented combinations of insertion and selection criterion illustrate that the overall tendency is comparable for the greater part to that of the previous test sets.

Although the waiting time is taken into account explicitly by the combinations of insertion and selection criterion, its influence cannot be clearly observed from the behavior of the parameters.

5.1.8 Parallel Assignment-based Insertion heuristic

The three parameters of the PAI heuristic give rise to 70 replications for each problem of test sets *T1* and *T2*.

Parallel initialization procedure (S). The results of the AID analyses of both test sets are comparable to those of the previous heuristics requiring the generation of seeds in this chapter.

The cone covering gives significantly better solutions than the circle covering, except for problems with hard time-windows. For these problems, the overestimated number of seeds of the circle covering is accidentally appropriate.

For the problems of the *clusters* and some of the *cones* pattern, with very loose time-windows, the circle covering is not significantly worse than the cone covering. The same observation has been made for test set *G1* (see section 3.1.8).

The cone covering with the lowest load fraction gives significantly worse solutions, in particular for problems with hard time-windows. As mentioned previously, this can be explained by the fact that seeds too close to the depot can prevent stops far from the depot to be inserted in a route because the vehicle cannot arrive there in time.

Nevertheless, neither seed generation method proposed is really appropriate for problems with time-windows.

Assignment procedure (E, R). Concerning parameter *E*, representing the combinations of weights ϵ_1 and ϵ_2 of assignment cost 2.10 of stop *i* to seed *k*, it is observed that the significantly better values for most problems are obtained with the combinations $\epsilon_1 \leq \epsilon_2$. This implies that the proximity of stop and seed is at least as important as the remoteness of both from the depot.

These findings correspond with those of the previous test sets.

With respect to the assignment procedure, it is observed that there is a significant difference between the assignment based on the minimal regret function and the direct assignment for about half of the problems of each test set. The significantly better solutions for these problems are almost always obtained with the assignment based on the minimal regret function. However, the common characteristics of these problems are not sufficiently clear to draw conclusion on an aggregate level.

If only the value of parameter R in the best solution is considered, the dominance of the assignment based on the minimal regret function can be observed.

The results show that the effect of time-windows is the most explicit on the parameter representing the initialization procedure for the PAI heuristic. Waiting time does not visibly affect the behavior of the parameters.

5.1.9 Generalized Assignment heuristic

Combining the four parameters of the GA heuristic results in 120 solutions for each problem of both test sets.

Parallel initialization procedure (S). The behavior of the parallel initialization procedure corresponds with that of the other heuristics of this chapter requiring seeds.

This means that the cone covering is significantly better than the circle covering, except for problems with hard time-windows.

For the problems of the *clusters* and some of the *cones* pattern, with very loose time-windows, the circle covering is not significantly worse than the cone covering. The same observation has been made for test set $G1$ (see section 3.1.9).

The cone covering with a too low load fraction positions its seeds too close to the depot. This can result in a great number of stops far from the depot to remain unrouted during the routing phase of the GA heuristic.

Assignment procedure (E, R). As far as assignment cost 2.10 is concerned, the significantly better combinations of its weights ϵ_1 and ϵ_2 correspond for the greater part with those observed in the previous test sets. The significantly worse solutions are obtained with the combination $\epsilon_1 = 1, \epsilon_2 = 0$. An assignment cost with these weights does not take any account of the proximity of stop and seed.

Further refinements cannot be made based on the results. The great number of stops to be rerouted by the post-processor can partly be held responsible for this.

As far as the assignment procedure is concerned, a significant difference between the two assignment procedures is observed for a negligible number of problems. The number of stops to be rerouted by the post-processor due to the

time-windows can be a possible reason for the almost insignificant effect of the assignment procedure.

Seed nature (P). Only for the GA heuristic, the cone covering procedure allows the generation of seed points which do not necessarily coincide with stops. Differences between solutions obtained with seed points and seed customers are significant for only a very small number of problems. Based on these problems, conclusions on a level higher than that of the individual problem cannot be made.

The GA heuristic is due to its two-phase nature not really appropriate for handling hard time-windows. For problems with hard time-windows, a large number of stops to be rerouted is most often the result.

5.1.10 Two Phase heuristic

For each problem, 50 different solutions are obtained by combining the values of the three parameters.

Initialization criterion (I). For about half the problems, the difference between the initialization with the stop farthest from the depot and the stop closest to the depot gives significantly different solutions.

A tendency observed indicates that the initialization with the stop farthest from the depot does not necessarily generate significantly better solutions for problems with a homogeneous spread of stops, as it is the case for test set *G1*. This is probably due to the binding effect of the time-windows.

If only the value of parameter *I* is considered in the best solution, neither initialization procedure has a higher frequency of occurrence.

Sequential selection criterion (G). Parameter *G*, representing the combinations of the weights γ_1 and γ_2 of sequential selection criterion 2.12, shows a comparable behavior to that of all previous test sets. This implies that stops are added to the current route on the basis of a criterion in which the proximity of unrouted stop *j* to initialization stop *i* is at least as important as the proximity of stop *j* to the depot ($\gamma_1 \leq \gamma_2$).

Parallel selection criterion (L). The consistency of behavior with respect to all previous test sets is also observed for parameter *L*, representing the combinations of weights λ_1 and λ_2 of parallel selection criterion 2.13.

In general, the insertion of stop *j* into the route of initialization stop *i* gives significantly better solutions if the proximity of stop *j* to initialization stop *i* is at least as important as the difference of remoteness of both stops from the depot.

The time-windows only seem to affect the initialization procedure of the TP heuristic. The waiting time has no visible effect on the behavior of the parameters.

5.1.11 Conclusions of the parametric analysis

The parametric analyses have shown that time-windows have an effect on the behavior of parameters. The magnitude of the effect depend upon the type of parameter. If the two groups of parameters distinguished previously (see section 3.1.11) are considered again, the following conclusions can be drawn.

The first group of parameters is hardly affected by the time-windows. The significant effect of all these parameters as well as their significantly better values remain for the greater part unaffected in comparison to all previous test sets. Although the SN, SI and PI heuristic use specific time-oriented criteria, the significantly better combinations of their weights remain basically the same in comparison to these of the previous test sets.

The parameters of the second group are subjected to some considerable shifts in comparison to their behavior for test set $G1$.

The main observation is that the circle and the cone covering procedures for seed point generation are inappropriate for problems with hard time-windows. Both take only account of capacity and demand-related constraints for generating seeds. A good alternative is offered by the procedure used by the TP heuristic. In the case of the PN or PS heuristics, the initialization procedure through merging can also be used with binding time-windows.

The behavior of the sequential initialization criterion has partly changed in comparison to that of test set $G1$. For the SN and SS heuristic the multiple initialization procedure with the stop farthest from the depot or through merging a pair of stops satisfying the current criterion is significantly better.

For the SI and TP heuristic, no deductions with respect to the initialization criterion can be made.

As far as the remaining parameters of the second group are concerned, no additional meaningful inferences can be made.

Related to this, a last remark concerns the possible effect on the solution value caused by the sometimes high number of stops to be rerouted by the post-processor.

5.2 Heuristic analysis

5.2.1 Heuristic analysis for test set T1

All problems. The better solution for the problems of test set *T1* are obtained with the PS heuristic. However, its solutions are not significantly better than those of the SI and SS heuristic. The good performances of the PS heuristic is mainly due to the use of the traditional procedure without seed points for problems with hard time-windows. The good performances of the SI and SS heuristic are expected, in so far that sequential heuristics never have unrouted stops left.

The significantly worse solution are obtained with the SN, PN, GS and SW heuristic. The performances of the SW heuristics are bad because only the proximity of stops in terms of their polar-angle and not in terms of their time-windows is considered. In order to overcome partly this problem, Solomon (1987) proposed to embed a time-oriented insertion heuristic in a SW heuristic.

Depot location. The time-windows tend to reinforce the dominance of the sequential heuristics as the depot is more decentralized. The performance of PS, SI and SS heuristics is significantly better than all other heuristics for a depot *outside*. For a depot *inside*, only the TP heuristic is not significantly worse than the PS, SI and SS heuristics.

The significant differences have faded for a *central* depot. Only the PN, SN and GS heuristics are significantly worse than the PS, SI and SS heuristics.

Grouping patterns. The PS, SS and SI heuristic give good results for the grouping patterns, except for the *clusters* and the *cones* pattern. Both patterns are characterized by loose time-windows. For these patterns, the behavior of the heuristics is comparable to that of test set *G1*.

Spreading patterns. If the spreading patterns, on which time-windows have an impact are considered, the following results are obtained. For the problems with a homogeneous spread of stops (*uniform* and *compressed* patterns), the dominance of the three heuristics, SS, SI and PS is confirmed. The good performances of the savings criterion for these patterns have to be noticed once more.

For the *concentric* and the *50% central* patterns, the good performances of the PS and sometimes the SI heuristic are observed. Nevertheless, the significant differences between the heuristics are limited for both patterns. However, the PN, SN, GS, SW, PI and PAI heuristics are significantly worse for most problems of both patterns.

Table 5.1 presents an overview of the relative performances of the 11 initial heuristics for the principal geographical pattern categories of test set *T1*.

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	--	--	-	-	--
PN	--	-+	--	--	--	--
SS	++	--	+	++	++	+
PS	+	+	++	+	+	+
GS	--	-+	-	--	--	--
SI	+	-+	+	++	++	++
PI	-+	-	--	-+	-+	-
PAI	-+	+	-	-+	-+	-
GA	-+	-+	+	-	-	-
TP	-+	+	+	-+	-+	+
SW	--	--	--	--	--	--

Table 5.1: Summary of the relative performances of the 11 heuristics for test set T1 (1 homogeneous time-window). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

Comparison of the CPU-times of the eleven heuristics reveal that their proportionality is preserved compared to these of test set *G1*. The shortest computing times are obtained with the sequential heuristics and the GA heuristic. Again, the SW, TP and GS heuristic have much longer computing times due to the large number of TSPs to be solved.

5.2.2 Heuristic analysis for test set T2

All problems. For the problems of test set *T2*, the SI heuristic performs significantly better than all other heuristics, except the TP heuristic. The dominance of the SI heuristic can be explained by its capability of handling waiting time.

The good performances of the SI heuristic have also been observed by Solomon (1987) and Potvin and Rousseau (1993). Latter authors also confirm the adequacy of the seed generation procedure used by the TP heuristic for time-window problems.

The fact that PS and SS heuristics are significantly worse than the SI heuristic for this test set can probably be explained by the fact that they do not handle waiting times explicitly. Therefore, Solomon (1987) and Van Landeghem (1988) propose to introduce a maximal waiting time for savings heuristics.

The SW heuristic is significantly worse than all other heuristics for the same reason as for test set *T1*.

Depot location. The performance of the SI heuristic improves with a more decentralized depot.

For the depot *outside*, the SI heuristic is significantly better than all other heuristics.

For both other depot patterns, the SI heuristic gives the best solutions, but the solutions of some other heuristics among which the PS, SS, GA, TP heuristics are not significantly worse.

Grouping patterns. The results of the grouping patterns are comparable to these of test set *T1* and *G1*, but significant differences are even more scarce. This could be explained by the fact that vehicles can drive instead of waiting at stops which decreases somewhat the binding effect of time-windows.

Spreading patterns. The TP heuristic performs very well for the *uniform* pattern. However, only the PN, SN, PI, GS and SW heuristic are significantly worse.

As noticed for previous test sets, the SI heuristic performs very well for the *compressed* pattern. The solutions of the TP and PS heuristics are not significantly worse for this pattern.

For problems with a *50% central* pattern, the SI heuristic is also significantly better than all other heuristics, except the SS heuristic. The dominance of the sequential heuristics for this pattern is explained by their ability to limit the number of links between central and peripheral stops.

For the *concentric* pattern, no dominant heuristics are distinguished. Only the PN, SN, PAI and SW heuristics are significantly worse than most other heuristics.

Table 5.2 presents an overview of the relative performances of the 11 initial heuristics for the principal geographical pattern categories of test set *T2*.

As far as the computing times are concerned, the same observations as for test set *T1* can be made.

Summarizing the results of the heuristic analyses of test sets *T1* and *T2*, the following findings have to be noticed.

In general, non-sequential heuristics requiring the generation of seeds give bad solutions for problems with binding time-windows. The SI, PS, SS and TP heuristics most often perform very well for this kind of problems. The SI heuristic is provided with time-oriented insertion and selection criterion for a good waiting-time management. The PS and SS heuristic perform well as far as no waiting time is involved and particularly if the spread of stops is homogeneous. The TP heuristic is characterized by a side-constraint-sensitive

	depot central/inside			depot outside		
	homo- geneous spread	clustered stops	central & peripheral stops	homo- geneous spread	clustered stops	central & peripheral stops
SN	--	--	--	-	--	-
PN	--	-	--	--	-	--
SS	+	--	-+	++	--	-+
PS	+	+	-+	+	+	+
GS	--	-+	-	-	+	--
SI	+	-+	+	++	++	++
PI	-+	-	-	-	-	-
PAI	-+	+	--	--	+	-
GA	-+	-+	++	-	-	-
TP	+	+	+	++	-+	-+
SW	--	--	--	--	--	--

Table 5.2: Summary of the relative performances of the 11 heuristics for test set T2 (2 homogeneous time-windows). Symbols: "++": very good; "+": good; "-+": moderate, "-": bad, "--": very bad.

procedure for generating seeds through a sequential heuristic.

The worst results are obtained with the SW, GS, SN and PN heuristic. Once again, the pathological problems are probably partly responsible for the bad performances of the SN and PN heuristics.

Chapter 6

Implementations of the improvement heuristics

Improvement heuristics for the VRP are aimed at enhancing an initial feasible solution through a search mechanism. Improvement heuristics can be classified following two criteria: the type and the search strategy.

The type of improvement heuristic refers to the number of routes involved in the improvement: a *within-route* procedure if only one route at a time is involved and a *between-routes* procedure if two routes are involved. The improvement heuristics proposed are designed for the VRP. These between-routes heuristics improve the initial solution by moving stops between a pair of routes.

The within-route improvement procedures are aimed at improving the sequence of a route by changing the order of stops within the route. They are typical for the TSP. As mentioned previously, all heuristics, both of the initial and of the improvement type, are provided with a 3-opt within-route heuristic, which is invoked at any time a change to a route is performed (see appendix B.2.2).

The search strategy is a procedure which indicates the order in which new solutions are searched. Among the VRP improvement heuristics, two categories can be distinguished based on the search strategy: local and global optimisation methods.

The local optimisation heuristic, the traditional "descent" method, finds a local minimum by performing only moves of stops which result in an improvement of the objective function value.

The search strategy of these local optimisation heuristics is blind. This means that the order by which new solutions are generated is only dependent on the information gathered during the execution of the heuristic (Osman (1991)).

These heuristics are halted if no further improvement of the objective function

value is possible. As a matter of fact, local optimisation heuristics are trapped in the local optimum in which they descend.

The implementation of the local optimisation heuristic proposed, is called the Local Improvement (LI) heuristic.

Global optimisation heuristics succeed in leaving the local optimum by temporarily accepting moves which cause a worsening of the objective function value. These heuristics are often called "metaheuristics" because the procedure used to generate a new solution out of the current one, is embedded in a heuristic which determines the search strategy.

The search strategy for the global optimisation heuristics is a directed search. This implies that information of the problem domain and the nature of the objective is used to direct the search procedure towards promising regions (Osman (1991)).

The main drawback of metaheuristics is that they have no definite stopping criterion defined. The longer the computing time, the higher the probability of finding the global minimum.

A Simulated Annealing (SA) and a Tabu Search (TS) metaheuristic are proposed. Other metaheuristics are Genetic Algorithms, Neural Networks, Great Deluge, Noise method, ... Most of these metaheuristics are based on principles of physical or biological processes.

In the following, at first, the different types of moves are proposed, which are required for the improvement heuristics to generate a neighborhood solution out of an existing solution. Subsequently, the implementations of the three improvement heuristics are described.

6.1 Types of move

The three improvement heuristics (LI, SA and TS) all use a same type of move. A move can be defined as the mutation of stops between routes in order to obtain a neighborhood solution out of an existing solution.

It is assumed that only feasible moves with respect to the side-constraints are performed.

Four types of moves are considered. They are called the String Cross, the String Exchange, the String Relocate and the String Mix by convention.

String Cross. The String Cross (SC) is a move in which two strings of stops are exchanged by crossing two arcs of two different routes. This corresponds to the exchange of entire route segments. An example of a SC move is given by figure 7.3.1.

The concept of this type of move has also been presented by Savelsbergh (1988)




Figure 6.1: Example of a String Cross.

and Potvin et al. (1992).

String Exchange. The move of type String Exchange (SE) generates a neighborhood solution by exchanging two strings of stops between two routes. Symbolically, this can be represented by (x_1, x_2) , where x_1 and x_2 are integers representing the length of the strings to be exchanged in both routes. If K denotes the maximum length of a string of stops, then $1 \leq x_1 \leq K$ and $1 \leq x_2 \leq K$ must be satisfied. The length of the strings x_1 and x_2 is not necessarily equal. Figure 6.2 contains an example of an SE move.

A similar type of move has been defined by Dror and Levy (1986) and Savelsbergh (1988). The two most common values for the maximum string length are $K = 1$ and $K = 2$.

String Relocation. The String Relocation (SR) can be described as the move of a string of stops from one route to another. Symbolically, this can be represented by $(x, 0)$ or $(0, x)$. The maximum number of stops to be relocated is bounded by a parameter called the maximal string length K , i.e. $1 \leq x \leq K$. This type of move is able to reduce the number of routes. See figure 6.3 for an example of a SR.

Dror and Levy (1986) and Savelsbergh (1988) proposed a similar type of move. The values $K = 1$ and $K = 2$ are commonly used.

String Mix. The String Mix (SM) is a mixture of the String Exchange and the String Relocate. When implemented, this move type selects the best between the String Exchange and the String Relocation. The number of routes can be reduced through the moves of the type String Relocation.




Figure 6.2: Example of a String Exchange.




Figure 6.3: Example of a String Relocation.

Another variant of a move generation mechanism is the λ -interchange mechanism described in Osman (1993). This mechanism considers a relocation as a special case of an exchange.

6.2 Local Improvement heuristic

The LI heuristic is very popular due to its simplicity and its relatively short computing time. Most published implementations of this heuristic for the VRP are similar. This implementation is aimed at evaluating the effect of a number of parameters on the final solution.

Procedure

- Step 1:* The initial solution is the current solution.
- Step 2:* Select the first move by selecting the first pair of routes of the current solution, based on their order number. Select the stops required for the type of move in both routes based on the route sequence.
- Step 3:* If the move is unfeasible or does not improve the objective function value, then go to step 5.
- Step 4:* If the selection strategy (cfr. *infra*) is first improvement, then perform the move. The new solution becomes the current solution. Go to step 2.
If the selection strategy is the best improvement, and the neighborhood solution is better than all previous solutions, then save the neighborhood solution and go to step 6.
- Step 5:* If an entire evaluation cycle has been completed, and the selection strategy is direct improvement, then the heuristic is stopped.
If an entire evaluation cycle has been completed, and the strategy is the best improvement, and the best neighborhood solution has been saved, then this neighborhood solution becomes the current solution. Go to step 2. If no improved neighborhood solution has been saved during a complete evaluation cycle for the best improvement strategy, then the heuristic is stopped.
- Step 6:* Select the next move by selecting the next stops in the route or, if needed, the next pair of routes to be evaluated. If no move can be selected go to step 5, else go to step 3.

Parameters

Initial solution (S)

1. A bad initial solution.
2. A good initial solution.

This parameter is provided especially for evaluating the effect of the initial solution on the final solution. The initial solution is generated with one of the initial heuristics (see section 7.1).

String Length (K)

1. String length of 1 stop.
2. String length of 2 stops.

With this parameter, the effect of the number of stops in the string to be moved can be considered. This parameter is only used if the type of move is SE, SR or SM.

Selection strategy (P)

1. First improvement: the first improving move is performed.
2. Best improvement: the best improving move after completion of an evaluation cycle is performed.

Evaluation procedure for string length $K > 1$ (F)

1. Evaluation of all possible string lengths between a pair of routes before selecting the next pair.
2. The string length is increased by one after having completed an evaluation cycle without improvement.

This parameter is irrelevant if the move type is SC.

The initial solution S and the string length K can be considered as problem-specific parameters. The other three parameters are of a more generic nature.

6.3 Simulated Annealing metaheuristic

The principle of the SA metaheuristic is deduced from the physical annealing process of solids. Kirckpatrick et al. (1983) and Cerny (1985) proposed the use of SA for combinatorial problems. Their work is based on the research of Metropolis et al. (1953) in the field of Statistical Mechanics. For an overview of the research and applications of SA, the reader is referred to Van Laarhoven and Aarts (1987), Aarts and Korst (1989), Collins et al. (1988) and Eglese (1990).

The analogy between physical annealing and SA is obvious. The purpose of physical annealing is to obtain a solid in its ground state with its atoms arranged into regular patterns. The aim of SA is to reach the optimal solution of the problem, in casu the VRP. In order to reach this state, the solid is heated followed by a gradual cooling. The cooling has to be slow enough to reach a thermal equilibrium at each temperature. The ground state corresponds with the minimal energy configuration of the atoms. A feasible solution to the VRP is considered as an atom configuration. The energy corresponding to an atom configuration is the objective function value in the case of SA.

In order to reach a thermal equilibrium at a given temperature, it is necessary for the temperature to remain stable long enough. As a result, the configurations will be distributed according to the Boltzmann distribution. Translated into SA terms, this implies that the number of feasible neighborhood solutions generated at a certain temperature must be sufficiently high. A neighborhood solution corresponds to a small displacement of an atom in the material in order to obtain a new configuration of the solid. The difference between the objective function value of a newly generated neighborhood solution and the current solution is defined by δ , which is the resulting change in energy in the case of the physical annealing process. If $\delta < 0$ the neighborhood solution becomes the current solution, but if $\delta \geq 0$ the newly generated feasible solution has a probability of $e^{-\delta/T}$ of being accepted as the new current solution. So, SA accepts with certain probability feasible solutions which also increase the value of the objective function value. Ideally, this acceptance probability must be close to one at high temperatures at the beginning of the cooling process and is nearly reduced to zero at a temperature close to zero near the end of the process. This prevents SA from being trapped in a local minimum.

A main problem inherent in the application of SA is the determination of the generic parameters to be used. This set of parameters is part of the so-called cooling scheme. The problems according to the choices of cooling schemes are extensively treated by Hajek (1988). Classifications of cooling schemes are proposed by Collins et al. (1988) and Van Laarhoven and Aarts (1987).

As far as our implementation is concerned, the following choices have been made. In order to determine the value of the *initial temperature*, we use the approach proposed by Johnson et al. (1989). The initial temperature T_{begin} is computed by solving the expression

$$P_a = e^{-\Delta C/T_{begin}}$$

and hence

$$T_{begin} = \frac{-\Delta C}{\ln P_a} \quad (6.1)$$

Here ΔC represents the average deterioration value, which is computed as the cumulated value of the values of all worsening moves possible from the initial

solution, divided by the number of moves which cause a deterioration of the objective function value. Parameter P_a represents the acceptance fraction, i.e. the ratio of the accepted to the total number of generated moves.

The *cooling function* we use for the reduction of the temperature is the simple geometric function. The temperature at iteration t , T_t , is obtained from the temperature of the previous iteration as follows:

$$T_t = R \cdot T_{t-1} \quad (6.2)$$

Here, R represents the cooling rate.

The principle of an epoch is used to determine the *thermal equilibrium* at each temperature (see Golden and Skiscim (1986) and Teodorovic and Pavkovic (1992)). Therefore, the value of the current solution is stored every s moves. If the difference between the value of the last epoch and all previously saved epochs is less than $t\%$, then thermal equilibrium is reached and the temperature can be decreased according to the cooling function 6.2. Here, the values of s and t are both set to 10.

The *stopping criterion* is satisfied if the percentage of accepted moves is inferior to a critical acceptance ratio for a predefined number of five consecutive temperature values. However, each time a new best solution is obtained, the counter is reset to zero. This stopping criterion has been proposed by Johnson et al. (1989).

In contradiction to the complete evaluation of moves performed in the implementations for the LI and TS heuristics proposed, the SA heuristic generates a neighborhood solution on a stochastic base. This implies that the routes and the stops required to perform a move are selected at random. In order to narrow the search space, we provide the option of a range delimiter to prevent the selection of too many bad moves with respect to the objective function value. The range delimiter is a travel time restriction between the stops of two different routes selected at random for the move. This travel time restriction is adapted to the problem considered. It is computed for each solution with K routes R_1, \dots, R_K , first by determining for each stop $i \in N \setminus \{0\}$, belonging to route R_k , the travel time to its nearest neighbor stop belonging to another route R_l .

$$NN_{i \in R_k} = \min_{j \in R_l, k \neq l} d_{ij} \quad (6.3)$$

Subsequently, the range delimiter D is set equal to the largest travel time of the set of travel times previously defined.

$$D = \max_{i \in N \setminus \{0\}} NN_i \quad (6.4)$$

This procedure assures the potential participation of each stop to a move.

Other implementations of SA for the VRP are proposed in the literature. Osman (1993) has proposed a hybrid SA/TS metaheuristic. This hybrid metaheuristic was adapted by Thangiah et al. (1994) to solve VRPs with time-windows and combined pick-ups and deliveries. In Teodorovic and Pavkovic (1992) SA is used for generating an initial solution as well as for improving it in the case of the VRP with stochastic demand. Robusté et al. (1990) and Alfa et al. (1991) integrated SA in an initial heuristic for the VRP.

The SA implementation for the VRP proposed by Van Breedam (1994) and Janssens and Van Breedam (1994) is for the greater part comparable to the one presented here. The good solutions obtained by the last authors for some classical test-problems confirm the quality of the SA heuristic proposed here.

Procedure

- Step 1:* The initial solution is the current solution and is saved.
- Step 2:* Determine the initial temperature T_{begin} by means of expression 6.1, taking account of the value for the acceptance ratio P_a . The current temperature T is set to the initial temperature T_{begin} .
- Step 3:* Generate a neighborhood solution by randomly selecting a move from the current solution.
- Step 4:* Compute δ as the difference in objective function value (total travel time) between the neighborhood solution and the current solution.
If $\delta > 0$ and $r \geq e^{-\delta/T}$, with r a pseudo-random number in $[0,1]$, then go to step 3.
- Step 5:* The neighborhood solution becomes the current solution. The neighborhood solution is saved if it is the best solution so far.
Check for thermal equilibrium, if required. Every 10 moves the value of the current solution, an epoch, is compared to all previously saved epoch values. If the deviation is superior to 10%, then no equilibrium is reached and go to step 3.
- Step 6:* Decrease the current temperature T by means of the cooling function 6.2. If the percentage of accepted moves is inferior to the critical acceptance ratio for more than 5 consecutive temperature reductions without obtaining a new best solution, then go to step 7 else go to step 3.
- Step 7:* The solution stored is the final best solution.

Parameters

Initial solution (S)

1. A bad initial solution.
2. A good initial solution.

As for the LI heuristic, this parameter is provided especially for evaluating the effect of the initial solution on the final solution. The initial solution is generated with one of the initial heuristics (see section 7.1).

String Length (K)

1. String length of 1 stop.
2. String length of 2 stops.

This parameter is irrelevant for the move of type SC.

Range delimiter (D)

1. No range delimiter.
2. Range delimiter for the random selection of a move. The range delimiter is computed as the largest travel time between two nearest neighbors of two different routes of the current solution.

Acceptance fraction (A)

1. 0.30
2. 0.50

Percentage of accepted moves of the initial solution. This parameter is used to determine the initial temperature (cfr. P_a in expression 6.1).

Temperature reduction fraction (R)

1. 0.70 (fast cooling).
2. 0.90 (slow cooling).

Fraction by which the temperature is reduced in the geometric temperature function 6.2.

Critical acceptance ratio (L)

1. 0.01
2. 0.05

Critical percentage of accepted moves beneath which the percentage of accepted moves has to drop for more than 5 temperature reductions without an improvement of the best solution, before the SA heuristic is stopped.

The initial solution S , the string length K and the range delimiter D can be considered as problem-specific parameters. The other three parameters are of a

more generic nature. Their values are chosen based on extreme but acceptable values published in literature.

6.4 Tabu Search metaheuristic

Tabu Search has been conceived by Glover (1986). Similar ideas were developed by Hansen (1986) who has proposed a steepest ascent/mildest descent heuristic. TS is based on the principles of intelligent problem solving. A fundamental element is the use of a flexible and dynamic memory structure i.e. the tabu list. The procedure of TS is simple. At each iteration, the best move is selected. If this move deteriorates the objective function value, it is only performed if the inverse move does not have the tabu status, i.e. if it is not in the tabu list. If it is in the list, then the next best move not in the tabu list is selected and performed. This process is repeated until a stopping criterion is reached.

The fundamental elements of the TS strategy are the tabu list, the aspiration criterion, the long-term memory and the stopping criterion.

The function of the *tabu list* as a short-term memory is to prevent cycling through repeated selection of inverse moves with the same stops. The type and the length of the tabu list largely depends on the problem considered. Too small a list cannot prevent cycling, while too big prohibits too many moves. Taillard (1991) proposes the use of a dynamic list length. A regression analysis to predict the list length for a VRP out of the number of stops, the number of vehicles and the capacity ratio is used by Osman (1991). This author relates the problem characteristics with the tabu list length.

The *aspiration criterion* is satisfied if the move performed on the current solution yields an improvement of the objective function value. The improving move is performed whether or not in the tabu list.

Other, less popular aspiration criteria are presented by Glover and Laguna (1993).

The *long-term memory* is aimed at diversifying the search procedure. Moves occurring frequently are penalized by subtracting a penalty value from the difference in objective function value between the current solution and the neighborhood solution. Besides a diversification, an intensification procedure can be used too.

The *stopping criterion* is most often determined as a fixed number of iterations or a fixed number of iterations without any improvement of the objective function value.

A good overview of TS and its applications is provided by Glover (1989), Glover (1990) and Glover and Laguna (1993).

With respect to the application of TS to the VRP, some publications are to

be noticed.

Pureza and França (1991), Osman (1991) and Osman (1993) use TS as an improvement heuristic with a type of move in which a relocation is considered as a special case of an exchange. Thangiah et al. (1994) implements a hybrid TS/SA metaheuristic for VRPs with time-windows using the λ -interchange mechanism described in Osman (1991). Taillard (1992) adds a diversification strategy to his implementation of a TS improvement heuristic for the VRP. In Semet and Taillard (1993) the use of a TS improvement heuristic is demonstrated for a practical case.

The implementation of Gendreau et al. (1992) is exceptional in so far that unfeasible solutions with respect to capacity and route-length constraints are temporarily allowed through the use of penalties. They use a move comparable to the SR type which has been embedded in a specific insertion heuristic.

Stewart et al. (1992) use a TS strategy for finding an initial solution through a repeated application of a generalized assignment heuristic.

The TS implementation proposed provides the option to use a static as well as a dynamic tabu list length. The tabu list contains records of three elements: the list position, the origin route and the string of stops moved. The list is implemented as a queue. At each iteration, the last move performed is added to the end of the list. Subsequently, the list is rearranged by removing the move at position 1 in the list and by pulling all subsequent moves one position forward in the list.

The main advantage of this type of list is its ease of adaptation to static and dynamic list lengths. Moreover, all possible lengths for the string of stops to be moved are allowed. Most existing TS implementations for the VRP (cfr. *supra*) of which we are aware cannot relocate or exchange more than one stop between routes. A drawback of this structure occurs when a move involves the exchange of two strings of stops, as it is the case with the SC and SE move type. Two separate list positions are required to administrate the move.

The TS implementation is provided with a long-term memory to allow diversification. To prevent stops from being too frequently selected for moves, a penalty function value is computed for each move. This amount is added to the objective function value of the neighborhood solution that would be obtained if the move was performed. The penalty function value for a move is determined first by computing the penalty value of each stop involved in the move. The penalty associated with a stop is computed with the method proposed by Taillard (1992). The penalty value for stop i is given by $W \cdot f_i$, in which f_i stands for the frequency of occurrence of a stop i in a move. A value for parameter W is obtained by selecting a value at random from the interval

$$(0.1 \cdot \Delta_t^{max} \sqrt{N \cdot k} ; 0.5 \cdot \Delta_t^{max} \sqrt{N \cdot k}) \quad (6.5)$$

Here, N and k represent the number of stops and the number or routes of the current solution, respectively. The factor $\sqrt{N \cdot k}$ can be considered as a norma-

lising factor because the frequency of occurrence of a stop in a move decreases with the size of the problem. The term Δ_t^{max} represents the largest improvement up to iteration t .

The use of a penalty function related to a stop is not appropriate for our implementation because the moves (2,0) or (2,2) would be a priori disfavored with respect to the moves (1,0) or (1,1). Consequently, we decided to associate a penalty function value with a move rather than with a stop. The penalty of a move is computed as the average of the penalties of all stops involved in a move.

The stopping criterion of our implementation is arbitrarily set to 500 iterations. This number is of secondary importance because the objective function value as a function of the computing time will be used as a basis for the heuristic analysis.

Procedure

- Step 1:* The initial solution is the current solution.
- Step 2:* Perform a complete evaluation cycle of moves of the current solution. Select the routes according to the route numbers and the stops according to the route sequence.
- Step 3:* If the best move found gives an improvement of the objective function value, then go to step 5.
- Step 4:* If a long-term memory is used, an additional penalty associated with the move must be taken into account when selecting the least deteriorating move. If the best move is in the tabu list then select the move which causes the least deterioration of the objective function value and which is not in the tabu list.
- Step 5:* Perform the move. The new solution becomes the current solution. If the new solution is better than the best solution so far, then save the new solution. Add the move to the tabu list if there is still some place left. Otherwise, discard the move.
- Step 6:* Rearrange the tabu list by removing the first move from the list and pushing the successive moves up forward by one position. Adapt the length of the tabu list, if required.
- Step 7:* Go to step 2 if the number of iterations is less than 500.
- Step 8:* The solution saved is the final best solution.

Parameters

Initial solution (S)

1. A bad initial solution.
2. A good initial solution.

As for the LI and SA improvement heuristics, this parameter is provided for

evaluating the effect of the initial solution on the final solution. The initial solution is generated with an initial heuristic.

String Length (K)

1. String length of 1 stop.
2. String length of 2 stops.

This parameter is irrelevant for the move of type SC.

Length of tabu list (L)

1. 10-20
2. 20-30
3. 30-40

In the case of a static tabu list length, only the lower bound of each interval is considered. For a dynamic list the length of the tabu list is randomly selected in the interval.

Iterations for adapting tabu list length (I)

1. $+\infty$ (static tabu list length)
2. 5
3. 25
4. 50

Long-term memory (G)

1. Implementation without long-term memory.
2. Implementation with long-term memory.

Two parameters, the initial solution S and the string length K are problem-specific parameters. The other parameters are more generic of nature. As for the SA implementation, the values chosen for the generic parameters are determined based on publications containing acceptable values.

Chapter 7

Analysis of the improvement heuristics

The analysis of improvement heuristics differs substantially from that of the initial heuristics. Improvement heuristics require an initial solution to start from. Consequently, the behavior of improvement heuristics depends to a greater extent on the structure of the neighborhood of the current solution than on the specific problem characteristics, as was the case for the initial heuristics. In the case of the improvement heuristics, the problem characteristics only have an indirect effect because they affect the neighbourhood structure.

The generic parameters of the improvement heuristics also have an influence on the neighborhood structure. Nevertheless, the relation between problem characteristics and the values of the generic parameters is hard to analyse. Generic parameters affect the choice of moves required to go from initial to final solution. The sequence of solutions between initial and final solution is called the trajectory (Pirlot (1992)). The sequential nature of the trajectory makes it difficult to trace the effect of the generic parameters on the final solution.

The above-mentioned arguments justify the use of a reduced test set instead of all problems out of the seven test sets used for the analyses of the initial heuristics. Moreover, the time required to realize the runs with the complete problem set for the improvement heuristics is most probably in disproportion to the transferability of the conclusions to other problems.

First, the reduced test set is presented. Subsequently, the results of the parametric analysis are described for each of the three improvement heuristics separately. Finally, comparative results of the three improvement heuristic are presented in the heuristic analysis.

7.1 Reduced test set

In order to build a reduced test set out of the seven test sets with a total of 420 test problems, the cluster analysis classification technique was used. For additional information on cluster analyses, the reader is referred to Anderberg (1973).

By means of the partitioning around medoids clustering method (see Kaufman and Rousseeuw (1989)), we succeeded to cluster the problems of each of the seven test sets around medoids. A medoid can be considered as the most representative problem of its cluster.

The dissimilarity measure used to compute the distance between the problems of a test set is composed of a measure of the spread across the total travel time of the solution and a measure of the spread on the total number of routes of the solution. This distance measure represents the dependent variable, i.e. the total travel time, as well as the structure of the solution, revealed by the number of routes. All solutions obtained with each heuristic for each of the 420 problems were used.

Satisfactory results for the cluster analysis were obtained by reducing the complete set of problems into a set of 15 problems. Each test set is represented by its two medoids. Only test set *P2* requires three problems for representing its 60 problems.

Table 7.1 contains the 15 representative problems. For each problem, the original test set is mentioned, together with a good and a bad objective function value of an initial solution. As mentioned in chapter 6, both solutions are required for the experiments to investigate the effect of the initial solution on the final solution.

The good initial solution is chosen as the first quartile of the distribution of the total travel time of the solutions obtained with all parameter combinations of the initial heuristics. The bad solution corresponds with the third quartile of this distribution.

The decision to choose neither the best nor the worst solution can be motivated as follows. On one hand, the best solution obtained with an initial heuristic can be optimal or close to optimal. Consequently, no further improvement can be found. On the other hand, the worst, or even a random solution contains for example such a high number of routes that the use of an improvement heuristic with a move of type SE or SC is disfavored a priori, due to the inability of both move types to reduce the number of routes.

Both the good and the bad initial solutions of problems 1 to 7 have the least possible number of routes with respect to the capacity constraints. This means that all vehicles are filled up to a 100% capacity utilization. This implies that a move of type SR cannot be performed, because no stop can be relocated from one route to another if no residual capacity is available in at least one route

of the solution. For the same problems, the SM gives the same solution as SE because only exchanges between routes are feasible. So, for problems 1 to 7 only moves of types SE and SC are meaningful.

For problems 8 to 15, all move types are allowed, because both the good and bad solutions of these problems have routes with residual capacity, allowing for relocation of stops between routes.

7.2 Parametric analyses of the improvement heuristics

For these parametric analyses, the AID-technique is used. In the following, the results of the AID-analyses of the three improvement heuristics are described separately.

7.2.1 Local Improvement heuristic

If the type of move is SE, SR or SM, 16 solutions are obtained by combining all four parameter values of the LI heuristic. Only 4 replications are available if the move is of type SC. The lower number of observations in the case of SC is caused by the fact that parameters K and F are of no relevance for this type of move. The AID-results for the SC move type are less reliable due to the insufficient number of replications.

Initial solution (S). The results of the AID-analyses reveal that the good initial solution gives the best final solutions for almost all problems. For most of these problems, the final solution obtained with the good initial solution is significantly better than the one obtained with the bad initial solution.

As far as the move of type SR and SM is concerned, a significantly better final solution is obtained with a bad initial solution only for problem 14. A possible explanation is that the minimum of the bad initial solution in which the heuristic descends is accidentally deeper than that of the good initial solution.

In general, the results indicate a high degree of dependency of the LI heuristic upon a good initial solution.

The processing time is larger in case of a bad initial solution because more evaluation cycles have to be completed.

String Length (K). For the type SE, a string of length $K=2$ gives the best solutions for all fifteen problems. For ten out of fifteen problems, $K=2$ is even significantly better. For the remaining problems 3, 6, 7, 8 and 15 some reasons can be mentioned for the fact that $K=2$ is better but not significantly better

Problem	Origin test set	Bad Solution	Good solution
1	G1	2562 (10) 3.1	1245 (10) 3.1
2	G1	2352 (10) 3.1	1703 (10) 24.2
3	G2	2726 (20) 43.0	1781 (20) 0.1
4	G2	3094 (20) 10.2	1671 (20) 2.0
5	G3	1266 (5) 61.0	1072 (5) 48.2
6	G3	1436 (5) 49.2	1105 (5) 64.0
7	P1	1332 (5) 44.0	1097 (5) 62.1
8	P1	1593 (7) 18.2	1251 (6) 43.0
9	P2	2544 (11) 15.1	2089 (11) 4.0
10	P2	1588 (11) 8.0	1202 (11) 4.2
11	P2	1948 (11) 6.2	1386 (11) 29.2
12	T1	2987 (20) 6.0	2212 (14) 21.0
13	T1	2129 (19) 2.1	1235 (12) 2.3
14	T2	3692 (20) 13.2	3011 (16) 2.2
15	T2	2175 (16) 16.1	1398 (12) 24.2

Table 7.1: The reduced test set including 15 problems. Besides the original test set, a good and a bad solution are mentioned. The solution cell shows the total travel time, the number of routes (between brackets) and CPU time in seconds on a second line.

than $K=1$. The route sequence imposed by mixed pick-ups and deliveries (problems 7 and 8) or the time windows (problem 15) can hinder the exchange of strings containing more than 1 stop. For problems 3 and 6 where only exchanges of equal string lengths are allowed, exchanges with $K=2$ are, in essence, equivalent to two exchanges with $K=1$.

Concerning the move of type SR, the substitutability of a relocation of $K=2$ stops with two relocations of $K=1$ stops is the primary reason for the insignificance of the effect of K on the final solution for all problems.

The results for the move of type SM, are a mixture of those of SE and SR.

The processing time required for $K=2$ amounts about twice the time necessary for $K=1$.

Selection strategy (P). The difference between the first and the best improvement hardly gives significantly different final solutions. No dominant strategy is observed even if only the value of P in the best solution is considered.

The observations made by Osman (1993) are comparable.

Evaluation procedure for string length $K > 1$ (F). No significant difference can be observed between the two procedures for the order of evaluation.

If the results of the four types of moves are compared, it is observed that for problems 8 to 15 the solutions of the SM type are always present in the significantly better move types. Moreover, for seven of these problems the best solution is obtained with the SM type. The SC and SE are mostly significantly worse for these problems.

For the problems 1 to 7, no significant differences are observed between the solutions of SC and SE. If only the move type in the best solution is considered, a dominance of the SE is noticed for five out of seven problems.

The results of the AID analyses for the LI heuristic indicates a clear tendency. The problem-specific parameters, the initial solution and the string length have an explainable behavior.

Conversely, the generic parameters have no significant effect on the final solution. No further conclusions can be drawn for this latter group of parameters.

7.2.2 Simulated Annealing metaheuristic

The AID-analyses for the SA heuristic are performed on 32 solutions per problem for move type SC and 64 solutions for move types SE, SR and SM. These solutions were obtained by combining the six parameters of the SA implementation. The different number of replications is explained by the fact the string length K is not relevant for move type SC.

Initial solution (S). The quality of the initial solution seems to be somewhat

less important for the SA heuristic than for the LI heuristic. The effect of the initial solution on the final solution is for a larger part of the problems significant. However, good initial solutions do not always produce the significantly better final solutions. For about a half of the problems, the significantly better final solutions are obtained with a good initial solution.

This demonstrates the greater independence of the SA heuristic of the quality of the initial solution. A reason for this phenomenon can be the great variability of the SA heuristic and its ability to escape from a local minimum fast.

String Length (K). The results for the string length K have to be moderated somewhat in the case of the SA metaheuristic. Due to the random selection of a move, there is no certainty on the proportion of moves with $K=1$ and with $K=2$. In an extreme case, all accepted moves could have the same string length.

The results reveal that for about half of the problems the string length $K=2$ gives significantly better solutions. Although the dominance of $K=2$ in the best solution is observed for most problems, any inferences are pure conjecture due to the above-mentioned uncertainty.

Range delimiter (D). A number of tendencies can be deduced from the AID-analyses of the 15 problems for the travel time restriction as range delimiter. A travel time restriction gives significantly better solutions in cases where the routes of the initial solution are well-separated and/or the only side-constraints are capacity constraints. In these cases, the travel time restriction somewhat prevents the selection of moves which cause a substantial deterioration of the objective function value.

An implementation without travel time restriction gives significantly better solutions for problems with a constraining geographic structure, which mostly contains clustered stops.

For most problems with time-windows, no travel time restriction is required to generate better and often even significantly better solutions. The time-windows play the role of range delimiter for these problems, because they hinder the selection of moves which considerably worsen the objective function value.

Acceptance fraction (A). The acceptance fraction has a significant effect on the final solution for only a very small part of the problems. No dominant value of A is observed, even if only the value of A in the best final solution is considered. The results are approximately comparable for the four types of move.

These results imply that the differences in the initial temperature caused by the different acceptance fractions, 0.30 en 0.50, only give significantly different solutions for a very limited number of problems.

Temperature reduction fraction (R). The significant effect of the cooling

rate R occurs only for a small number of problems. Nevertheless, by considering only the value of R in the best final solution, a dominance of $R=0.90$ is clearly observed. This indicates that a slow cooling mostly yields better solutions.

Our observations confirm the findings of Johnson et al. (1989).

Obviously, the processing time required for a slow cooling ($R=0.90$) is considerably higher than that for a fast cooling ($R=0.70$).

Critical acceptance ratio (L). There is no single problem of the four move types for which the critical acceptance ratio gives significantly different solutions. Moreover, for a large number of problems, both values of L give the same best solutions. This means that for these problems, the percentage of accepted moves during the last five temperature reductions of the process was already inferior to the lower of both values of L , i.e. 0.01.

As far as the type of move is concerned, it can be observed that the types SR and SM are significantly better for problems 8 to 15. This can be explained by the ability of both move types to reduce the number of routes. The advantage of the SM over the SR occurs when the number of routes is reduced to the minimal number where no further relocations can be performed. The SM can continue with the exchange of stops in such situations.

For problems 1 to 7, significantly better solutions are obtained with SE for problems 1 to 4 and with SC for problems 5 to 7. The significantly better solutions of the SC for problems 5 to 7 can possibly be related to the long routes of these problems. The exchange of entire route segments as is done with the SC can possibly increase the probability of obtaining a greater improvement in the objective function value.

As a conclusion, we can state that the problem-specific parameter, the initial solution, the string length and the range delimiter, usually have a significant effect on the final solution. Moreover, their effects can more or less be related to the problem characteristics.

As far as the generic parameters are concerned, no consistent significant effect is observed. Only for the parameter representing the cooling rate, it can be observed that a slow cooling is most often better. No clear dominant values can be observed either for the parameters representing the acceptance ratio for determining the initial temperature, or for the critical acceptance ratio.

7.2.3 Tabu Search metaheuristic

Combining the values of the five parameters of the TS heuristic gives rise to 96 solutions per problem if the move type is SE, SR or SM. For the move type SC, 48 solutions are available for the AID-analyses because the parameter string length, K , is not relevant.

Initial solution (S). For almost all problems of the four move types, the good initial solution gives significantly better final solutions. This high degree of dependency of the TS heuristic on the good quality of the initial solution can be explained by the traditional path followed by the objective function of the TS heuristic. First, the objective function value descends as far as possible. Once it arrived at the point where no further improvements can be performed, the TS heuristic tends to stagnate and starts to oscillate. Consequently, the quality of the initial solution is determining for the quality of the final solution, just like for the LI heuristic.

String Length (K). The analyses reveal that a string length of $K=2$ stops is significantly better than $K=1$ stop if a move with $K=2$ can, in essence, not be substituted by two moves with $K=1$. This means that $K=2$ gives significantly better solutions for the SE, particularly for problems where exchanges of unequal string lengths are allowed (problems 8 to 15).

For the SR move type, the string length has a significant effect on half of the problems. For most of these problems $K=1$ is significantly better. This is explained by the above-mentioned principle of substitution.

The string length $K=2$ is significantly better for five out of eight problems in the case of move type SM.

On average, the processing time for $K=2$ is twice as much as the one for $K=1$.

Length of tabu list (L). The effect of parameter L is significant only for a minority of problems. Even if only the value of L in the best solution is considered, no meaningful deduction can be made with respect to the length of the tabu list.

Iterations for adapting tabu list length (I). The problems for which the different values of parameter I give significantly different solutions are not numerous. Even by considering only the value of I in the best solution, no inferences can be made. For the move of type SC, a light preference in favor of the use of a dynamic tabu list is observed.

The limited effect of the tabu list length and its adaptation confirms the findings of Pureza and França (1991), Semet and Taillard (1993), Osman (1991) and Osman (1993).

Long-term memory (G). For almost all problems, it can be observed that an implementation with a long-term memory gives significantly worse solutions than the implementation without it. This can be explained by the fact that the use of a long-term memory hinders the selection of more favorable stops because they carry a high associated penalty. Consequently, the least favorable stops are involved in the move, which mostly cause a considerable worsening of the objective function value. This diversification strategy induces a greater variety in the path of the objective function value.

Gendreau et al. (1992) suggest that bad results obtained with a long-term function can possibly be caused by the inappropriateness of the constants 0.1 and 0.5 in formula 6.5.

The computing time for the implementation with a long-term memory is somewhat more time-consuming than the one without because of the calculation of the penalty function for every stop involved in a move.

With respect to the type of move, it can be observed that the SM is mostly included in the significantly best move types for the problems 8 to 15. For the problems 1 to 7, a light dominance in favor of the SE is observed.

The conclusion for the TS is much like these of the two previous heuristics. The significant effect of the problem-specific parameters can mostly be related to the problem characteristics. The most important deduction for this group of parameters is the dependency of the quality of the final solution upon that of the initial solution.

Only for one generic parameter reliable conclusions can be drawn. The use of the long-term memory gives significantly worse solutions for the majority of problems.

7.3 Heuristic analysis of the improvement heuristics

The heuristic analysis of the improvement heuristics is different from that of the initial heuristics. Initial heuristics have a clear stopping criterion, i.e. when all stops are routed and a feasible solution is obtained. Consequently, a comparison of the best solutions, independently of the computing times, is permitted for initial heuristics.

As far as the stopping criterion of the improvement heuristics is concerned, only the LI has one. The TS and SA metaheuristics have no unambiguous stopping criteria. The computing time of both heuristics highly depends on the value assigned to the parameters. Nevertheless, it remains difficult to estimate the processing time of the SA and TS heuristics. Moreover, the probability of finding a better final solution increases with the run time. A simple comparison of the final solution of the three improvement heuristics without taking into account the run time is not appropriate.

A point of comparison with the heuristic analysis of the initial heuristics, is that only the best final solution of each of the three heuristics for each problem is considered.




Figure 7.1: Example of the path of the objective function value for the LI, TS and SA heuristics. The dashed vertical lines represent the minima attained by one of the heuristics.

An alternative for comparing the improvement heuristics dynamically is required. The specific feature of the dynamic analysis is that not only the final solutions of the three improvement heuristics are compared, but also their intermediary solutions at various time points. Three time points are considered, corresponding to the time at which the best final solution of each individual heuristic is obtained.

The dynamic heuristic analysis is performed with the LI, SA and TS improvement heuristic for each of the four types of move and for each problem separately.

A statistical analysis of the different solutions is not meaningful, because it is not clear at which time point to perform the analysis without (dis)favoring one of the three heuristics.

An important analysis tool for the dynamic heuristic analysis is the graphical representation of the path of the objective function value of each heuristic versus computing time. An example is given in figure 7.1.

The path of the objective function of the LI heuristic is represented by a straight line connecting the initial with the final solution value.

With respect to the SA, the objective function value of the best solution and its run time are registered at each temperature reduction. Consequently, the SA is represented by a monotonic descending curve. The recording of the best solution found during a temperature reduction instead of the current solution at the time of the temperature reduction can be justified as follows. If the current solution at the time of the temperature reduction were saved, it would not be possible to know the best solution if it was reached before the moment of the temperature reduction. The drawback of this approach is that deteriorations of the objective function value are not registered. The only alternative to this

approach would have been to record every move performed. However, this would make the graph less clear. Moreover, we are primarily interested in the improvement instead of the worsening of the objective function value of the SA heuristic.

The objective function value of the TS heuristic is recorded at each iteration. In the graph, the path of the objective function value can be identified by its oscillating nature.

The three dashed vertical lines represent the three time points corresponding to the minima obtained by the individual heuristics. Beside the graph, a tabular representation is also proposed.

In what follows, the dynamic heuristic analysis is described for each type of move separately.

7.3.1 String Cross

Table 7.2 on page 119 presents the objective function values of the three heuristics, with move type SC, observed at the three time points considered. A number of interesting findings can be deduced from this table and the graphs C.1 to C.15 on pages 145 to 149.

It is observed that the LI heuristic never gives the best final solution. The LI heuristic reaches its minimum value for all problems at the first minimum and this final solution is better than the best solutions obtained by the SA and TS heuristics at that point.

The path of the objective function of the TS heuristic is similar to that of the LI heuristic, but delayed with a few seconds. The TS performs moves which improve the objective function value as long as possible. Consequently, the objective function value of the TS equals the final solution of the LI heuristic a few seconds later. This delay is caused by the updating of the tabu list and, if required, by the calculations for the long-term memory. The largest differences at the first minimum are limited to 7% if the same initial solution is used, but can rise to 98% (cfr. problem 4) if a different initial solution is used.

At the time of the second minimum, both metaheuristics give better solutions than the LI heuristic. For nine problems, the SA gives the best solution, for the remaining six the TS. For thirteen problems, the final solution attained at the second minimum corresponds to the best solution found at that time.

The SA gives the best solution for six problems at the time of the second minimum, the TS for eight. An ex-aequo is observed for problem 1.

The same indecisive situation is observed at the time of the third minimum. For problems 10 and 14 the best solution is obtained with the SA at minimum 2

and with the TS at minimum 3. For problems 5, 6 and 8 it is vice-versa. For the remaining ten problems, the metaheuristic which yields the best solution at the second minimum, also yields the best solution at the time of the third minimum.

These results do not permit us to indicate which metaheuristic is best for move type SC, even if only the final solutions are considered. For eight problems, the difference between both metaheuristics is 1% or less. For the remaining problems this difference never exceeds 4%.

By considering the graph of the path of the objective function value for the three heuristics, the following observation can be made.

The path for the TS is for most problems similar. First a sharp decline and then stagnation. The stagnation phase is explained by the fact that in absence of improving moves, the least worsening moves are selected. Consequently, radical changes in the path of the objective function value are not observed. The variations observed in the TS path of problems 1, 6 and 8 are caused by the diversification strategy through the use of the long-term memory.

The path of the objective function value of the SA heuristic is similar for the problems 5, 6, 8, 9 and 11. The common feature is the long time after which the current solution drops beneath the value of the initial solution. This can be explained by the fact that the moves executed at the begin of the process cause such a high deterioration of the objective function value that it takes quite a lot of moves before the current solution again approaches the initial solution. For the problems 13 to 15, the solution of the SA drops very fast. This can possibly be explained through the binding effect of time-windows, which prevents crossings inducing a large increase in travel time.

Table 7.2 illustrates the dynamic comparison of the objective function values of the LI, SA and TS heuristics with move type SC at the three time points corresponding with the minimum of each heuristic. The symbol '*' indicates which heuristic attains its minimal value after the given run time. The best solution of the three heuristics at each time point is printed in bold face. The column at the right of each cell contains the relative difference with respect to the best solution at that time point. Times are given in seconds on a 80486DX/33Mhz. processor. Tables 7.3 to 7.5 (cfr. infra), for move types SE, SR and SM, respectively, have a same legend.

7.3.2 String Exchange

Table 7.3 on page 121 contains the numerical dynamic comparison of the three improvement heuristics for the 15 problems. A number of observations can be deduced based on this table and on the graphs C.16 to C.30 on pages 150 to 154.

Problem		Initial	Minimum 1	Minimum 2	Minimum 3
1	Time		5	240	1006
	LI	1245	1233 (*)	1233	1233
	TS	1245	1245	1213	1192 (*)
	SA	1245	1245	1213 (*)	1213
2	Time		30	142	384
	LI	1703	1610 (*)	1610	1610
	TS	1703	1637	1580 (*)	1580
	SA	2352	2220	1882	1613(*)
3	Time		2	113	480
	LI	1781	1774 (*)	1774	1774
	TS	1781	1781	1768	1763 (*)
	SA	1781	1781	1770(*)	1770
4	Time		7	17	35
	LI	1671	1590 (*)	1590	1590
	TS	1671	1602	1577 (*)	1577
	SA	3094	3094	2365	1639(*)
5	Time		59	846	1413
	LI	1072	1019 (*)	1019	1019
	TS	1072	1022	988 (*)	988
	SA	1266	1266	1192	975 (*)
6	Time		83	3076	4286
	LI	1105	1051 (*)	1051	1051
	TS	1105	1083	1015 (*)	1015
	SA	1105	1105	1023	979 (*)
7	Time		67	502	2766
	LI	1097	1070 (*)	1070	1070
	TS	1097	1097	1040	1032 (*)
	SA	1332	1332	1053(*)	1053
8	Time		50	390	2184
	LI	1251	1198 (*)	1198	1198
	TS	1251	1247	1154 (*)	1154
	SA	1251	1251	1251	1106 (*)
9	Time		5	219	303
	LI	2089	2009 (*)	2009	2009
	TS	2089	2089	1940	1841(*)
	SA	2089	2089	1818 (*)	1818
10	Time		7	223	340
	LI	1202	1141 (*)	1141	1141
	TS	1202	1182	1106	1097 (*)
	SA	1588	1588	1102 (*)	1102
11	Time		37	148	680
	LI	1386	1303 (*)	1303	1303
	TS	1386	1313	1298	1261(*)
	SA	1386	1386	1215 (*)	1215
12	Time		42	151	157
	LI	2212	1830 (*)	1830	1830
	TS	2212	1876	1754 (*)	1754
	SA	2212	1904	1789	1784(*)
13	Time		6	64	395
	LI	1235	1127 (*)	1127	1127
	TS	1235	1198	1102	1064(*)
	SA	1235	1188	1062 (*)	1062
14	Time		11	83	380
	LI	3011	2857 (*)	2857	2857
	TS	3692	3692	2915	2640 (*)
	SA	3692	3692	2663 (*)	2663
15	Time		35	108	226
	LI	1398	1236 (*)	1236	1236
	TS	1398	1270	1179	1169(*)
	SA	1398	1331	1165 (*)	1165

Table 7.2: Heuristic analysis of LI, TS and SA with move type SC.

For thirteen out of fifteen problems, the LI heuristic is the first minimum. For problems 3 and 4, the final solution of the SA heuristic corresponds to the first minimum. For both problems, the LI heuristic departs from the bad initial solution. Consequently, the processing time of the LI heuristic is large due to the large number of moves to be performed. However, the differences observed between the SA and LI heuristics at the time of the first minimum are almost negligible (<1%).

For problems 12 and 14, the final solution of the LI heuristic is not the first minimum. At that time point, the intermediary solution of the SA is already better than the final solution of the LI heuristic. This can possibly be explained by the presence of time-windows in both problems, which prevent moves increasing the total travel time considerably.

In case the final solution of the LI heuristic is the best solution at the time of minimum 1, it is observed that the results of both metaheuristics can be up to 7% worse with the same initial solution and up to 56% with a different initial solution.

The final solution of the SA heuristic is obtained for ten problems at the time of the second minimum. Only for three of these problems, the best final solution of minimum 2 remains the best final solution at the third minimum time point.

The best final solution, independently of the computing time, is obtained with the SA heuristic for only five problems. The best overall final solution for all remaining problems is obtained with the TS heuristic. However, for most problems the difference between the final solutions of the SA and TS heuristics does not exceed 5%.

The best solution for the TS is mostly obtained after the longest run time of the three heuristics. For most problems the specific path for the objective function of the TS heuristic is observed, i.e. descending as long as possible, followed by stagnation and oscillation. For most of these problems, a solution slightly worse than the final solution has already been reached a large number of iterations before. Consequently, a more efficient stopping criterion for the TS heuristic could be the termination after a fixed number of iterations without improvement of the current best solution.

The path of the TS for the problems 9 and 12 is highly variable due to the use of the long-term memory.

For the time-window problems 12 to 15, the specific path of the objective function of the SA is observed. The SA reaches its final solution very fast, probably due to the time-windows which prevent the selection of very deteriorating moves. The final solutions for two out of four time-windows problems are obtained with the SA. The TS heuristic yields the best final solution for the two remaining problems. However, the time required for the TS to reach its best final solution is high in comparison to that of the SA.

Problem		Initial	Minimum 1		Minimum 2		Minimum 3
1	Time		56		306		4120
	LI	1245	1158 (*)		1158	0.01	1158
	TS	1245	1199	0.03	1155	0.01	1143 (*)
	SA	1245	1245	0.07	1144 (*)		1144
2	Time		68		421		3503
	LI	1703	1609 (*)		1609	0.01	1609
	TS	2352	2193	0.36	1903	0.20	1591(*)
	SA	1703	1703	0.06	1585 (*)		1585
3	Time		138		177		638
	LI	2726	1756	0.01	1756(*)	0.01	1756
	TS	1781	1762	0.01	1762	0.01	1751 (*)
	SA	1781	1752 (*)		1752		1752
4	Time		169		193		265
	LI	3049	1487	0.01	1487(*)	0.01	1487
	TS	1671	1495	0.01	1495	0.01	1476 (*)
	SA	1671	1486 (*)		1486		1486
5	Time		249		4417		5346
	LI	1072	1030 (*)		1030	0.03	1030
	TS	1072	1037	0.01	998 (*)		998
	SA	1072	1072	0.04	1072	0.07	1014(*)
6	Time		172		2501		4595
	LI	1105	1037 (*)		1037	0.05	1037
	TS	1105	1067	0.03	984 (*)		984
	SA	1105	1105	0.06	1105	0.11	1037(*)
7	Time		69		888		1164
	LI	1097	1096 (*)		1096	0.04	1096
	TS	1097	1097	0.01	1067	0.01	1038 (*)
	SA	1097	1097	0.01	1055 (*)		1055
8	Time		680		2081		2667
	LI	1251	1200 (*)		1200	0.05	1200
	TS	1251	1209	0.01	1157	0.02	1114 (*)
	SA	1251	1251	0.04	1137 (*)		1137
9	Time		72		207		387
	LI	2089	1852 (*)		1852	0.03	1852
	TS	2543	1907	0.03	1800 (*)		1800
	SA	2089	2544	0.37	2207	0.11	1799 (*)
10	Time		33		323		2229
	LI	1202	1098 (*)		1098	0.01	1098
	TS	1202	1141	0.04	1114	0.02	1052 (*)
	SA	1588	1564	0.42	1094 (*)		1094
11	Time		131		359		372
	LI	1386	1250 (*)		1250	0.03	1250
	TS	1386	1269	0.01	1238	0.02	1174 (*)
	SA	1948	1948	0.56	1217 (*)		1217
12	Time		105		259		5021
	LI	2212	1955(*)	0.04	1955	0.05	1955
	TS	2212	2035	0.06	1988	0.06	1877(*)
	SA	2212	1926		1866 (*)		1866
13	Time		48		215		2765
	LI	1235	1108 (*)		1108	0.05	1108
	TS	1235	1168	0.05	1116	0.07	1019 (*)
	SA	1235	1149	0.04	1046 (*)		1046
14	Time		36		280		3022
	LI	3011	2868(*)	0.03	2868	0.12	2868
	TS	3011	2890	0.04	2769	0.08	2443 (*)
	SA	3011	2781		2558 (*)		2558
15	Time		116		2078		1843
	LI	1398	1219(*)	0.05	1219	0.07	1219
	TS	1398	1259	0.08	1215	0.07	1149(*)
	SA	1398	1161		1135 (*)		1135

Table 7.3: Heuristic analysis of LI, TS and SA with move type SE.

7.3.3 String Relocation

For the move of type SR, only problems 8 to 15 yield meaningful solutions. Based on the graphs C.31 to C.38 on pages 155 to 157 and on table 7.4 on page 123, a number of deductions can be made.

Again, the dominance of the LI heuristic is observed at the time of the first minimum. The beginning of the path associated with the TS heuristic closely resembles that of the LI, i.e. a descending path. However, the path of the TS is delayed in time in comparison with that of the LI due to the tabu list management. The final solution of the LI is on average up to 2% better than the intermediary solution of the TS. Exceptionally for problem 12, this difference amounts 7%.

Special mention needs to be made for problem 9. The TS heuristic is blocked after performing a move which reduces the number of routes from 11 to 10. No move of type SR can be performed any more with 10 routes. Consequently, the TS heuristic stopped prematurely. The LI with move type SR can avoid this exceptional situation by using the first improvement selection strategy. The SA heuristic also avoids the ultimate selection of the move which reduces the number of routes.

For problems 10 and 11, the same phenomenon is observed, be it that the TS heuristic becomes blocked at a much later time point.

For most problems, the LI yields the best solution at minimum 1. For problem 13 and 14, however, the intermediary solution of the SA heuristic is already better than the final solution of the LI heuristic. This can probably be explained by the effect of time-windows which prevent the selection of highly deteriorating moves. This is even more true in the case of SR than in the cases of SC and SE, because a relocation is difficult to perform in the case of time-windows. The path of the objective function value for the SA is comparable for all time-window problems 12 to 15.

For six out of eight problems, the final solution at minimum 2 corresponds with that of the SA. For five of the six problems mentioned (11 to 15), this solution is the best at that second time point. This not true for problems 8 and 10. By analyzing the graphical representation, it is observed that it takes quite a long time for the objective function value to descend beneath the initial solution value. For problem 8 this can probably be explained by the route sequence imposed by the mixed pick-ups and deliveries, which makes relocation more difficult.

The membership of the best final solution at the third minimum is perfectly balanced between the SA and TS heuristics. The computing time at minimum 3 is usually higher if the final solution is obtained by the TS heuristic. As mentioned earlier, a stopping criterion in the sense of a fixed number of stops without improvement or a range delimiter for limiting the number of moves evaluated would reduce the run time for the TS heuristic. The idea of a range delimiter for the TS has also been proposed by Semet and Taillard (1993). The stagnation phase, successive to the decline phase observed in the path of the objective function of the TS heuristic is not observed for the SA. The stopping criterion used for the SA heuristic prevents stagnation.

Problem		Initial	Minimum 1	Minimum 2	Minimum 3
8	Time		99	823	2279
	LI	1251	1128 (*)	1128	1128
	TS	1251	1132	0.01	1067 (*)
	SA	1251	1251	0.11	1251
9	Time		9	10	142
	LI	2089	1969 (*)	1969	1969
	TS	2089	2017	0.02	2017(*)
	SA	2089	2089	0.06	2089
10	Time		50	261	268
	LI	1202	1108 (*)	1108	1108
	TS	1202	1120	0.02	1097 (*)
	SA	1202	1202	0.08	1108(*)
11	Time		67	167	372
	LI	1386	1252 (*)	1252	1252
	TS	1386	1267	0.01	1263
	SA	1386	1343	0.07	1220 (*)
12	Time		47	234	381
	LI	2212	1844 (*)	1844	1844
	TS	2212	1968	0.07	1699
	SA	2212	2078	0.013	1616 (*)
13	Time		23	141	349
	LI	1235	1134(*)	0.00	1134
	TS	1235	1156	0.02	1045
	SA	1235	1133		983 (*)
14	Time		244	804	1312
	LI	3692	2681(*)	0.11	2681
	TS	3011	2691	0.12	2437
	SA	3692	2382		2350 (*)
15	Time		49	131	712
	LI	1398	1248 (*)	1248	1248
	TS	1398	1271	0.02	1170
	SA	2162	1301	0.03	1120 (*)

Table 7.4: Heuristic analysis of LI, TS and SA with move type SR.

7.3.4 String Mix

Similar to the SR, the SM is only meaningful if applied to problems 8 to 15. For problems 1 to 7, the move of type SM gives the same solutions as obtained with the SE. Following deductions result from the analysis of table 7.5 on page 127

as well as the graphs C.39 to C.46 on pages 158 to 160.

For six out of eight problems, the first minimum corresponds to the final solution of the LI heuristic. At that moment the final solution of the LI heuristic is the best solution of the three heuristics for only four out of eight problems. For problems 12 and 13 the SA solution is already better.

For problems 14 and 15, the final solution of the SA heuristics is obtained at minimum 1. This solution is also better than that of the LI heuristic at that moment.

These observations for problems 12 to 15 can probably be explained by the binding effect of time-windows, which limits the selection of moves, which cause a high deterioration of the objective function value in the beginning of the SA process.

Moreover, the run time of the LI heuristic is more affected in cases where many improvements have to be performed. Due to the random selection of moves, the run time of the SA heuristic is less affected.

The differences between the solutions of the three heuristics at the first time point amount at most to 11% if a same initial solution is used and 49% for different initial solutions.

At minimum 2, the SA has reached its final solution for seven out of eight problems. The exception is problem 8, for which the TS heuristic obtains its final solution. The path of the SA heuristic for problem 8 shows that it takes a long time for the objective function value to descend beneath the initial solution value. As mentioned for the SR move type, this is probably due to the route sequence imposed by the mixed pick-ups and deliveries, which hinders easy random selections of moves.

While considering minimum 3, it is observed that the TS heuristic yields the final best solution for six problems. Otherwise, the SA solution is better.

Again, the typical path of the objective function value is observed for the TS, i.e. a descent followed by a stagnation. The same observations for the TS with a move of type SM have been made by Pureza and França (1991).

Again, the path of objective function value of the TS permits a formulation of recommendations with respect to the stopping criterion and the range delimiter in order to speed up the TS heuristic.

The differences between the final solutions of the SA and TS heuristics, never exceed 4%.

The conclusions by Osman (1993) of the implementation of the three improvement heuristics equipped with his λ -mechanism, are not different from ours.

7.4 Conclusions for the analyses of the improvement methods

The analyses of the improvement heuristics were performed on a reduced test set of fifteen problems. The limited number of problems is justified by the limited external validity of the results, caused by the generic parameters of the improvement heuristics.

The behavior of the parameters of the improvement methods turned out to be polarised.

The observations made for the problem-specific parameters revealed that their effect is mostly significant and can be related to the problem-characteristics.

Important conclusions are the high dependency of the TS and LI heuristics on the quality of a good initial solution and the significantly better solutions obtained with a move with a string of two stops for situations where it cannot be substituted by two moves with a string of one stop each.

Especially for the SA heuristic, the range delimiter gives better solutions when used for problems with a homogeneous spread of stops and/or well separated routes without time-windows.

With reference to the type of move used, a slight dominance of the SM is observed. For problems 1 to 7, where no relocations were possible, the SM occurs in the shape of the SE.

As far as the behavior of the genetic parameters is concerned, it is observed that their effect is mostly not significant and that their values can hardly be related to the problem characteristics. However, for two parameters some deductions could be made. The first, the cooling rate of the SA heuristic gives better results with slow cooling.

Second, the results demonstrate that the use of the long-term memory predominantly gives significantly worse results.

The relative behavior of the three heuristics has been analyzed by means of a dynamic comparison of the objective function values at three time points. These time points correspond to the minimal value of the objective function value of each of the three heuristics. The main findings of this analysis can be summarized as follows.

The final solution of the LI heuristic is usually the first minimum encountered. This LI solution is mostly better than the intermediary SA and TS solution at that time point. Only for some problems with time-windows, the solution of the SA is already better at the first minimum.

For the time points corresponding to minima 2 and 3 the best solutions are obtained with the TS and SA heuristics. The results of the analysis do not show a preference of one over the other.

The SA heuristic tends to produce a final solution in less run time than the TS heuristic. The specific path of the objective function value of the TS value indicates that the heuristic can be speeded up by stopping after a fixed number of iterations without improvement and by using a range delimiter for restricting the number of moves to be evaluated at each iteration. Moreover, the difference of the final solutions between the SA and TS rarely exceeds 4%.

A final remark is the comparison of the final results of the initial with those of the improvement heuristics for the fifteen problems of the reduced test set. A remarkable observation is that the best final solution obtained with one of the three improvement heuristics is better than the best solution obtained with one of the initial heuristics only for eight out of the fifteen problems. This indicates that the quality of the initial heuristics remains very important. Moreover, the initial heuristic generates an initial solution in a very short time as opposed to that of the improvement heuristic. In addition, the TS and LI heuristics are highly dependent on the quality of an initial solution. All these arguments allow us to conclude that it is worth spending much effort on conceiving a good initial heuristic, while subsequently this good initial solution can further be improved by an improvement heuristic if sufficient run time and computer resources are affordable.

Problem		Initial	Minimum 1	Minimum 2	Minimum 3	
8	Time		150	735	2624	
	LI	1251	1125 (*)	1125	1125	
	TS	1251	1137	0.01	1069 (*)	1069
	SA	1251	1251	0.11	1251	0.17
9	Time		25	119	1181	
	LI	2089	1876 (*)	1876	0.03	1876
	TS	2089	1966	0.05	1856	0.02
	SA	2089	2086	0.11	1818 (*)	1818
10	Time		85	364	669	
	LI	1202	1064 (*)	1064	0.02	1064
	TS	1202	1106	0.04	1049	1045 (*)
	SA	1588	1588	0.49	1063(*)	0.02
11	Time		119	349	2261	
	LI	1386	1205 (*)	1205	0.01	1205
	TS	1386	1261	0.05	1225	0.02
	SA	1386	1386	0.15	1187 (*)	1187
12	Time		251	350	2654	
	LI	2212	1770(*)	0.04	1770	0.04
	TS	2212	1809	0.06	1770	0.04
	SA	2212	1700		1696 (*)	1696
13	Time		82	165	349	
	LI	1235	1058(*)	0.01	1058	0.06
	TS	1235	1088	0.04	1050	0.06
	SA	1235	1044		998 (*)	998
14	Time		171	183	1915	
	LI	3692	2645	0.09	2645(*)	0.09
	TS	3011	2784	0.14	2725	0.12
	SA	3011	2435 (*)		2435	2435
15	Time		119	131	1903	
	LI	1398	1182	0.06	1182(*)	0.06
	TS	1398	1242	0.11	1202	0.08
	SA	1398	1118 (*)		1118	1118

Table 7.5: Heuristic analysis of LI, TS and SA with move type SM.

Chapter 8

Synthesis and guidelines for further research

The primary contribution of this work is the analysis of a set of heuristics for solving the VRP. The behavior of eleven initial and three improvement heuristics has been analysed on a parametric and on a heuristic level. The findings can be summarized as follows.

For the analysis of the behavior of the *initial heuristics*, seven test sets each comprising sixty problems were produced, based on a geographical basic set of sixty problems. This basic set was constructed on three criteria: the location of the depot, the grouping of the customers, and the spreading of the customers. The addition of customer-related, vehicle-related and time-related side-constraints to this basic set, gave rise to seven test sets of each sixty problems.

The test sets were kept as deterministic as possible to ensure the internal validity of the results, i.e. analysing what has to be analysed. The main disadvantage of the use of a deterministic problem set turned out to be the existence of pathological problems. These pathological problems are a threat to the internal validity of the results.

Automatic Interaction Detection has proven to be a valuable analytical tool for the *parametric analyses*. AID allows the detection of the significant effect of a parameter on the solution value. Moreover, the sets of significantly better values of a parameter can be distinguished.

Table 8.1 contains an overview of the parametric analysis of the initial heuristics for the seven test sets. The main conclusion of the parametric analysis is that, in essence, the parameters of the initial heuristics can be classified into

two separate groups.

The first group of parameters represents parameters containing the combinations of the continuous weights of the selection criteria of the initial heuristics. The selection criterion refers to either the savings criteria, the insertion and selection criteria or the assignment costs.

The common characteristics of these parameters is that their significantly better values show consistency throughout the seven test sets. The influence of the geographical structure and/or the type of side-constraint on the significantly better values is rather limited. This indicates a certain problem independency of the values of these parameters.

The transferability of the findings for these parameters to other comparable problems is not excluded. Moreover, the significantly better values are consecutive values for a number of these parameters. This observation allows us to suggest that the intermediary values between the significantly better values are also significantly better. Before stating this, however, further research is required.

Most parameters of the second group demonstrate a higher degree of problem dependency. The significant effect and the significantly better values of these parameters vary with the side-constraints and/or with the geographical structure of the VRP. This group contains, among others, parameters representing the initialization procedures for sequential and non-sequential heuristics.

Due to its nature, the sequential initialization procedure tends to have less influence on the solution value than the simultaneous initialization procedure. Sequential heuristics pursue maximal capacity utilization of the vehicles, while non-sequential heuristics act upon the structures of the routes and thus more directly on total travel time. This can also be deduced from the analyses: the simultaneous initialization procedures have a significant effect on the solution value for almost all problems. However, their significantly better values do depend on the problem characteristics and/or the side-constraints. Nevertheless, conclusions can be drawn at a higher level of aggregation than at the individual problem level.

As far as the sequential initialization procedure is concerned, not only its significantly better values but also the significance of its effect is dependent on the problem characteristics and/or the side-constraints. However, this dependency-relation is not always clear.

At worst, conclusions for the other parameters of this second group are only possible at the lowest aggregation level, i.e. the individual problem level.

In essence, the findings of the parameters of this second group are less transferable than those of the first group. Transferability is limited to similar problems with respect to the problem characteristics and/or the side-constraints.

The *heuristic analysis* of the initial heuristics is essentially a statistical ana-

Heuristics	Parameters	G1	G2	G3	P1	P2	T1	T2
SN	I	-	-	-	-	-	+	+
	P	--	--	--	--	-+	-+	-+
	A	--	--	--	+	--	--	--
	C						+	+
PN	S	-+	+	-	-	-+	-+	-+
SS	I	-+	-+	-+	+	+	+	+
	P	-	-	-	-	+	-+	-+
	A	--	--	--	+	--	-	-
	C	++	++	++	++	++	++	++
PS	S	+	+	+	+	+	+	+
	C	++	++	++	++	++	++	++
GS	P	--	--	--	--	--	--	--
SI	I	--	-	-	-	-+	-+	-+
	M,L,E	++	++	++	++	++	+	+
PI	S	+	+	+	+	+	+	+
	M,L,E	++	++	++	++	++	+	+
PAI	S	+	+	+	+	+	+	+
	R	-+	-+	-+	-+	-+	-+	-+
	E	++	++	++	++	++	++	++
GA	P	-+	-+	-+	-+	-+	-	-
	S	+	+	+	+	+	+	+
	R	-+	-+	-+	-+	-+	-+	-+
	E	++	++	++	++	++	++	++
TP	I	+	+	+	+	+	+	+
	G	++	++	++	++	++	++	++
	L	++	++	++	++	++	++	++

Table 8.1: Summary of the parametric analysis of the initial heuristics for the seven test sets. Symbols:

"--": significant effect for a very limited number of problems;

"-": significant effect for a number of problems without common characteristics;

"-+": significant effect for a number of problems with common characteristics;

"+": significant effect for most problems with common characteristics;

"++": consistent significant effect for all problems.

lysis by means of the ordinal Friedman test. Only the best solution for each problem obtained with a heuristic is used. The design of the heuristic analysis permits us to indicate the significantly better heuristics for some problems with common characteristics. As a result, one can get an idea of the connection between the performance of the heuristics and the problem characteristics and/or the side-constraints.

The tables of chapters 3, 4 and 5 summarize the heuristic analysis for different problem characteristics and various side-constraints.

The findings of these analyses are transferable to problems with comparable characteristics. However, the condition is that the intermediary combinations of the values of the parameters, which we did not consider, do not result in significantly different rankings of the heuristics.

The choice made for the side-constraints, the geographical criteria and the values of the parameters are a threat to the external validity of the results. Therefore, more research is needed on problems with other characteristics and side-constraints. This includes problems with stochastic elements and heterogeneous side-constraints. Moreover, this will also afford an insight into the biasing effect of the pathological problems on the significance of parameters.

As far as the results of the analyses of the behavior of the *improvement heuristics* is concerned, the primary purpose was not the transferability of the results to other problems. This constraint was set by the generic parameters of the improvement heuristics. It proved to be difficult to relate the behavior of the generic parameters to the problem characteristics. These arguments, but also the high computing times, justified the use of a reduced test set containing only fifteen problems. The reduced test set was obtained by a cluster analysis applied to the set of all solutions obtained with the eleven initial heuristics for the 420 problems of the seven test sets.

The *parametric analysis* for the improvement heuristics with the AID technique revealed a classification of the parameters into two main groups. Table 8.2 summarizes the results of the parametric analysis.

The first group contains the problem-specific parameters, like the quality of the initial solution, the string length for the three improvement heuristics and the range delimiter for the SA heuristic. Even the type of move could be considered as a problem-specific parameter. The significant effect and the significantly better values of these parameters could more or less be related to the problem characteristics.

Therefore, the transferability of the findings for these parameters to other problems with comparable characteristics cannot be excluded.

The second group of parameters includes the generic parameters of the im-

Heuristics	Parameters	SC	SE	SR	SM	Aggregation
LI	S	+	+	++	++	+
	P	--	--	--	--	--
	F		--	--	--	--
	K		+	-+	-+	
SA	S	+	+	+	++	+
	A	--	--	--	--	--
	L	--	--	--	--	--
	R	-+	-+	-+	-+	-+
	D	+	+	-+	-+	-+
	K		+	-+	-+	
TS	S	++	++	++	++	++
	L	--	--	--	-	-
	I	--	--	--	--	--
	G	++	++	++	++	++
	K		+	-+	+	

Table 8.2: Summary of the parametric analysis of the improvement heuristics for the four types of move and the aggregated solutions. Symbols:
 ”--”: significant effect for a very limited number of problems;
 ”-”: significant effect for a number of problems without common characteristics;
 ”-+”: significant effect for a number of problems with common characteristics;
 ”+”: significant effect for most problems with common characteristics;
 ”++”: consistent significant effect for all problems.

provement heuristics. For a large number of problems, most of the generic parameters have no significant effect on the final objective function value. The portability of these findings to other problems is speculative, due to the fact that the values for the generic parameters can hardly be related to the problem characteristics. The parameter values were chosen as extreme as possible, taking account of what has not been categorized as bad in the literature. The results revealed that the variability of the final objective function value due to the values chosen was very limited.

Therefore, the purpose of the parametric analysis for the generic parameters was limited to the evaluation of differences between the values proposed.

Useful recommendations resulting from the analysis of the generic parameters are the slow cooling in the case of SA and the bad solutions obtained with the long-term memory in the case of the TS heuristic.

The *heuristic analysis* of the improvement heuristics is a dynamic analysis. The path of the objective functions of the three improvement heuristics are compared at three different time points, corresponding with the time at which

the best solution is obtained for each heuristic.

The main conclusion of the dynamic analysis is that the available run time determines the choice of the best improvement heuristic.

The heuristic analysis did not permit to indicate the better among both metaheuristics, SA or TS. The difference between their final solution never exceeded 4%.

The concept of the dynamic analysis proposed is transferable to other problems. The results obtained are only transferable to problems with almost identical characteristics. Nevertheless, it must be derived that due to the generic parameters, problems which are only slightly different with respect to their geographical structure and side-constraints, can induce a considerably different path of the objective function value of an improvement heuristic.

Describing the behavior of the generic parameters has proved to be a difficult task due to the indirect relation with the problem characteristics. Therefore, there is no guarantee that more extensive research would make their behavior more clear.

The dynamic heuristic analysis can be a useful tool for improving the implementation of the heuristics by analysing the path of the objective function value produced by the heuristics. This has been illustrated by the recommendations we proposed for speeding up the TS heuristic. Moreover, the opportunities for building hybrid metaheuristics, combining the best features of both metaheuristics can be evaluated. The TS provides the certainty of finding at least a local minimum, while the SA permits a fast shift of the search process towards other minima.

A final remark concerns the practical use of this research. This work can be considered as a first step towards the development of a diagnostic system for selecting the appropriate heuristic for solving the VRP. Such a diagnostic system must be conceived to determine the best heuristic or heuristics to solve a particular VRP.

For choosing the most appropriate heuristic, the results of the heuristic analysis can be used. Once a heuristic is found, the findings of the parametric analysis can be used to propose good sets of values for the parameters.

Finally, the results of the heuristic and parametric analyses of the improvement heuristics can be used in order to advise a good heuristic for enhancing the initial solution.

It is beyond all doubt that in order to design a workable and reliable diagnostic system many experiments are required. A way of doing this is by tracing the behavior of the parameters and heuristics when perturbing the deterministic problems proposed here increasingly. It is obvious that the construction of a diagnostic system may lead to further research.

Appendix A

The basic set of geographical problems

The basic set of 60 geographical problems is represented graphically. The set is constructed on three criteria: the location of the depot, the grouping of customers and the spreading of customers. Three different depot locations are considered: central, inside and outside. For the grouping of customers, five patterns are distinguished: singleton, clusters, 50% clusters, cones and 50% cones. Four patterns were taken into account for the spreading of customers: uniform, 50% central, concentric and compressed. Combining these patterns exhaustively give rise to a basic set of 60 geographic problems.

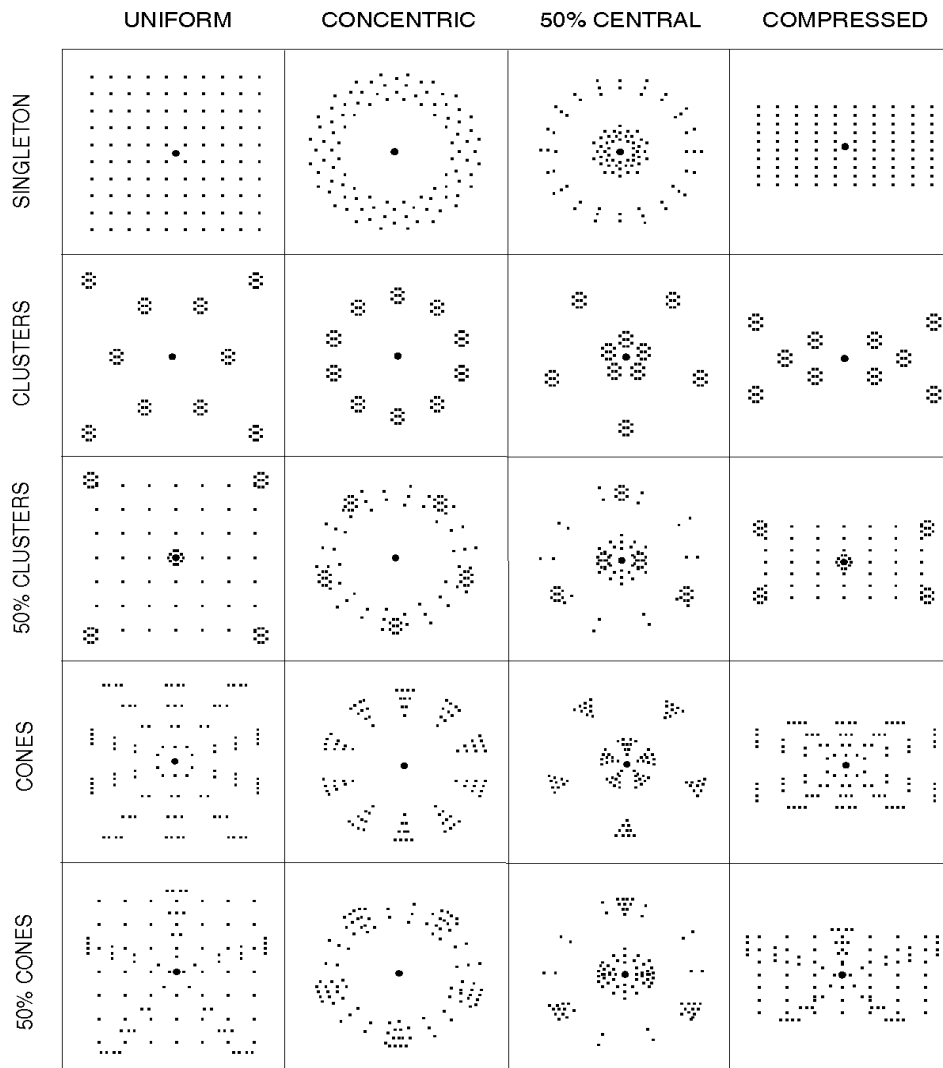


Figure A.1: Geographical problems of the basic set. Across: spreading patterns; down: grouping patterns. Depot location: central.

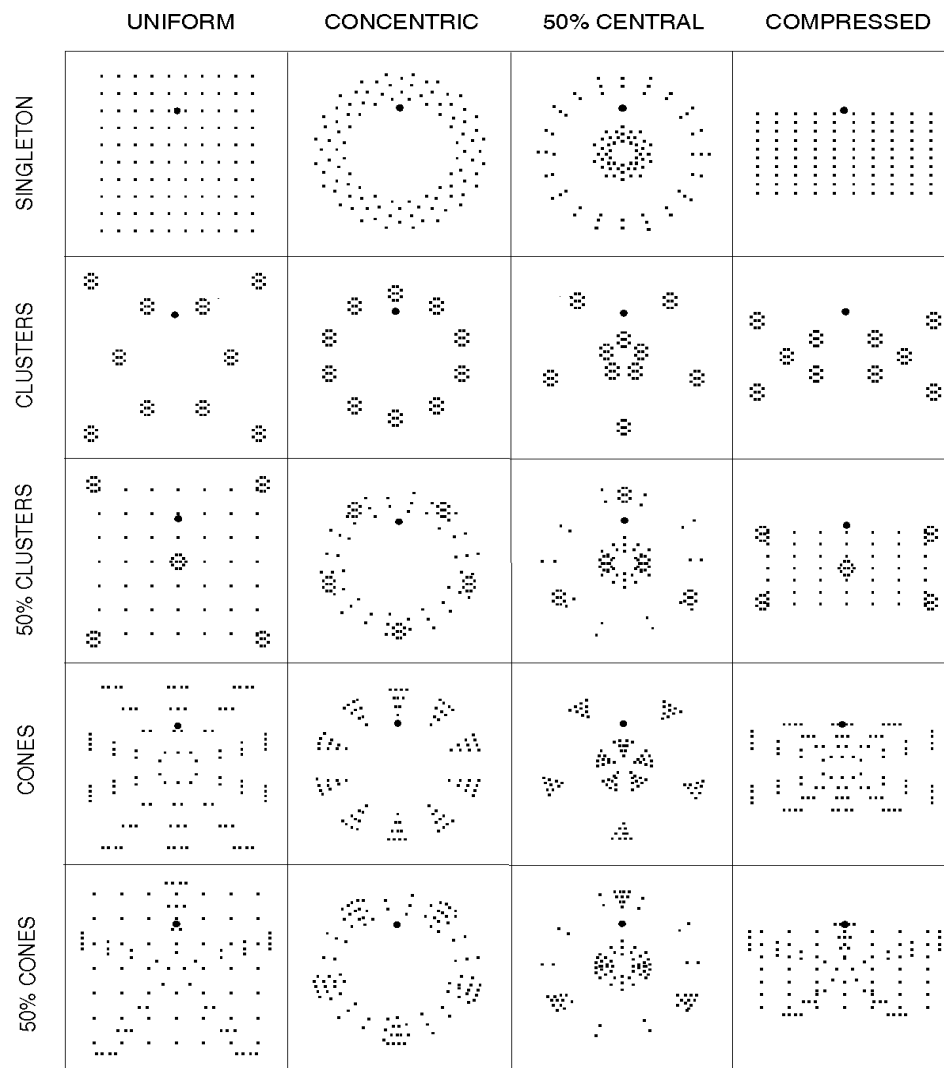


Figure A.2: Geographical problems of the basic set. Across: spreading patterns; down: grouping patterns. Depot location: inside.

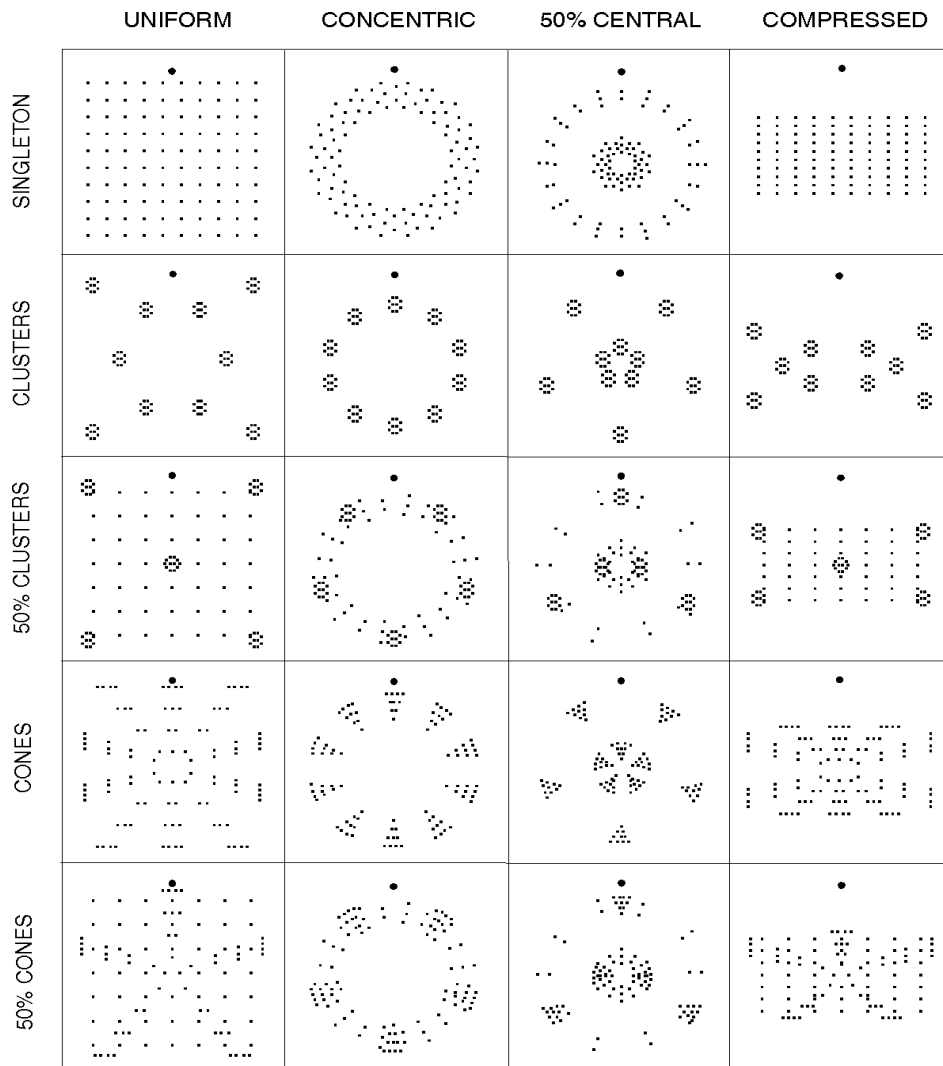


Figure A.3: Geographical problems of the basic set. Across: spreading patterns; down: grouping patterns. Depot location: outside.

Appendix B

Descriptions of some procedures and tests used

B.1 Seed point generation methods

A seed point can be considered as the geographical representation of a vehicle. Consequently, a capacity is associated with each seed point.

Two seed point generation methods are considered: the cone covering and the circle covering method. For this research, we consider only automatic seed point generation methods. This implies that seeds are generated without previously stating the number of seeds required.

Beside the automatic methods, procedures exist to generate a determined number of seed points (see among others Savelsbergh (1988), Fisher and Jaikumar (1981)).

An important drawback of the automatic seed point generation methods considered, is that they only take account of the capacity constraints. As a result, for heuristics like the GA heuristic, with separated assigning and route sequencing phases, TSPs can be unsolvable during the sequencing phase in the case of hard time-windows.

We are not aware of publication on automatic seed point generation methods for time-windows. The usual approach in such cases is to generate and position seed points iteratively. The number and the location of the seed points are determined iteratively by repeatedly solving the VRP (see Nygard et al. (1988)).

B.1.1 Circle Covering

The Circle Covering method has been conceived by Savelsbergh (1990b). It generates seed points that coincide with location of stops, i.e. seed customers. Essentially, it is designed to perform well on problems with clustered stops.

Procedure

Step 1: Associate a circle with each stop in such a way that the sum of the demands of the stops within the circle approximates as much as possible the vehicle capacity.

Step 2: Make a list with all stops in increasing (decreasing) order of the radius of their associated circle.

Step 3: Pick the next stop out of the list.

Step 4: If not all stops are covered by the circles selected, then go to step 3.

Step 5: The seeds are the centres of the circles providing the covering.

Parameters

Seed generation method (S)

1. Circle covering with selection of circles in order of their increasing radius.
2. Circle covering with selection of circles in order of their decreasing radius.

B.1.2 Cone covering

The Cone Covering procedure is originally based on the ideas of Fisher and Jaikumar (1981), and was also used by some other authors, among which Nygard et al. (1988).

A specific characteristics of the procedure is that seeds do not have to coincide with the location of stops. A drawback is the dependency of the method on the location of the depot.

Procedure

Step 1: Rank all stops in increasing order of their polar angle with respect to the depot.

Step 2: Locate the two successive stops in the list with the largest difference between their polar angles. The stop with the largest angle of this pair is selected to be linked with the depot. Hence, the starting position for a counterclockwise sweep procedure is determined. Rearrange the list

- of stops with reference to this stop selected. This stop is the first stop of the first cone.
- Step 3:* If the rearranged list of polar angles is empty, then the sweep procedure is finished and go to step 5. Otherwise, pick the next stop out of the list.
- Step 4:* If the stop selected can be added to the current cone without violating the capacity constraint, then add it and go to step 3. Otherwise, the stop selected becomes the first stop of a new current cone and go to step 3.
- Step 5:* Determine the bisectrice of each cone. The exact location of the seed on the bisectrice of each cone is determined based on the load fraction. The load fraction gives the proportion of the total demand of all stops in the cone which is positioned between seed point and depot.
- Step 6:* If seed points must coincide with customer locations, then the stop closest to the seed point is selected to become a seed customer.

Parameters

Seed generation method (S)

1. Cone covering with load fraction=0.05
2. Cone covering with load fraction=0.25
3. Cone covering with load fraction=0.50
4. Cone covering with load fraction=0.75
5. Cone covering with load fraction=1

B.2 Heuristic for the TSP

Some initial heuristics, among which the SW, TP, GA, PAI and GS heuristics require a heuristic for solving a TSP. The heuristic proposed is a simple insertion heuristic combined with a 3-opt improvement heuristic.

B.2.1 Insertion heuristic

The insertion heuristic yields a feasible solution for the TSP. The reader is referred to Rosenkrantz et al. (1977) for additional information on TSP heuristics. The heuristic starts with the depot as first stop and iteratively adds stops making use of an insertion and selection criterion. The insertion and selection criterion are equivalent and insert stop u after stop i in the route if the follo-

wing expression is satisfied:

$$\min_{u,i} [t_{iu} + t_{i+1u} - t_{ii+1}]$$

B.2.2 3-opt improvement heuristic

The solution for the TSP is improved using a *3-opt* branch exchange heuristic. A TSP route is 3-optimal if no better route sequence can be found by replacing 3 branches by 3 other ones of the same route. The *3-opt* heuristic is based on the *k-opt* procedure introduced by Croes (1958), Lin (1965) and Lin and Kernighan (1973).

Experiments showed that the *3-opt* gives the best trade-off between quality of solution and computing time (Christofides and Eilon (1969)). For a complete overview on the *3-opt* procedure with side-constraints, the reader is referred to Solomon et al. (1988) and Savelsbergh (1990a).

B.3 Post-processor

The post-processor handles unrouted stops which were not assigned to a route by a non-sequential initial heuristic.

The post-processor starts by selecting the unrouted stop farthest from the depot. An attempt is made to insert this stop in one of the routes, using the criterion of the minimal increase in route time. If the stop cannot be added to a route, a new route is started with the stop considered. All unrouted stops are iteratively assigned to routes using this procedure.

The priority given to the rerouting of stops far from the depot is inspired by the consideration that using a new route for servicing stops close to the depot is usually less expensive than for servicing stops far from the depot.

B.4 Automatic Interaction Detection

The AID technique has been conceived by Morgan and Sonquist (1963) and further refined by Sonquist et al. (1971). The technique is aimed at discovering the structure of the relation between variables. The dependent variable is continuous of nature, while the independent ones had to be nominally-scaled.

An AID solution can be represented by a tree structure containing all binary splits. At each phase, an analysis of variance is performed on each two groups of values of the same variable. The split with the highest significant *F*-value




Figure B.1: Example of a tree produced by the AID-analyses. The path of the best solution is indicated with a "*".

among the analyses of variance is selected. The splitting process is halted if no more binary split with a significant F -value can be found.

Figure B.1 contains an example of a solution of an AID analysis. Assume a heuristic with three parameters A , B and C . The symbol "*" refers to the path of the best solution. So, the significantly better values of each parameter are: $A=1$; $B=1,3$; $C=2,3$.

B.5 Friedman test

The Friedman test can be considered as the ordinal variant of the two-way analysis of variance. The test determines whether the rank totals of k treatments are significantly different. The data can be represented in a tabular form. Rows represent the b subjects (the test problems) and the columns the k treatments (the heuristics).

The hypothesis H_0 states that the k treatments all have the same effect. In case of the alternative hypothesis H_1 , there is at least one treatment with a larger rank total than at least one other treatment.

The test statistic is given by

$$T_2 = \frac{(b-1)([B_2 - bk(k+1)^2/4]}{A_2 - B_2}$$

with

$$A_2 = \sum_{i=1}^b \sum_{j=1}^k R_{ij}^2$$

$$B_2 = \frac{1}{b} \sum_{j=1}^k R_j^2$$

$$R_j = \sum_{i=1}^b R_{ij}$$

Here, R_{ij} stands for treatment j of subject i . The test statistic T_2 is approximately F -distributed with $k_1 = k - 1$ and $k_2 = (b - 1)(k - 1)$ degrees of freedom.

If H_0 is rejected, the rank sums are mutually compared. Treatments i and j are considered to be different if the following inequality is satisfied:

$$|R_i - R_j| > t_{1-\alpha/2} \sqrt{\frac{2b(A_2 - B_2)}{(b-1)(k-1)}}$$

where $t_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of the t distribution with $(b - 1)(k - 1)$ degrees of freedom.

The reader is referred to Conover (1980) and Siegel and Castellan (1989) for additional information on the Friedman test.

Appendix C

Graphs for the improvement heuristics

The following graphs represent the paths of the objective function values $f(R)$ of the three improvement heuristics, LI, TS and SA versus CPU-time T in seconds on an 80486DX processor at 33 Mhz.

The three dashed lines correspond with the time points at which the best final solution were obtained for each heuristic.




Figure C.1: String Cross for problem 1




Figure C.2: String Cross for problem 2




Figure C.3: String Cross for problem 3




Figure C.4: String Cross for problem 4




Figure C.5: String Cross for problem 5




Figure C.6: String Cross for problem 6




Figure C.7: String Cross for problem 7

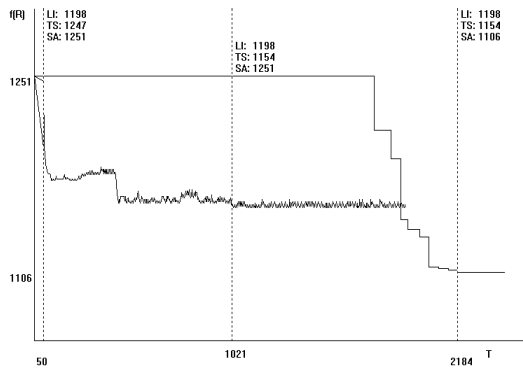


Figure C.8: String Cross for problem 8

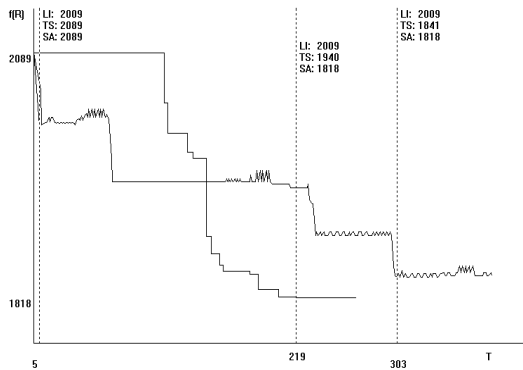


Figure C.9: String Cross for problem 9

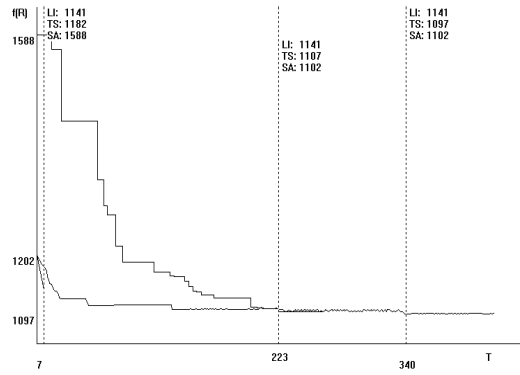


Figure C.10: String Cross for problem 10




Figure C.11: String Cross for problem 11




Figure C.12: String Cross for problem 12




Figure C.13: String Cross for problem 13

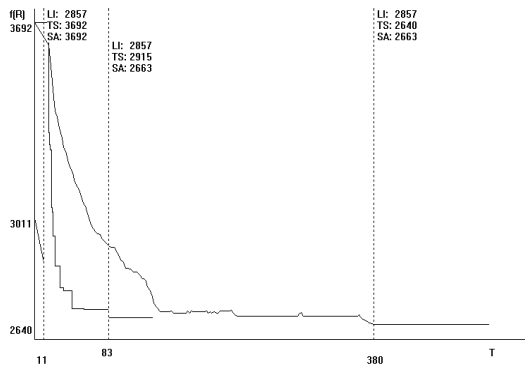


Figure C.14: String Cross for problem 14




Figure C.15: String Cross for problem 15




Figure C.16: String Exchange for problem 1




Figure C.17: String Exchange for problem 2




Figure C.18: String Exchange for problem 3




Figure C.19: String Exchange for problem 4




Figure C.20: String Exchange for problem 5

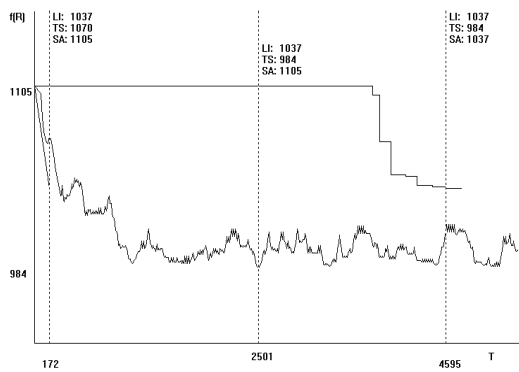


Figure C.21: String Exchange for problem 6

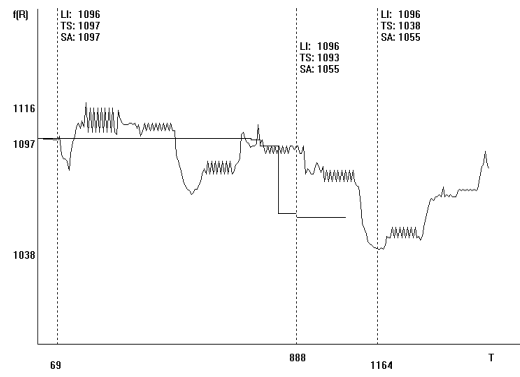


Figure C.22: String Exchange for problem 7




Figure C.23: String Exchange for problem 8




Figure C.24: String Exchange for problem 9




Figure C.25: String Exchange for problem 10




Figure C.26: String Exchange for problem 11




Figure C.27: String Exchange for problem 12




Figure C.28: String Exchange for problem 13




Figure C.29: String Exchange for problem 14




Figure C.30: String Exchange for problem 15




Figure C.31: String Relocation for problem 8




Figure C.32: String Relocation for problem 9




Figure C.33: String Relocation for problem 10




Figure C.34: String Relocation for problem 11




Figure C.35: String Relocation for problem 12




Figure C.36: String Relocation for problem 13




Figure C.37: String Relocation for problem 14




Figure C.38: String Relocation for problem 15




Figure C.39: String Mix for problem 8




Figure C.40: String Mix for problem 9




Figure C.41: String Mix for problem 10




Figure C.42: String Mix for problem 11




Figure C.43: String Mix for problem 12




Figure C.44: String Mix for problem 13




Figure C.45: String Mix for problem 14




Figure C.46: String Mix for problem 15

Bibliography

- Aarts, E. and J. Korst (1989). *Simulated annealing and Boltzmann machines*. Wiley, Chicester.
- Alfa, A., S. Heragu, and M. Chen (1991). A 3-opt based simulated annealing algorithm for vehicle routing problems. *Computers and Industrial Engineering* 21, 635–639.
- Altinkemer, K. and B. Gavish (1991). Parallel savings based heuristic for the delivery problem. *Operations Research* 39(3), 456–469.
- Anderberg, M. (1973). *Cluster analysis for applications*. Academic Press, New York.
- Baker, B. (1992). Further improvements to vehicle routing heuristics. *Journal of the Operational Research Society* 43(10), 1009–1012.
- Baker, E. and J. Schaffer (1986). Solution improvement heuristics for the vehicle routing and scheduling problem with time-window constraints. *American Journal of Mathematical and Management Sciences* 6(3-4), 261–300.
- Balakrishnan, N. (1993). Simple heuristics for the vehicle routing problem with soft time-windows. *Journal of the Operational Research Society* 44(3), 279–287.
- Bodin, L. (1983). Solving large vehicle routing and scheduling problems in small core. In *Proceedings of the ACM '83 Conference*, pp. 27–37.
- Bodin, L., 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(2), 63–211.
- Buxey, G. (1979). The vehicle scheduling problem and monte carlo simulation. *Journal of the Operational Research Society* 30, 563–573.
- Cerny, V. (1985). A thermodynamical approach to the travelling salesman problem: an efficient simulation algorithm. *Journal of Optimization Theory Appl.* 45, 41–51.
- Christofides, N. (1985). Vehicle routing. In E. Lawler, J. Lenstra, A. Rinnooy Kan, and D. Shmoys (Eds.), *The Traveling Salesman Problem. A Guided Tour of Combinatorial Optimization*, pp. 431–448. Wiley, Chicester.

- Christofides, N. and S. Eilon (1969). An algorithm for the vehicle-dispatching problem. *Operational Research Quarterly* 20(3), 309–318.
- Christofides, N., A. Mingozzi, P. Toth, and C. Sandi (1979). *Combinatorial Optimization*. Wiley, Chicester.
- Clarke, G. and J. Wright (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research* 12(4), 568–581.
- Collins, N., R. Eglese, and B. Golden (1988). Simulated annealing - an annotated bibliography. *American Journal of Mathematical and Management Sciences* 8, 209–307.
- Conover, W. (1980). *Practical nonparametric statistics*. Wiley, New York.
- Croes, A. (1958). A method for solving traveling salesman problems. *Operations Research* 5, 791–812.
- Dantzig, G. and J. Ramser (1959). The truck dispatching problem. *Management Science* 6, 81–91.
- Desrochers, M. and T. Verhoog (1989). A matching-based savings algorithm for the vehicle routing problem. Technical report, Cahier du Gerad G-89-04, Ecole des Hautes Etudes Commerciales de Montréal.
- Dror, M. and L. Levy (1986). A vehicle routing improvement algorithm comparison of a greedy and a matching implementation for inventory routing. *Computers and Operations Research* 13(1), 33–45.
- Eglese, R. (1990). Simulated annealing: a tool for operational research. *European Journal of Operational Research* 46, 271–281.
- Eilon, S., C. Watson-Gandy, and N. Christofides (1971). *Distribution Management: Mathematical modelling and practical analysis*. Griffin, London.
- Fisher, M. and R. Jaikumar (1981). A generalized assignment heuristic for the vehicle routing problem. *Networks* 11, 109–124.
- Gaskell, T. (1967). Bases for vehicle fleet scheduling. *Operational Research Quarterly* 18, 281–295.
- Gendreau, M., A. Hertz, and G. Laporte (1992). A tabu search heuristic for the vehicle routing problem. Technical report, Working Paper ORWP 92/14, Département de Mathématiques, Ecole Polytechnique Fédérale de Lausanne, Suisse.
- Gillet, B. and L. Miller (1974). A heuristic algorithm for the vehicle-dispatch problem. *Operations Research* 22, 340–349.
- Glover, F. (1986). The general employee scheduling problem: An integration of management science and artificial intelligence. *Computers and Operations Research* 15, 563–593.
- Glover, F. (1989). Tabu search: Part i. *ORSA Journal on Computing* 1, 190–206.

- Glover, F. (1990). Tabu search: Part ii. *ORSA Journal on Computing* 2, 4–32.
- Glover, F. and M. Laguna (1993). Tabu search. In C. Reeves (Ed.), *Modern Heuristic Techniques for Combinatorial Problems*. Blackwell Scientific Publications, Oxford.
- Goetschalckx, M. and C. Jacobs-Blecha (1986). The vehicle routing problem with backhauls. *European Journal of Operational Research* 42, 39–51.
- Golden, B. (1977). Evaluating a sequential vehicle routing algorithm. *AIIE Transactions* 9, 204–208.
- Golden, B. and C. Skiscim (1986). Using simulated annealing to solve routing and location problems. *Naval Research Logistics Quarterly* 33, 261–279.
- Golden, B. and W. Stewart (1985). Empirical analysis of heuristics. In E. Lawler, J. Lenstra, A. Rinnooy Kahn, and D. Shmoys (Eds.), *The Traveling Salesman Problem. A Guided Tour of Combinatorial Optimization*, pp. 207–246. Wiley, Chicester.
- Hajek, B. (1988). Cooling schedules for optimal annealing. *Mathematics of Operations Research* 13, 311–329.
- Hansen, P. (1986). The steepest ascent mildest descent heuristic for combinatorial programming. In *Presented at the Congress on Numerical Methods in Combinatorial Optimization, Capri, Italy*. presented at the Congress on Numerical Methods in Combinatorial Optimization, Capri, Italy.
- Holmes, R. and R. Parker (1976). A vehicle scheduling procedure based upon savings and a solution perturbation scheme. *Operational Research Quarterly* 27(1), 83–92.
- Janssens, G. and A. Van Breedam (1994). A simulated annealing post-processor for the vehicle routing problem. In C. Reeves (Ed.), *Applications of Modern Heuristics*. Alfred Waller Ltd.
- Johnson, D., C. Aragon, L. McGeoch, and C. Schevon (1989). Optimization by simulated annealing: an experimental evaluation; part i: Graph partitioning. *Operations Research* 37, 865–892.
- Kaufman, L. and P. Rousseeuw (1989). *Finding groups in data. An introduction to Cluster Analysis*. Wiley, New York.
- Kirckpatrick, S., C. Gelatt Jr., and M. Vecchi (1983). Optimization by simulated annealing. *Science* 220, 671–680.
- Knowles, K. (1967). The use of a heuristic tree-search algorithm for the vehicle routing and scheduling. In *Operational Research Conference, Exeter, England*.
- Koskodis, Y., W. Powell, and M. Solomon (1992). An optimization-based heuristic for vehicle routing and scheduling with soft time-windows. *Transportation Science* 26, 69–85.

- Laporte, G. and Y. Nobert (1987). Exact algorithms for the vehicle routing problem. *Annals of Discrete Mathematics* 31, 147–184.
- Lenstra, J. and A. Rinnooy Kan (1981). Complexity of vehicle routing and scheduling problems. *Networks* 11, 221–227.
- Lin, S. (1965). Computer solutions to the traveling salesman problem. *Bell System Technical Journal* 44, 2245–2269.
- Lin, S. and B. Kernighan (1973). An effective heuristic algorithm for the traveling salesman problem. *Operations Research* 21, 498–516.
- Martello, S. and P. Toth (1981). An algorithm for the generalized assignment problem. In J. Brans (Ed.), *Operations Research '81*, pp. 589–603. North-Holland, Amsterdam.
- McDonald, J. (1972). Vehicle scheduling - a case study. *Operational Research Quarterly* 23(4), 433–444.
- Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 21, 1087–1092.
- Mole, R. (1979). A survey of local delivery vehicle routing methodology. *Journal of the Operational Research Society* 30, 245–252.
- Mole, R. and S. Jameson (1976). A sequential route-building algorithm employing a generalised savings criterion. *Operational Research Quarterly* 27, 503–511.
- Morgan, J. and J. Sonquist (1963). Problems in the analysis of survey data and a proposal. *Journal of the American Statistical Association* 58, 415–434.
- Nelson, M., K. Nygard, J. Griffin, and W. Shreve (1988). Implementation techniques for the vehicle routing problem. *Computers and Operations Research* 12(3), 273–283.
- Nygard, K., P. Greenberg, W. Bolkan, and E. Swenson (1988). Generalized assignment methods for the deadline vehicle routing problem. In B. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 107–126. North-Holland, Amsterdam.
- Osman, I. (1991). *Metastrategy simulated annealing and tabu search algorithms for combinatorial optimization problems*. Ph. D. thesis, The Management School, University of London.
- Osman, I. (1993). Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research* 41, 421–451.
- Paessens, I. (1988). The savings algorithm for the vehicle routing problem. *European Journal of Operational Research* 34, 336–344.

- Pirlot, M. (1992). General local search heuristics in combinatorial optimization: a tutorial. *Jorbel* 32(2), 7–67.
- Potvin, J.-Y., T. Kervahut, B. Garcia, and J.-M. Rousseau (1992). A tabu search heuristic for the vehicle routing problem with time windows. Technical report, Technical Report, CRT-855, Centre de Recherche sur les Transports, Université de Montréal, Montréal, Canada.
- Potvin, J.-Y. and J.-M. Rousseau (1993). A parallel route building algorithm for the vehicle routing and scheduling problem with time windows. *European Journal of Operational Research* 66, 331–340.
- Pureza, V. and P. França (1991). Vehicle routing problems via tabu search metaheuristic. Technical report, Technical Report, CRT-747, Centre de Recherche sur les Transports, Université de Montréal, Montréal, Canada.
- Robusté, F., C. Daganzo, and R. Souleyrette (1990). Implementing vehicle routing models. *Transportation Research B* 24, 263–286.
- Rosenkrantz, D., R. Sterns, and P. Lewis (1977). An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing* 6, 571–581.
- Savelsbergh, M. (1988). *Computer Aided Routing*. Ph. D. thesis, Centrum voor Wiskunde en Informatica, Amsterdam.
- Savelsbergh, M. (1990a). An efficient implementation of local search algorithms for constrained routing problems. *European Journal of Operational Research* 47, 75–85.
- Savelsbergh, M. (1990b). A parallel insertion heuristic for vehicle routing with side-constraints. *Statistica Neerlandica* 44(3), 139–148.
- Semet, F. and E. Taillard (1993). Solving real-life vehicle routing problems efficiently using taboo search. *Annals of Operations Research* 41, 469–488.
- Siegel, S. and N. Castellan (1989). *Nonparametric statistics*. McGraw-Hill, Singapore.
- Solomon, M. (1987). Algorithms for the vehicle routing and scheduling problems with time-window constraints. *Operations Research* 35(2), 254–265.
- Solomon, M., E. Baker, and J. Schaffer (1988). Vehicle routing and scheduling with time-window constraints: efficient implementations of solution improvement procedures. In B. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 85–105. North-Holland, Amsterdam.
- Sonquist, J., E. Baker, and J. Morgan (1971). *Searching for structure*. Ann Arbor, Survey Research Center, University of Michigan.
- Stewart, W., J. Kelly, and M. Laguna (1992). Solving vehicle routing problems using generalized assignments and tabu search. Technical report, College of Business Administration, University of Colorado, Boulder.

- Taillard, E. (1991). Robust taboo search for the quadratic assignment problem. *Parallel Computing* 17, 433–455.
- Taillard, E. (1992). Parallel iterative search methods for vehicle routing problems. Technical report, Working paper ORWP 92/03, Département de Mathématiques, Ecole Polytechnique Fédérale de Lausanne.
- Teodorovic, D. and G. Pavkovic (1992). A simulated annealing technique approach to the vehicle routing problem in the case of stochastic demand. *Transportation Planning and Technology* 16, 261–273.
- Thangiah, S., I. Osman, and T. Sun (1994). Hybrid genetic algorithm, simulated annealing and tabu search methods for vehicle routing problem with time windows. Technical report, Technical Report SRU-CpSc-TR-94-27, Computer Science Department, Slippery Rock University, Slippery Rock, PA, U.S.A.
- Thangiah, S., I. Osman, R. Vinayagamoorthy, and T. Sun (1994). Algorithms for vehicle routing problems with time deadlines. *American Journal of Mathematical and Management Sciences* 13(3-4).
- Thangiah, S., T. Sun, and J.-Y. Potvin (1994). Heuristic approaches to vehicle routing with backhauls and time windows. Technical report, Technical Report SRU-CpSc-TR-94-28, Computer Science Department, Slippery Rock University, Slippery Rock, PA, U.S.A.
- Tillman, F. and H. Cochran (1968). A heuristic approach for solving the delivery problem. *Journal of Industrial Engineering* 19, 354–358.
- Tyagi, M. (1968). A practical method for the truck dispatching problem. *Journal of the Operations Research Society of Japan* 10, 76–92.
- Van Breedam, A. (1995). Improvement heuristics for the vehicle routing problem based on simulated annealing. *European Journal of Operational Research* 86, 480–490.
- Van Laarhoven, P. and E. Aarts (1987). *Simulated Annealing: Theory and Practice*. Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Van Landeghem, H. (1988). A bi-criteria heuristic for the vehicle routing problem with time-windows. *European Journal of Operational Research* 36, 217–226.
- Webb, M. (1972). Relative performance of some sequential methods of planning multiple delivery journeys. *Operational Research Quarterly* 23(2), 361–372.
- Wren, A. and A. Holliday (1972). Computer scheduling of vehicles from one or more depots to a number of delivery points. *Operational Research Quarterly* 23, 333–344.
- Yellow, P. (1970). A computational modification to the savings method of vehicle scheduling. *Operational Research Quarterly* 21(2), 281–283.

Ziegler, H.-J., G. Binder, H.-V. Niemeier, and H. Schrenk (1988). *Computer-gestützte Transport- und Tourenplanung*. Expert Verlag, Ehningen.

Index

- AID, 11, 141
- basic test set, 10, 134
 - criteria, 10
 - depot locations, 10, 134
 - grouping patterns, 10, 134
 - spreading patterns, 10, 134
- experimental plan, 10
 - analyses, 11
 - objectives, 10
 - treatments, 10
 - units, 10
- Friedman test, 12, 142
- Generalized Assignment heuristic, 25
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 41, 61, 71, 86
- generalized assignment problem, 24, 25
- Generalized Savings heuristic, 20
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 36, 58, 68, 82
- global optimisation heuristic, 94
- heuristic analysis, 12
 - improvement heuristics, 115
 - test set G1, 45
 - test set G2, 47
 - test set G3, 48
 - test set P1, 62
 - test set P2, 73
 - test set T1, 89, 90
- improvement heuristic
 - within-route, 141
- improvement heuristics
 - between-routes, 93
 - conclusions, 131
 - within-route, 93
- initial heuristics, 14
 - conclusions, 128
 - route-building heuristics, 14
 - parallel, 14
 - sequential, 14
 - two-phase heuristics, 14
- Local Improvement heuristic, 97
 - heuristic analysis, 115
 - parametric analysis, 109
- local optimisation heuristic, 93
- metaheuristics, 94
- mixed pick-up and deliveries, 53
- Parallel Assignment-based Insertion heuristic, 24
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 40, 60, 70, 85
- Parallel Insertion heuristic, 22
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 38, 59, 69, 84

- Parallel Nearest Neighbor heuristic,
 - 16
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 33, 55, 66, 79
- Parallel Savings heuristic, 18
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 35, 57, 67, 81
- parameter transformation procedure, 11
- parametric analysis, 11, 109
 - test set G1, 31
 - test set G2, 31
 - test set G3, 31
 - test set P1, 54
 - test set P2, 64
 - test set T1, 77
 - test set T2, 77
- post-processor, 141

- reduced test set, 108

- search strategy, 93
- seed generation, 138
 - Circle covering, 139
 - Cone covering, 139
- Sequential Insertion heuristic
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 37, 58, 68, 82
- Sequential Nearest Neighbor heuristic,
 - 15
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 32, 54, 65, 77
- Sequential Savings heuristic, 17
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 33, 56, 66, 80
- Simulated Annealing, 98
 - Boltzmann distribution, 99
 - cooling function, 99
 - epoch, 100
 - generic parameters, 99
 - ground state, 99
 - heuristic analysis, 115
 - initial temperature, 99
 - parametric analysis, 111
 - range delimiter, 100
 - stopping criterion, 100
 - thermal equilibrium, 99, 100
- Sweep heuristic, 28
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91

- Tabu Search, 102
 - aspiration criterion, 103
 - heuristic analysis, 115
 - long-term memory, 103, 104
 - parametric analysis, 114
 - stopping criterion, 103, 104
 - tabu list, 103
- test set G1, 30
- test set G2, 30
- test set G3, 30
- test set P1, 53
- test set P2, 54
- test set T1, 76
- test set T2, 76
- Traveling Salesman Problem, 140
 - 3-opt, 15, 141
 - heuristics, 140
- Two Phase heuristic, 26
 - heuristic analysis, 46, 48, 50, 63, 74, 90, 91
 - parametric analysis, 42, 61, 72, 87
- types of move, 94
 - String Cross, 94
 - String Exchange, 95
 - String Mix, 95

String Relocation, 95

Vehicle Routing Problem, 8
 assumptions, 8
 definition, 8
 mathematical programming for-
 mulation, 9